The Geometry of Sameness: An ε-Equivalence of Translation and Distance

Categories of Realizations, Error-Budget Transfer, and a Hidden Variable

Bee Rosa Davis

NASA Mission Systems Engineer

bee_davis@alumni.brown.edu

Abstract

At an intuitive level, the problems we study are like trying to understand the Earth from maps. The underlying *semantic sameness structure* S is the globe: a hidden world of entities and their relationships. One engineering tradition starts from the globe (a manifold): learn a space where geodesic distance reflects semantic change and read decisions off geodesic gaps—this is the geometry-first view used by systems like HERALD and VIDAR. A second tradition starts from overlapping flat maps: heterogeneous observation spaces with translators between them, stitched together just enough to reconcile telemetry or transactions—this is the translation-first view used by systems like AMC and PRISM. Our central claim is that, whenever the "maps" are smooth and consistent enough to be sewn into a globe, these are not competing descriptions but dual realizations of the same S: tools and guarantees can be moved across the seam.

Formally, for a fixed semantic sameness structure S we define two categories of realizations: $\mathsf{SamTrans}(S)$, whose objects are *translation-based* systems on heterogeneous observation spaces with bounded translator drift, and $\mathsf{SamGeom}(S)$, whose objects are *manifold-based* systems with Riemannian metrics, path families, and configuration margins. We introduce *ε-equivalence of categories* tailored to detection, together with explicit functors $F_S : \mathsf{SamTrans}^0(S) \to \mathsf{SamGeom}^0(S)$ and $G_S : \mathsf{SamGeom}^0(S) \to \mathsf{SamTrans}^0(S)$ that act like "stitching the maps into a globe" and "unfolding the globe back into maps" on well-behaved subcategories $\mathsf{SamTrans}^0(S), \mathsf{SamGeom}^0(S)$ that satisfy a *smooth chartability* condition. Smooth chartability is the critical, falsifiable hypothesis of the framework: it is exactly the requirement that the local translators around S can be approximated by a compatible atlas whose transition maps behave like local diffeomorphisms, so that a globe M actually exists.

Our main result is an *Error Budget Transfer Theorem.* Each realization (translator or manifold) comes with a four-term detection error budget $(E_{\text{geom}}, E_{\text{link}}, \xi, \zeta)$ for geometry, linkage, calibration, and abstention, together with an independence slack δ_{indep} and a finite-variance Cantelli term B_{Cantelli} linking statistic separation to posterior correctness. We show that for any $T \in \mathsf{SamTrans}^0(S)$ and its geometric realization $M = F_S(T)$,

$$E^{\text{TSP}}_{\text{tot}}(T) \;=\; 1 - (1 - E^{\mathbf{T}}_{\text{geom}})(1 - E^{\mathbf{T}}_{\text{link}})(1 - \xi^{\mathbf{T}})(1 - \zeta^{\mathbf{T}}) + \delta^{\mathbf{T}}_{\text{indep}}$$

and

$$E^{\text{MSP}}_{\text{tot}}(M) \;=\; 1 - (1 - E^{\mathbf{M}}_{\text{geom}})(1 - E^{\mathbf{M}}_{\text{link}})(1 - \xi^{\mathbf{M}})(1 - \zeta^{\mathbf{M}}) + \delta^{\mathbf{M}}_{\text{indep}}$$

differ by at most an explicit, first-order slack $|E^{\text{TSP}}_{\text{tot}}(T) - E^{\text{MSP}}_{\text{tot}}(M)| \le C_F(\varepsilon_{\text{trans}}, \varepsilon_{\text{dist}}, \delta_{\text{chart}})$, and the associated Cantelli bounds on high-risk alerts satisfy an analogous relation $|B_{\text{Cantelli}}(T) -$

$B_{\text{Cantelli}}(M)| \le C'_F(\cdots)$. A symmetric statement holds for $M \in \mathsf{SamGeom}^0(S)$ and $T = G_S(M)$. In other words, once S passes the smooth chartability test, *translator-based and manifold-based realizations of S share the same detection guarantees up to controlled slack*: the choice of "flat maps vs. globe" becomes an engineering decision, not a conceptual one.

Practically, this ε-equivalence has two consequences. First, it lifts theorems across the bridge: guarantees proved in a geometry-first system (e.g., Davis-style path-based bounds for HERALD or VIDAR) can be pulled back to translation-first architectures like AMC via G_S, and vice versa. Second, it lifts tools: audits developed for translator graphs (e.g., AMC's spectral stability checks) become curvature or chart-overlap diagnostics on manifolds via F_S. We instantiate this framework on four systems—AMC and PRISM as translation-first, HERALD and VIDAR as geometry-first—to illustrate how chartability can be audited empirically and how error budgets and techniques transfer in practice.

1 Introduction

Many modern systems need to decide whether two heterogeneous observations refer to the same underlying entity or event: is this telemetry packet semantically equivalent to that one after compression and format switching? Does this payment record match that ledger entry across rails? Do these viral variants represent the same antigenic state? Does this video clip belong to the same physical person as that enrollment image?

Two families of architectures answer these questions in different languages:

- **Translator-based systems** (e.g., AMC-style adaptive format pipelines, multi-rail reconciliation) build a graph of learned translators between embedding formats or modalities and assert sameness when translated representations agree within a drift budget.
- **Manifold-based systems** (e.g., Davis-style HERALD and VIDAR) learn or inherit a Riemannian manifold on which geodesic distance along identity-preserving paths encodes semantic change, and assert sameness when points remain within a configuration margin in that geometry.

Informally, one view speaks in *translations* and drift; the other in *distances* and curvature. Both solve the same kind of problem, but their guarantees, diagnostics, and engineering tools are typically developed in isolation.

This work starts from a simple question:

> *For a fixed notion of semantic sameness, are translator-based and manifold-based detection systems two views of the same underlying problem? If so, can guarantees and tools proven in one view be transported to the other without re-deriving them from scratch?*

We answer this in the affirmative under explicit regularity assumptions. The key move is to make the *problem* itself explicit.

1.1 Semantic sameness structures and realizations

We model the underlying task as a *semantic sameness structure S*: latent identities, observation spaces, rendering maps, and a sameness relation on observations. For a fixed S, engineers can build many concrete systems.

- $\mathsf{SamTrans}(S)$: a category whose objects are *translator-based realizations* of S (observation encoders, format/rail translators, drift budgets, and detection logic in translator space), and whose morphisms are reparameterizations preserving translators up to controlled error.
- $\mathsf{SamGeom}(S)$: a category whose objects are *manifold-based realizations* of S (Riemannian manifolds, chart atlases, path families, and Davis-style detection logic), and whose morphisms are quasi-isometries respecting configuration structure and charts.

Crucially, these are *two* design spaces for the same S: translator systems and manifold systems are not being unified with each other directly, but as two families of realizations of the same hidden variable.

On well-behaved subcategories $\mathsf{SamTrans}^0(S) \subset \mathsf{SamTrans}(S)$ and $\mathsf{SamGeom}^0(S) \subset \mathsf{SamGeom}(S)$, defined by smooth chartability, bounded distortion, and non-vacuous configuration margins, we construct functors

$$F_S : \mathsf{SamTrans}^0(S) \to \mathsf{SamGeom}^0(S), \qquad G_S : \mathsf{SamGeom}^0(S) \to \mathsf{SamTrans}^0(S),$$

that turn translators into manifolds and manifolds into translator graphs. These functors are not magic; they succeed precisely when the chartability and distortion assumptions are satisfied.

1.2 Main contributions

Our contributions are organized around three load-bearing claims and their practical implications.

(0) Language: ε-equivalence of realization categories. We formalize translator and manifold realizations of a fixed semantic sameness structure S as categories $\mathsf{SamTrans}(S)$ and $\mathsf{SamGeom}(S)$ and introduce an *ε-equivalence* notion tailored to detection. Rather than claiming a strict equivalence of categories, we construct functors F_S and G_S and show that $G_S \circ F_S \simeq \mathrm{Id}_{\mathsf{SamTrans}^0(S)}$ and $F_S \circ G_S \simeq \mathrm{Id}_{\mathsf{SamGeom}^0(S)}$ *up to quasi-isometry and bounded morphism error*, with all slack terms explicit. Categorical language is not decorative here: it is the minimal formalism needed to express how *error budgets* transform under representation change.

(1) Error Budget Transfer Theorem (main result). Our central technical result shows that detection guarantees transfer across the two realizations. Conceptually, we prove that the correctness bound for a translator system (T) is ε-equivalent to the correctness bound for the manifold system (M) that realizes the same problem S:

$$\mathrm{LB}(T) \cong_\varepsilon \mathrm{LB}(F_S(T))$$

This main result, the Error Budget Transfer Theorem, is stated formally below. It shows that the full, multi-part error budgets for both systems are equivalent up to explicit, first-order slack terms.

Error Budget Transfer Theorem (informal preview). *Fix a semantic sameness structure S and consider well-behaved realizations $T \in \mathsf{SamTrans}^0(S)$ and $M \in \mathsf{SamGeom}^0(S)$.*

Translator-side bound. The translator realization T admits a Davis-style correctness lower bound

$$\mathrm{LB}_{\mathrm{T}}(T) = B_{\mathrm{Cantelli}}(T)\left[(1-E^{\mathrm{T}}_{\mathrm{geom}})(1-E^{\mathrm{T}}_{\mathrm{link}})(1-\xi^{\mathrm{T}})(1-\zeta^{\mathrm{T}}) - \delta^{\mathrm{T}}_{\mathrm{indep}}\right]_{+} - \varepsilon^{\mathrm{T}}_{\mathrm{est}}.$$

Manifold-side bound. The manifold realization M admits an analogous bound

$$\mathrm{LB}_{\mathrm{M}}(M) = B_{\mathrm{Cantelli}}(M)\left[(1-E^{\mathrm{M}}_{\mathrm{geom}})(1-E^{\mathrm{M}}_{\mathrm{link}})(1-\xi^{\mathrm{M}})(1-\zeta^{\mathrm{M}}) - \delta^{\mathrm{M}}_{\mathrm{indep}}\right]_{+} - \varepsilon^{\mathrm{M}}_{\mathrm{est}}.$$

Transfer. The functors F_S and G_S transport these guarantees:

$$\left|\mathrm{LB}_{\mathrm{T}}(T) - \mathrm{LB}_{\mathrm{M}}(F_S(T))\right| \;\leq\; C_F(\varepsilon_{\mathrm{trans}}, \varepsilon_{\mathrm{dist}}, \delta_{\mathrm{chart}}),$$

$$\left|\mathrm{LB}_{\mathrm{M}}(M) - \mathrm{LB}_{\mathrm{T}}(G_S(M))\right| \;\leq\; C_G(\varepsilon_{\mathrm{trans}}, \varepsilon_{\mathrm{dist}}, \delta_{\mathrm{chart}}),$$

where C_F and C_G are explicit first-order slack functions.

(2) Characterization via smooth chartability. The functor $F_S : \mathsf{SamTrans}^0(S) \to \mathsf{SamGeom}^0(S)$ is not automatic; it succeeds only on translator systems with *smooth chartability.* We isolate this as the central, falsifiable hypothesis of the framework: locally linear approximants to each translator must exist with bounded Jacobian and small residual drift on in-regime data. This assumption is testable via concrete audit protocols (e.g., Jacobian-based linearization error histograms on AMC-style translators) and is empirically supported in systems whose architectures already regularize for smoothness and stability.

(3) Constructive functors and ε-equivalence. Under smooth chartability and bounded distortion, we give explicit constructions of F_S (path-metric completion of a translator graph followed by a Riemannian smoothing) and G_S (chart-transition translators induced by overlapping coordinate maps). We show that on $\mathsf{SamTrans}^0(S)$ and $\mathsf{SamGeom}^0(S)$ these functors form an ε-equivalence of categories: they are quasi-inverse up to explicitly bounded drift, curvature, and chart-overlap slack. This formalizes the intuition that translator and manifold views are dual coordinatizations of the same semantic sameness problem S.

Implications: tool and theorem transfer. The practical consequence of this ε-equivalence is that *tools and theorems* proven in one view can now be reused in the other.

- *Theorems transfer.* A guarantee proven in a manifold system (e.g., HERALD's Davis-style dominance lower bound, VIDAR's deepfake correctness bounds) can be pulled back via G_S to a translator-based realization of the same S (e.g., AMC-style telemetry translators), with only first-order changes in the error budget.
- *Tools transfer.* Engineering techniques developed for translator graphs (spectral stability audits, drift-triggered retraining, lineage-aware rollback) can be pushed forward via F_S to become curvature and chart-stability diagnostics on manifolds, providing new audit tools for geometry-first systems.

Engineers starting from a single semantic sameness structure S therefore have a principled choice of working in translator or manifold coordinates, knowing that both design spaces are linked by functors that preserve detection guarantees up to explicitly quantified slack.

1.3 Running examples

We ground the theory in four concrete systems, which we treat as realizations of different semantic sameness structures S.

AMC (telemetry format switching). AMC [?] is a telemetry-native embedding architecture for adaptive format switching and semantic translation in high-volume telemetry pipelines. It maintains a directed graph of learned format translators constrained by an ε-bounded semantic drift loss and uses spectral properties of a format-compatibility graph to select stable routes and fallbacks. In our language, AMC is a canonical object of $\mathsf{SamTrans}(S_{\text{tele}})$ for a telemetry sameness structure S_{tele}.

PRISM (payment reconciliation). PRISM [?] performs multi-rail payment reconciliation with learned translators between heterogeneous transaction formats, semantic alignment constraints, and a central latent reconciliation space. It naturally lives in $\mathsf{SamTrans}(S_{\text{pay}})$ but exposes an implicit manifold structure over reconciled transaction identities.

HERALD (viral antigenic drift). HERALD [?], introduced in prior work on Davis manifolds [?], builds a pullback Riemannian manifold on viral sequences where geodesic distance along epitope-restricted mutation paths approximates antigenic distance; a PIT-fused drift statistic and Cantelli bounds map geometric drift to dominance risk with an explicit error budget. HERALD is a prototypical object of $\mathsf{SamGeom}(S_{\text{antigen}})$.

VIDAR (deepfake detection via identity trajectories). VIDAR treats face-recognition embeddings as points on the hypersphere, models short clips as identity trajectories on an identity manifold, and builds RIM features plus auxiliary detectors as inputs to a Davis-style scalar statistic and bound.[1] VIDAR is a canonical object of $\mathsf{SamGeom}(S_{\text{id}})$.

Across these examples, the same pattern recurs: translator systems like AMC and PRISM live in $\mathsf{SamTrans}(S)$; geometry-first systems like HERALD and VIDAR live in $\mathsf{SamGeom}(S)$; and our functors F_S and G_S explain how to move guarantees and tools between these views.

1.4 Scope and non-goals

This manuscript is deliberately theoretical. It:

- defines semantic sameness structures S and the categories $\mathsf{SamTrans}(S)$ and $\mathsf{SamGeom}(S)$;
- introduces smooth chartability and bounded-distortion path geometry as explicit, falsifiable assumptions defining subcategories $\mathsf{SamTrans}^0(S)$ and $\mathsf{SamGeom}^0(S)$;
- constructs F_S and G_S and proves an ε-equivalence of realization categories;

[1] See Davis, "The Davis Manifold: Geometry-First Detection with Compositional Error Budgets," for the VIDAR instantiation.

- proves the Error Budget Transfer Theorem (Theorem 6.3) showing that Davis-style correctness bounds transfer across these views with first-order slack; and
- sketches audit protocols for checking smooth chartability and error-budget components in systems like AMC, PRISM, HERALD, and VIDAR.

We do *not* report new empirical performance, deployable calibration curves, or production latencies; those belong in system-specific papers. Our goal here is to expose the common hidden variable S, make the representation change between translator and manifold views explicit and functorial, and characterize when and how detection guarantees can be transported between them.

1.5 Roadmap

The paper develops in three acts.

Act I (Framework). Section 2 formalizes semantic sameness structures S and defines translator and manifold realizations, together with Davis-style error budgets and the notion of smooth chartability. Section 3 introduces the categories $\mathsf{SamTrans}(S)$ and $\mathsf{SamGeom}(S)$, constructs the functors F_S and G_S at a high level, and defines ε-equivalence of realization categories.

Act II (Constructions and transfer). Section 4 builds F_S concretely by endowing translator graphs with a path metric and a smoothed Riemannian structure under smooth chartability, and analyzes the resulting distortion and chart-overlap slack. Section 5 constructs G_S from manifold charts and transition maps and analyzes its stability. Section 6 proves the Error Budget Transfer Theorem (Theorem 6.3) by tracking how geometric, linkage, calibration, and abstention errors transform under F_S and G_S.

Act III (Case studies and outlook). Section 7 instantiates the framework on AMC, PRISM, HERALD, and VIDAR, including concrete chartability audits and error-budget estimates. Section 8 discusses limitations, open problems, and the broader “Davis universality” conjecture for geometry-first detection in temporal semantic sameness problems.

2 Semantic sameness structures and realizations

Throughout this section we fix a *semantic sameness problem* and make precise what it means to “realize” that problem either as a translator graph or as a Riemannian manifold. The hidden object is the sameness structure itself; translation-based and manifold-based systems are two different realizations of the same underlying task.

We use boldface symbols such as $\mathbf{T}$ and $\mathbf{M}$ for concrete systems (a particular TSR or MSR) and reserve plain letters for abstract spaces and maps.

2.1 Semantic sameness structures

Intuitively, a semantic sameness structure records: (i) a latent space of entities u (viruses, identities, vehicles, . . .), (ii) how each modality i renders those entities as observations x_i, (iii) when two observations in possibly different modalities “are the same entity”, and (iv) which latent trajectories count as benign identity-preserving evolution.

Definition 1 (Semantic sameness structure)**.** *A* semantic sameness structure *is a tuple*

$$S := (I,\ \mathcal{I},\ \{X_i\}_{i\in\mathcal{I}},\ \{\pi_i\}_{i\in\mathcal{I}},\ \approx,\ \{\mathcal{P}_S(L)\}_{L>0}),$$

where:

(i) *I is a (typically high-dimensional) set of* latent entities *or states. A point $u \in I$ represents the underlying "object of interest" (e.g., true antigenic state, true human identity, true vehicle configuration).*

(ii) *$\mathcal{I}$ is a finite index set of* modalities *(e.g., file formats, sensors, assays, views).*

(iii) *For each $i \in \mathcal{I}$, X_i is an* observation space *for modality i. We write*

$$X := \bigsqcup_{i\in\mathcal{I}} X_i$$

for the disjoint union of all observations.

(iv) *For each $i \in \mathcal{I}$, $\pi_i : I \to X_i$ is a (possibly partial)* rendering map *that produces an ideal observation of u in modality i when defined.*[2] *When $\pi_i(u)$ is undefined, modality i simply does not observe u.*

(v) *$\approx$ is a* semantic sameness relation *on X, defined by*

$$x \approx x' \quad\Longleftrightarrow\quad \exists u \in I,\ \exists i,j \in \mathcal{I} \text{ such that } x \text{ is an observation "of" } u \text{ in } X_i,\ x' \text{ is an observation "of" } u \text{ in}$$

In the noiseless idealization this reduces to: $x \approx x'$ if there exist $u \in I$ and (i,j) with $x = \pi_i(u)$ and $x' = \pi_j(u)$. In practice we blur small observation noise and treat x as an observation "of" u whenever it lies in a pre-specified neighborhood of $\pi_i(u)$.

(vi) *For each $L > 0$, $\mathcal{P}_S(L)$ is a family of* benign latent paths *$\gamma : [0,1] \to I$ indexed by a* semantic horizon *L. Each $\gamma \in \mathcal{P}_S(L)$ represents an identity-preserving evolution of latent state (e.g., realistic antigenic drift over one season, smooth identity motion over a short clip). We require that $L \mapsto \mathcal{P}_S(L)$ is monotone: if $L_1 \le L_2$ then $\mathcal{P}_S(L_1) \subseteq \mathcal{P}_S(L_2)$.*

Examples (semantic structures).

- **Format-heterogeneous telemetry (AMC [?]).** I is the space of underlying system states; each X_i is a particular log or telemetry format; π_i renders a state into that format; $\gamma \in \mathcal{P}_S(L)$ is a short-time state trajectory.
- **Antigenic drift (HERALD).** I are viral lineages; X_{seq} are amino-acid sequences; X_{neut} are neutralization measurements; π_{seq} encodes the canonical genome for a lineage; π_{neut} renders idealized assay readouts; $\mathcal{P}_S(L)$ are mutation paths with at most L epitope substitutions.
- **Identity trajectories (VIDAR).** I are person identities; X_{video} are short talking-head clips; X_{audio} are speech segments; each π_i renders a clip of the same person; $\mathcal{P}_S(L)$ are natural motion trajectories over windows of bounded duration.

[2]Formally, one may model π_i as a Markov kernel from I to X_i, capturing stochastic observation noise; for notational simplicity we write π_i as a map and treat noise separately in the detector error budget.

We will later view *realizations* of a fixed S as objects of two categories: a translation-based category $\mathsf{SamTrans}(S)$ and a geometry-based category $\mathsf{SamGeom}(S)$.

2.2 Translation-based realizations (TSRs)

A translation-based realization starts from modality-specific feature spaces and implements sameness comparisons via learned translators between those spaces.

Definition 2 (Translation-based semantic realization (TSR))**.** *Let S be a semantic sameness structure as in Definition* **??***. A* translation-based semantic realization *(TSR) of S is a tuple*

$$\mathbf{T} := (\{V_i\}_{i\in\mathcal{I}}, \{\varphi_i\}_{i\in\mathcal{I}}, G_T, \{T_{ij}\}_{(i,j)\in E_T}, \varepsilon_{\text{trans}}(\cdot)),$$

where:

(i) *For each $i \in \mathcal{I}$, V_i is a finite-dimensional inner-product space (feature space for modality i).*

(ii) *$\varphi_i : X_i \to V_i$ is a feature extractor (e.g., a neural encoder). We denote ideal features by*

$$v_i(t) := \varphi_i(\pi_i(\gamma(t))) \qquad \text{for } \gamma \in \mathcal{P}_S(L),\ t \in [0,1],$$

whenever $\pi_i(\gamma(t))$ is defined.

(iii) *$G_T = (\mathcal{I}, E_T)$ is a directed graph whose vertices are modalities and whose edges $i \to j$ indicate the availability of a learned* translator *$T_{ij} : V_i \to V_j$.*

(iv) *For any path $p = (i_0, \ldots, i_k)$ in G_T we write*

$$T_p := T_{i_{k-1}i_k} \circ \cdots \circ T_{i_0 i_1} \quad \text{and} \quad |p| := k.$$

(v) *$\varepsilon_{\text{trans}} : \mathbb{N} \to [0,\infty)$ is a* translator drift profile *such that, for any benign latent path $\gamma \in \mathcal{P}_S(L)$, any time $t \in [0,1]$, any modalities i, j with both $\pi_i(\gamma(t))$ and $\pi_j(\gamma(t))$ defined, and any G_T-path $p : i \to j$ of length k, we have*

$$\|T_p(v_i(t)) - v_j(t)\| \ \le\ \varepsilon_{\text{trans}}(k). \tag{1}$$

This inequality measures representation error *of the translators at a single latent time, decoupled from any semantic drift along γ.*

We write TSR(S) *for the collection of all TSRs that realize S.*

Examples (TSRs).

- **AMC-style telemetry [?].** V_i are format-specific embedding spaces; φ_i encode logs into these spaces; T_{ij} are learned format translators; $\varepsilon_{\text{trans}}$ bounds accumulated translator drift over paths in the format graph.
- **PRISM-style multi-signal systems [?].** V_i are detector-specific representations; T_{ij} fuse or reconcile outputs between detectors; $\varepsilon_{\text{trans}}$ captures how far translations can drift while still preserving semantic identity.

2.3 Manifold-based realizations (MSRs)

A manifold-based realization represents semantic sameness by embedding latent entities into a Riemannian manifold and requiring that feature-space coordinates act as approximate charts.

Definition 3 (Manifold-based semantic realization (MSR))**.** *Let S be as in Definition* **??***. A* manifold-based semantic realization *(MSR) of S is a tuple*

$$\mathbf{M} := (\{V_i\}_{i\in\mathcal{I}}, \{\varphi_i\}_{i\in\mathcal{I}}, (\mathcal{M}, g), \{\psi_i\}_{i\in\mathcal{I}}, \varepsilon_{\text{dist}}(\cdot)),$$

where:

(i) $\{V_i\}_{i\in\mathcal{I}}$ *and* $\{\varphi_i : X_i \to V_i\}_{i\in\mathcal{I}}$ *are feature spaces and encoders as in a TSR.*

(ii) $(\mathcal{M}, g)$ *is a d-dimensional smooth Riemannian manifold of* semantic states *with geodesic distance d_g.*

(iii) *For each $i \in \mathcal{I}$, $\psi_i : V_i \to \mathcal{M}$ is a smooth* parameterization *whose image $U_i := \psi_i(V_i)$ is an open subset of $\mathcal{M}$. We require that the inverse maps*

$$\chi_i := \psi_i^{-1} : U_i \to V_i$$

exist and are smooth; the family $\{(U_i, \chi_i)\}_{i\in\mathcal{I}}$ forms a (possibly partial) chart atlas on the portion of $\mathcal{M}$ we actually use.

(iv) *There exists a* latent state map $F : I \to \mathcal{M}$ *such that, whenever $\pi_i(u)$ is defined,*

$$F(u) = \psi_i(\varphi_i(\pi_i(u))).$$

This ensures that, on ideal data, all modalities agree on the same manifold point for a given latent entity.

(v) $\varepsilon_{\text{dist}} : (0, \infty) \to [0, \infty)$ *is a* metric distortion profile *such that, for any chart (U_i, χ_i), any two points $z_a, z_b \in U_i$ with $d_g(z_a, z_b) > 0$, and chart distance*

$$\delta_i(z_a, z_b) := \|\chi_i(z_a) - \chi_i(z_b)\|_{V_i},$$

we have

$$\left|\frac{\delta_i(z_a, z_b)}{d_g(z_a, z_b)} - 1\right| \;\le\; \varepsilon_{\text{dist}}(d_g(z_a, z_b)). \tag{2}$$

This $\varepsilon_{\text{dist}}$ measures representation error *of the charts χ_i relative to the underlying metric g. We write $\mathsf{MSR}(S)$ for the collection of all MSRs that realize S.*

Examples (MSRs).

- **HERALD-style geometry.** $(\mathcal{M}, g)$ is an antigenic manifold; V_i is the latent space of a sequence encoder; ψ_i is the learned map into $\mathcal{M}$; $\varepsilon_{\text{dist}}$ captures geodesic–Euclidean distortion along antigenic trajectories.
- **VIDAR-style identity manifold.** $(\mathcal{M}, g)$ is the hypersphere $\mathbb{S}^{d-1}$ with its standard metric; V_{video} is the pre-normalized ArcFace embedding space; ψ_{video} is ℓ_2-normalization; $\varepsilon_{\text{dist}}$ is small in

the small-angle regime.

2.4 Semantic drift profile

Once an MSR $\mathbf{M}$ is fixed, latent benign paths in $\mathcal{P}_S(L)$ induce paths on $\mathcal{M}$ via the latent map F.

Definition 4 (Semantic drift profile). *Let S be a semantic sameness structure and $\mathbf{M}$ an MSR of S as in Definition **??**. For $L > 0$ and any $\gamma \in \mathcal{P}_S(L)$ define the induced manifold path*

$$z_\gamma(t) := F\big(\gamma(t)\big) \in \mathcal{M}.$$

The semantic drift profile *of S under $\mathbf{M}$ is*

$$\sigma_S^{\mathbf{M}}(L) := \sup_{\gamma \in \mathcal{P}_S(L)} \sup_{0 \le s \le t \le 1} d_g\big(z_\gamma(s), z_\gamma(t)\big).$$

Thus $\sigma_S^{\mathbf{M}}(L)$ records, in purely geometric terms, how far benign latent paths are allowed to move in $\mathcal{M}$ over horizon L. It is a property of the pair $(S, \mathbf{M})$, independent of translators or observation noise.

2.5 Chart overlap slack and translator asymmetry

When a single S admits both a TSR $\mathbf{T}$ and an MSR $\mathbf{M}$, we will compare them via two structural diagnostics: a chart-overlap slack that measures how well translators agree with chart transitions, and a translator asymmetry that measures round-trip drift.

Definition 5 (Chart overlap slack). *Let S be fixed, $\mathbf{T}$ a TSR and $\mathbf{M}$ an MSR of S sharing the same feature spaces $\{V_i\}_{i \in \mathcal{I}}$ and encoders $\{\varphi_i\}_{i \in \mathcal{I}}$. For each i, let $U_i := \psi_i(V_i) \subseteq \mathcal{M}$ and $\chi_i := \psi_i^{-1} : U_i \to V_i$.*

For any overlap region $U_i \cap U_j \neq \varnothing$ and point $z \in U_i \cap U_j$, the ideal coordinate change from modality i to modality j maps

$$v_i := \chi_i(z) \quad \mapsto \quad v_j^\star := \chi_j(z).$$

The translator T_{ij} instead produces $T_{ij}(v_i)$. The chart overlap slack *is*

$$\delta_{\text{chart}}(\mathbf{T}, \mathbf{M}) := \sup_{(i,j) \in E_T} \sup_{z \in U_i \cap U_j} \big\| T_{ij}\big(\chi_i(z)\big) - \chi_j(z) \big\|.$$

This quantity measures how far translator-based coordinate changes deviate from the chart transitions implied by $\mathbf{M}$.

Definition 6 (Translator asymmetry). *Fix a TSR $\mathbf{T}$ with graph $G_T = (\mathcal{I}, E_T)$ and feature spaces $\{V_i\}$. Let $L_0 > 0$ be a chosen operating horizon and define, for each ordered pair (i, j) with both π_i and π_j defined along $\mathcal{P}_S(L_0)$,*

$$\Omega_{ij}(\mathbf{T}) := \Big\{ v_i(t) = \varphi_i\big(\pi_i(\gamma(t))\big) \;\Big|\; \gamma \in \mathcal{P}_S(L_0),\ t \in [0,1],\ \pi_j(\gamma(t)) \text{ defined} \Big\} \subseteq V_i.$$

Assume that G_T contains edges $i \to j$ and $j \to i$. The translator asymmetry *of* $\mathbf{T}$ *is*

$$\delta_{\text{asymm}}(\mathbf{T}) := \sup_{(i,j)\in E_T,\ (j,i)\in E_T} \ \sup_{v\in\Omega_{ij}(\mathbf{T})} \|T_{ji}(T_{ij}(v)) - v\|.$$

This is a structural parameter: it measures worst-case round-trip drift on in-regime features, not a probability. It will enter later as a diagnostic and, when needed, as a slack term in error-transfer bounds.

2.6 Sameness detectors and error budgets

Finally, we formalize detectors that operate on a fixed realization $\mathbf{R}$ of S and attach a decomposed error budget. Here $\mathbf{R}$ may be either a TSR $\mathbf{T}$ or an MSR $\mathbf{M}$.

Definition 7 (Sameness detector and error budget)**.** *Let S be a semantic sameness structure and $\mathbf{R}$ a realization of S (TSR or MSR). A* sameness detector *on $\mathbf{R}$ is a tuple*

$$\mathbf{D}_{\mathbf{R}} := (\Sigma_{\mathbf{R}},\ A_{\mathbf{R}},\ E^{\mathbf{R}}_{\text{geom}},\ E^{\mathbf{R}}_{\text{link}},\ \xi^{\mathbf{R}},\ \zeta^{\mathbf{R}},\ \delta^{\mathbf{R}}_{\text{indep}},\ \tau_{\text{vac}}),$$

with components:

(i) $\Sigma_{\mathbf{R}}$ *is a scalar* detection or similarity statistic *built from the representation induced by* $\mathbf{R}$*:*

- *for a TSR $\mathbf{T}$, $\Sigma_{\mathbf{T}}$ acts on feature-space objects (e.g., paths or pairs in $\bigsqcup_i V_i$);*
- *for an MSR $\mathbf{M}$, $\Sigma_{\mathbf{M}}$ acts on points or paths in $\mathcal{M}$.*

(ii) $A_{\mathbf{R}}$ *is a decision rule mapping ($\Sigma_{\mathbf{R}}$, context) to a finite action set such as* $\{\mathsf{same}, \mathsf{changed}, \mathsf{ambiguous}, \mathsf{abstain}\}$ *or* $\{\mathsf{info}, \mathsf{warn}, \mathsf{high\_risk}, \mathsf{insufficient\_signal}\}$*.*

(iii) $E^{\mathbf{R}}_{\text{geom}}$ *is the* geometric / representation error*: the probability (on in-regime, non-abstaining inputs) that the structural assumptions of $\mathbf{R}$ fail on ideal data:*

- *for a TSR $\mathbf{T}$, $E^{\mathbf{T}}_{\text{geom}}$ is the probability that the drift bound (**??**) is violated for some benign path $\gamma \in \mathcal{P}_S(L_0)$, time t, and translator path p of length at most the operating horizon;*
- *for an MSR $\mathbf{M}$, $E^{\mathbf{M}}_{\text{geom}}$ is the probability that the distortion bound (**??**) fails on some pair of induced manifold points $z_\gamma(s), z_\gamma(t)$ coming from benign paths, or that curvature / chart regularity assumptions used to define $\mathbf{M}$ are violated.*

By construction, $E^{\mathbf{R}}_{\text{geom}}$ depends only on S and the ideal maps (φ_i, T_{ij}) or (φ_i, ψ_i, g), not on observation noise.

(iv) $E^{\mathbf{R}}_{\text{link}}$ *is the* linkage / observation error*: the probability that the detection pipeline fails* even when geometry is ideal*. It captures the effect of observation noise and feature misspecification:*

- *for a TSR, this includes deviations between noisy inputs $\tilde{x}_i$ and ideal renders $\pi_i(u)$ that cause $\varphi_i(\tilde{x}_i)$ to move far enough from $\varphi_i(\pi_i(u))$ to flip $\Sigma_{\mathbf{T}}$ across a decision boundary;*
- *for an MSR, it includes analogous noise-induced shifts of manifold points $z \in \mathcal{M}$ and any failure of $\Sigma_{\mathbf{M}}$ to track d_g as intended.*

(v) $\xi^{\mathbf{R}}$ *is a* calibration error *term quantifying how far the calibrated mapping from $\Sigma_{\mathbf{R}}$ to probabilities departs from the true conditional $P(Y = 1 \mid \Sigma_{\mathbf{R}})$ in the high-risk region.*

(vi) $\zeta^{\mathbf{R}}$ *is an* abstention failure *term: the probability that the system fails to abstain when it should (e.g., when distortion audits, feature-range checks, or support conditions indicate that* $\mathbf{R}$ *is being used outside its validated regime).*

(vii) $\delta^{\mathbf{R}}_{\text{indep}} \in [0,1]$ *is an* independence slack *that upper-bounds the gap between the joint failure probability of the four error sources (geometry, linkage, calibration, abstention) and the product of their marginals. It quantifies deviations from an idealized independence model.*

(viii) $\tau_{\text{vac}} \in (0,1)$ *is a* vacuity threshold*. When the combined error budget (e.g., a sum or multiplicative form in later theorems) exceeds* τ_{vac}*, we declare the resulting correctness bound non-informative and treat the detector as operating outside its validated regime.*

Three-layer view. Definitions **??**–**??** separate three distinct contributors to detection performance:

- *Semantic drift* $\sigma^{\mathbf{M}}_S(L)$ (Definition **??**) is how far latent entities move in geometry under benign evolution. It is a property of $(S, \mathbf{M})$ and the path family.
- *Representation error* $\varepsilon_{\text{trans}}$ and $\varepsilon_{\text{dist}}$ (Definitions **??**–**??**) measure how imperfectly translators and charts realize the intended semantics on ideal data. These feed into $E^{\mathbf{R}}_{\text{geom}}$.
- *Observation noise and linkage error* are captured by $E^{\mathbf{R}}_{\text{link}}$: even if geometry were perfect, noisy or shifted observations can still flip decisions.

This three-layer decomposition is the backbone of the later Error Budget Transfer Theorem: it lets us move guarantees between $\mathsf{SamTrans}(S)$ and $\mathsf{SamGeom}(S)$ while keeping semantic drift, representation quality, and observation noise conceptually distinct.

3 Realization categories and ε-equivalence

Section **??** introduced semantic sameness structures S and two ways of realizing them in concrete systems: translation-based realizations (TSRs, Definition **??**) and manifold-based realizations (MSRs, Definition **??**). In this section we package these realizations into two categories,

$$\mathsf{SamTrans}(S) \quad \text{and} \quad \mathsf{SamGeom}(S),$$

equipped with coarse "error gauges" on objects, and we state an approximate equivalence between well-behaved subcategories. This makes precise the idea that TSRs and MSRs are two ε-equivalent *categories of realizations* of the same underlying sameness problem S.

Throughout this section, $\|\cdot\|$ denotes the natural inner-product norm in the relevant feature or ambient space, and $\mathcal{I}$ is the finite index set of modalities from Definition **??**. We fix a sameness structure S once and for all and suppress it from the notation when convenient.

3.1 Category of translation-based realizations

We first define the category whose objects are translation-based realizations of S.

Definition 8 (Category $\mathsf{SamTrans}(S)$)**.** *Fix a semantic sameness structure* S *as in Definition* **??***, with latent space* I*, modality index set* $\mathcal{I}$*, observation spaces* $\{X_i\}_{i\in\mathcal{I}}$*, rendering maps* $\{\pi_i : I \to X_i\}_{i\in\mathcal{I}}$*, sameness relation* $\approx$*, and benign path family* $\mathcal{P}_S(L)$*.*

Objects. *An object of* $\mathsf{SamTrans}(S)$ *is a translation-based semantic realization (TSR)*

$$\mathbf{T} = \Big(S,\, \{V_i^{\mathbf{T}}\}_{i\in\mathcal{I}},\, \{\varphi_i^{\mathbf{T}} : X_i \to V_i^{\mathbf{T}}\}_{i\in\mathcal{I}},\, \{T_{ij}^{\mathbf{T}} : V_i^{\mathbf{T}} \to V_j^{\mathbf{T}}\}_{i,j\in\mathcal{I}},\, G_{\mathbf{T}},\, \varepsilon_{\mathrm{trans}}^{\mathbf{T}}(\cdot)\Big)$$

as in Definition **??**. *For categorical purposes we implicitly identify each feature space* $V_i^{\mathbf{T}}$ *with a fixed Euclidean space* $\mathbb{R}^{d_i}$ *up to isometry, so that different TSRs share a common coordinate system per modality.*

Morphisms. *Given two TSRs* $\mathbf{T}, \mathbf{T}' \in \mathsf{SamTrans}(S)$*, a morphism*

$$R : \mathbf{T} \to \mathbf{T}'$$

with tolerances $(\eta_{\mathrm{enc}}, \eta_{\mathrm{trans}}) \geq 0$ *is a family of reparameterizations*

$$R = \{R_i : V_i^{\mathbf{T}} \to V_i^{\mathbf{T}'}\}_{i\in\mathcal{I}}$$

such that for all in-regime latent entities u *in a fixed subset* $\Omega^{\mathrm{ideal}} \subseteq I$ *(e.g., the union of points visited by benign paths in* $\mathcal{P}_S(L_\star)$ *for some operational horizon* $L_\star$*) the following hold:*

1. ***Encoder compatibility.*** *The* R_i *approximately intertwine the encoders:*

$$\big\|R_i(\varphi_i^{\mathbf{T}}(\pi_i(u))) - \varphi_i^{\mathbf{T}'}(\pi_i(u))\big\| \;\leq\; \eta_{\mathrm{enc}} \quad \textit{for all } i \in \mathcal{I}. \tag{3}$$

2. ***Translator compatibility.*** *The* R_i *approximately commute with translators:*

$$\big\|R_j(T_{ij}^{\mathbf{T}}(\varphi_i^{\mathbf{T}}(\pi_i(u)))) - T_{ij}^{\mathbf{T}'}(R_i(\varphi_i^{\mathbf{T}}(\pi_i(u))))\big\| \;\leq\; \eta_{\mathrm{trans}} \quad \textit{for all } i,j \in \mathcal{I}. \tag{4}$$

Composition and identities. *If* $R : \mathbf{T} \to \mathbf{T}'$ *and* $R' : \mathbf{T}' \to \mathbf{T}''$ *are morphisms, their composition* $R' \circ R : \mathbf{T} \to \mathbf{T}''$ *is defined componentwise:*

$$(R' \circ R)_i := R'_i \circ R_i, \quad i \in \mathcal{I}.$$

The identity morphism on $\mathbf{T}$ *is the family of identity maps*

$$\mathrm{id}_{\mathbf{T}} := \{\mathrm{id}_{V_i^{\mathbf{T}}}\}_{i\in\mathcal{I}}.$$

These operations make $\mathsf{SamTrans}(S)$ *a category.*

Intuitively, objects of $\mathsf{SamTrans}(S)$ are full translator systems built on a fixed sameness structure S; morphisms are approximate changes of feature coordinates that preserve the encoders and translators up to small tolerances. The in-regime set Ω^{ideal} encodes where these comparisons are intended to be meaningful.

3.2 Category of manifold-based realizations

We now define the category whose objects are manifold-based realizations of S.

Definition 9 (Category $\mathsf{SamGeom}(S)$)**.** *Fix the same sameness structure* S *as above.*

Objects. *An object of* $\mathsf{SamGeom}(S)$ *is a manifold-based semantic realization (MSR)*

$$\mathbf{M} = \Big(S,\, \{V_i^{\mathbf{M}}\}_{i\in\mathcal{I}},\, \{\varphi_i^{\mathbf{M}} : X_i \to V_i^{\mathbf{M}}\}_{i\in\mathcal{I}},\, (\mathcal{M}^{\mathbf{M}}, g^{\mathbf{M}}),\, \{\psi_i^{\mathbf{M}} : V_i^{\mathbf{M}} \to \mathcal{M}^{\mathbf{M}}\}_{i\in\mathcal{I}},\, \varepsilon_{\mathrm{dist}}^{\mathbf{M}}(\cdot)\Big)$$

as in Definition **??**. *For each* $i \in \mathcal{I}$ *the smooth map* $\psi_i^{\mathbf{M}} : V_i^{\mathbf{M}} \to \mathcal{M}^{\mathbf{M}}$ *is a parameterization whose image* $U_i^{\mathbf{M}} := \psi_i^{\mathbf{M}}(V_i^{\mathbf{M}})$ *is an open subset of the manifold, and*

$$\chi_i^{\mathbf{M}} := (\psi_i^{\mathbf{M}})^{-1} : U_i^{\mathbf{M}} \to V_i^{\mathbf{M}}$$

is the associated chart map. The distortion profile $\varepsilon_{\mathrm{dist}}^{\mathbf{M}}(\cdot)$ *controls how the chart metrics compare to the Riemannian metric* $g^{\mathbf{M}}$.

Given a latent entity $u \in I$ *for which modality* i *observes* u *(i.e.,* $\pi_i(u)$ *is defined), we write*

$$v_i^{\mathbf{M}}(u) := \varphi_i^{\mathbf{M}}(\pi_i(u)) \in V_i^{\mathbf{M}}, \qquad z^{\mathbf{M}}(u) := \psi_i^{\mathbf{M}}(v_i^{\mathbf{M}}(u)) \in \mathcal{M}^{\mathbf{M}}.$$

By consistency of the MSR, $z^{\mathbf{M}}(u)$ *does not depend on the choice of modality* i *when multiple modalities observe the same* u.

Morphisms. *Given two MSRs* $\mathbf{M}, \mathbf{M}' \in \mathsf{SamGeom}(S)$*, a morphism*

$$H : \mathbf{M} \to \mathbf{M}'$$

with tolerances $(\eta_{\mathrm{enc}}, \eta_{\mathrm{chart}}) \geq 0$ *is a pair*

$$H = (h, R), \qquad h : \mathcal{M}^{\mathbf{M}} \to \mathcal{M}^{\mathbf{M}'}, \quad R = \{R_i : V_i^{\mathbf{M}} \to V_i^{\mathbf{M}'}\}_{i\in\mathcal{I}},$$

where h *is smooth and* $\{R_i\}$ *are reparameterizations of the feature spaces, satisfying two compatibility conditions on an in-regime latent subset* $\Omega^{\mathrm{ideal}} \subseteq I$ *and the corresponding in-regime feature sets*

$$\Omega_i^{\mathrm{reg}} := \varphi_i^{\mathbf{M}}\big(\pi_i(\Omega^{\mathrm{ideal}})\big) \subseteq V_i^{\mathbf{M}}.$$

1. ***Encoder compatibility.*** *For all* $u \in \Omega^{\mathrm{ideal}}$ *and* $i \in \mathcal{I}$,

$$\big\|R_i(\varphi_i^{\mathbf{M}}(\pi_i(u))) - \varphi_i^{\mathbf{M}'}(\pi_i(u))\big\| \;\leq\; \eta_{\mathrm{enc}}. \tag{5}$$

2. h***-induced reparameterization and chart compatibility.*** *For each modality* $i \in \mathcal{I}$*, define the* h-induced reparameterization

$$R_h^i : V_i^{\mathbf{M}} \to V_i^{\mathbf{M}'}, \qquad R_h^i(v) := (\chi_i^{\mathbf{M}'} \circ h \circ \psi_i^{\mathbf{M}})(v).$$

This map transports a feature vector $v \in V_i^{\mathbf{M}}$ *onto the manifold* $\mathcal{M}^{\mathbf{M}}$*, applies* h*, and then pulls back to* $V_i^{\mathbf{M}'}$ *using the chart of* $\mathbf{M}'$.

We require that R *and the* h*-induced maps agree up to* η_{chart} *on in-regime features:*

$$\sup_{v\in\Omega_i^{\mathrm{reg}}} \big\|R_i(v) - R_h^i(v)\big\| \;\leq\; \eta_{\mathrm{chart}} \quad \text{for all } i \in \mathcal{I}. \tag{6}$$

Composition and identities. *If $H = (h, R) : \mathbf{M} \to \mathbf{M}'$ and $H' = (h', R') : \mathbf{M}' \to \mathbf{M}''$ are morphisms, their composition is*

$$H' \circ H := \big(h' \circ h,\ \{R'_i \circ R_i\}_{i \in \mathcal{I}}\big).$$

The identity morphism on $\mathbf{M}$ is

$$\mathrm{id}_{\mathbf{M}} := (\mathrm{id}_{\mathcal{M}^{\mathbf{M}}}, \{\mathrm{id}_{V_i^{\mathbf{M}}}\}_{i \in \mathcal{I}}).$$

These operations make $\mathsf{SamGeom}(S)$ a category.

Intuitively, morphisms in $\mathsf{SamGeom}(S)$ consist of a manifold map h (acting in state space) and a family of feature-space maps R_i; the condition (**??**) says that R_i must be close, on in-regime representations, to the reparameterization induced by h via the charts. This separates the definition of the induced map R^i_h from the compatibility requirement and makes the geometry of morphisms transparent.

3.3 Error gauges on realization categories

The categories $\mathsf{SamTrans}(S)$ and $\mathsf{SamGeom}(S)$ contain many realizations that differ in ways irrelevant to detection. To compare them coarsely, we introduce object-level *error gauges* that act as pseudo-metrics on the object sets. These gauges focus on the components that enter the detection theory: drift/distortion profiles and geometric behavior on in-regime entities.

Definition 10 (Error gauges on $\mathsf{SamTrans}(S)$ and $\mathsf{SamGeom}(S)$)**.** *Fix an in-regime latent subset $\Omega^{\mathrm{ideal}} \subseteq I$, for example the union of points visited by benign paths in $\mathcal{P}_S(L_\star)$, and let $\sigma_S^{\mathbf{M}}(L)$ be the semantic drift profile from Definition **??** whenever an MSR $\mathbf{M}$ exists.*

Gauge on $\mathsf{SamTrans}(S)$. *For two TSRs $\mathbf{T}, \mathbf{T}' \in \mathsf{SamTrans}(S)$, define*

$$\begin{aligned}\Delta_{\mathsf{SamTrans}}(\mathbf{T}, \mathbf{T}') := {} & \sup_{k \in \mathbb{N}} \big|\varepsilon^{\mathbf{T}}_{\mathrm{trans}}(k) - \varepsilon^{\mathbf{T}'}_{\mathrm{trans}}(k)\big| \\ & + \sup_{i,j \in \mathcal{I}} \sup_{u \in \Omega^{\mathrm{ideal}}} \big\|T^{\mathbf{T}}_{ij}(v^{\mathbf{T}}_i(u)) - T^{\mathbf{T}'}_{ij}(v^{\mathbf{T}'}_i(u))\big\|, \end{aligned} \tag{7}$$

*where $v^{\mathbf{T}}_i(u) := \varphi^{\mathbf{T}}_i(\pi_i(u))$ and $v^{\mathbf{T}'}_i(u) := \varphi^{\mathbf{T}'}_i(\pi_i(u))$. We tacitly identify $V^{\mathbf{T}}_i$ and $V^{\mathbf{T}'}_i$ with a common $\mathbb{R}^{d_i}$ up to isometry, so that the norm in (**??**) is well-defined.*

Gauge on $\mathsf{SamGeom}(S)$. *For two MSRs $\mathbf{M}, \mathbf{M}' \in \mathsf{SamGeom}(S)$, define*

$$\begin{aligned}\Delta_{\mathsf{SamGeom}}(\mathbf{M}, \mathbf{M}') := {} & \sup_{\ell > 0} \big|\varepsilon^{\mathbf{M}}_{\mathrm{dist}}(\ell) - \varepsilon^{\mathbf{M}'}_{\mathrm{dist}}(\ell)\big| \\ & + \sup_{u, u' \in \Omega^{\mathrm{ideal}}} \Big|d_{g^{\mathbf{M}}}(z^{\mathbf{M}}(u), z^{\mathbf{M}}(u')) - d_{g^{\mathbf{M}'}}(z^{\mathbf{M}'}(u), z^{\mathbf{M}'}(u'))\Big|, \end{aligned} \tag{8}$$

where $z^{\mathbf{M}}(u)$ and $z^{\mathbf{M}'}(u)$ are the induced manifold states defined above.

Both $\Delta_{\mathsf{SamTrans}}$ and $\Delta_{\mathsf{SamGeom}}$ are symmetric, non-negative, and vanish when the corresponding drift/distortion profiles and geometric behaviors agree exactly, but they may not separate all distinct objects (they are pseudo-metrics rather than metrics). They deliberately ignore implementation details that do not affect geometric behavior on in-regime entities.

3.4 ε-equivalence of realization categories

We next formalize the sense in which $\mathsf{SamTrans}(S)$ and $\mathsf{SamGeom}(S)$ will be treated as approximately equivalent. The key idea is that we do *not* expect strict categorical equivalence, but rather an equivalence up to small slack in the error gauges.

Definition 11 (ε-equivalence of categories of realizations)**.** *Let* $(\mathcal{C}, \Delta_{\mathcal{C}})$ *and* $(\mathcal{D}, \Delta_{\mathcal{D}})$ *be categories equipped with object-level pseudo-metrics (error gauges) on* $\mathrm{Ob}(\mathcal{C})$ *and* $\mathrm{Ob}(\mathcal{D})$*, respectively. Let* $\mathcal{C}^0 \subseteq \mathcal{C}$ *and* $\mathcal{D}^0 \subseteq \mathcal{D}$ *be full subcategories.*

We say that $\mathcal{C}^0$ *and* $\mathcal{D}^0$ *are* ε-equivalent *if there exist functors*

$$F : \mathcal{C}^0 \to \mathcal{D}^0, \qquad G : \mathcal{D}^0 \to \mathcal{C}^0,$$

and non-negative slack functions C_F, C_G *such that:*

1. *For every object* $C \in \mathrm{Ob}(\mathcal{C}^0)$*,*

$$\Delta_{\mathcal{C}}\big(G(F(C)), C\big) \;\le\; C_G(\varepsilon_C),$$

where ε_C *collects the relevant distortion parameters of* C *(e.g.,* $\varepsilon^{\mathbf{T}}_{\mathrm{trans}}$*,* $\varepsilon^{F(\mathbf{T})}_{\mathrm{dist}}$*, chart slack, etc.).*

2. *For every object* $D \in \mathrm{Ob}(\mathcal{D}^0)$*,*

$$\Delta_{\mathcal{D}}\big(F(G(D)), D\big) \;\le\; C_F(\varepsilon_D),$$

with ε_D *defined analogously.*

3. *(*First-order property.*) The slack functions are* first order *in the distortion parameters: when all distortion and slack terms (e.g.,* $\varepsilon_{\mathrm{trans}}$*,* $\varepsilon_{\mathrm{dist}}$*, chart-overlap slack* δ_{chart}*, translator asymmetry* δ_{asymm}*) tend to zero,* C_F *and* C_G *also tend to zero.*

In our setting, $\mathcal{C}^0$ and $\mathcal{D}^0$ will be well-behaved subcategories $\mathsf{SamTrans}^0(S) \subseteq \mathsf{SamTrans}(S)$ and $\mathsf{SamGeom}^0(S) \subseteq \mathsf{SamGeom}(S)$ on which smooth chartability and bounded-distortion conditions hold. The functors F and G will be the forward and reverse constructions F_S and G_S that move between translators and manifolds.

3.5 Structural equivalence theorem and implications

We are now ready to state the main structural theorem of this section in informal form. It says that, under a smooth chartability assumption, the TSR and MSR design spaces for a fixed sameness structure S are ε-equivalent in the sense of Definition **??**.

Theorem 1 (Realization categories are ε-equivalent (informal))**.** *Fix a semantic sameness structure* S *and an in-regime latent subset* $\Omega^{\mathrm{ideal}} \subseteq I$*. There exist:*

- *well-behaved subcategories*

$$\mathsf{SamTrans}^0(S) \subseteq \mathsf{SamTrans}(S), \qquad \mathsf{SamGeom}^0(S) \subseteq \mathsf{SamGeom}(S),$$

consisting of TSRs and MSRs that satisfy smooth chartability and bounded-distortion conditions;

- *functors*

$$F_S : \mathsf{SamTrans}^0(S) \to \mathsf{SamGeom}^0(S), \qquad G_S : \mathsf{SamGeom}^0(S) \to \mathsf{SamTrans}^0(S),$$

  that:

  1. *preserve the underlying sameness structure S and feature spaces $\{V_i, \varphi_i\}_{i \in \mathcal{I}}$;*
  2. *send translators to manifolds by quotienting and gluing charts (Section 4);*
  3. *send manifolds to translators by extracting chart transitions (Section 5);*

- *slack functions C_F, C_G which are first order in the distortion and chartability parameters (translator drift $\varepsilon_{\text{trans}}$, metric distortion $\varepsilon_{\text{dist}}$, chart-overlap slack δ_{chart}, and translator asymmetry δ_{asymm}),*

such that:

$$\Delta_{\mathsf{SamTrans}}(\mathbf{T},\ G_S(F_S(\mathbf{T}))) \ \le\ C_G(\varepsilon^{\mathbf{T}}_{\text{trans}}, \varepsilon^{F_S(\mathbf{T})}_{\text{dist}}, \delta_{\text{chart}}(\mathbf{T}, F_S(\mathbf{T})), \delta_{\text{asymm}}(\mathbf{T})) \quad \textit{for all } \mathbf{T} \in \mathsf{SamTrans}^0(S),$$

$$\Delta_{\mathsf{SamGeom}}(\mathbf{M},\ F_S(G_S(\mathbf{M}))) \ \le\ C_F(\varepsilon^{\mathbf{M}}_{\text{dist}}, \varepsilon^{G_S(\mathbf{M})}_{\text{trans}}, \delta_{\text{chart}}(G_S(\mathbf{M}), \mathbf{M}), \delta_{\text{asymm}}(G_S(\mathbf{M}))) \quad \textit{for all } \mathbf{M} \in \mathsf{SamGeom}^0(S)$$

In particular, as all distortion and slack terms tend to zero, the round-trip errors in the gauges $\Delta_{\mathsf{SamTrans}}$ and $\Delta_{\mathsf{SamGeom}}$ tend to zero: $\mathsf{SamTrans}^0(S)$ and $\mathsf{SamGeom}^0(S)$ are ε-equivalent categories of realizations of S.

The detailed constructions of F_S and G_S and the proof of Theorem **??** are given in Sections 4–5. At a conceptual level:

- For a fixed sameness structure S, TSRs and MSRs are two different *realizations* of the same problem: $\mathsf{SamTrans}^0(S)$ and $\mathsf{SamGeom}^0(S)$ are two design spaces for the same semantic sameness structure.
- The functors F_S and G_S witness an ε-equivalence between these design spaces, showing that—under smooth chartability—they are dual up to controlled slack. An engineer who starts from only S can choose either a translator-based or manifold-based architecture, with the assurance that the theoretical guarantees of the resulting detector are comparable.
- Because the equivalence is functorial and comes with explicit slack bounds, *theorems and tools transfer*: risk bounds or stability results proven in the manifold view can be pulled back to translator systems via G_S, and translator-side diagnostics (e.g., spectral audits on format graphs) can be pushed forward to furnish new manifold-side checks via F_S. This transfer principle underlies the Error Budget Transfer Theorem in Section 6.

4 Smooth chartability and the functor F_S

Fix a semantic sameness structure S and the associated categories $\mathsf{SamTrans}(S)$ and $\mathsf{SamGeom}(S)$ from Section **??**. Throughout this section, we work with ideal latent entities $u \in \Omega^{\text{ideal}} \subseteq I$ and their encoder outputs $v_i(u) := \varphi_i(\pi_i(u)) \in V_i$ as in Section **??**. We write $\|\cdot\|$ for the natural inner-product norm in the relevant feature or ambient space.

Our goal is to construct a functor

$$F_S : \mathsf{SamTrans}^0(S) \longrightarrow \mathsf{SamGeom}^0(S)$$

on a well-behaved subcategory $\mathsf{SamTrans}^0(S) \subseteq \mathsf{SamTrans}(S)$, sending a translation-based realization (TSR) $\mathbf{T}$ to a manifold-based realization (MSR) $F_S(\mathbf{T})$, and a morphism $R : \mathbf{T} \to \mathbf{T}'$ to a morphism $F_S(R) : F_S(\mathbf{T}) \to F_S(\mathbf{T}')$. Crucially, F_S must be *constructed* from the TSR data, not defined by fiat.

We proceed in three steps:

1. We build a canonical *translator metric quotient* $(\widetilde{\mathcal{M}}^{\mathbf{T}}, \tilde{d}_{\mathbf{T}})$ from any TSR $\mathbf{T}$ (Section **??**).
2. We define an *intrinsic* notion of smooth chartability for TSRs that depends only on $\mathbf{T}$ itself and its latent coverage (Section **??**), and use it to carve out the subcategory $\mathsf{SamTrans}^0(S)$.
3. We impose a single, explicit assumption of *Riemannian realizability* for these metric spaces and use it to define F_S on objects and morphisms (Section **??**).

This isolates the genuinely geometric hypothesis—that the translator-induced metric space is quasi-isometric to a smooth manifold— from the intrinsic regularity properties of the TSR itself.

4.1 The translator metric quotient

Let $\mathbf{T} \in \mathsf{SamTrans}(S)$ be a TSR with components

$$\mathbf{T} = \Big(S,\ \{V_i^{\mathbf{T}}\}_{i\in\mathcal{I}},\ \{\varphi_i^{\mathbf{T}}\}_{i\in\mathcal{I}},\ \{T_{ij}^{\mathbf{T}}\}_{(i,j)\in E_{\mathbf{T}}},\ G_{\mathbf{T}},\ \varepsilon_{\text{trans}}^{\mathbf{T}}(\cdot)\Big),$$

as in Definition **??**. For each modality $i \in \mathcal{I}$, recall the in-regime feature subset

$$\Omega_i^{\text{reg}}(\mathbf{T}) := \{v_i(u) = \varphi_i^{\mathbf{T}}(\pi_i(u)) : u \in \Omega^{\text{ideal}},\ \pi_i(u) \text{ defined}\} \subseteq V_i^{\mathbf{T}},$$

and define the disjoint union

$$X^{\mathbf{T}} := \bigsqcup_{i\in\mathcal{I}} \Omega_i^{\text{reg}}(\mathbf{T}), \qquad x = (i, v_i) \in X^{\mathbf{T}}.$$

Local edge costs. We endow $X^{\mathbf{T}}$ with two families of elementary *edges*:

- *Within-chart edges.* For any $i \in \mathcal{I}$ and any $v_i, v_i' \in \Omega_i^{\text{reg}}(\mathbf{T})$ we introduce an edge $e : (i, v_i) \to (i, v_i')$ with cost
$$c_{\text{intra}}(e) := \|v_i - v_i'\|.$$
- *Translator edges.* For any $(i,j) \in E_{\mathbf{T}}$ and any $v_i \in \Omega_i^{\text{reg}}(\mathbf{T})$ such that $T_{ij}^{\mathbf{T}}(v_i)$ lies in $\Omega_j^{\text{reg}}(\mathbf{T})$, we introduce an edge
$$e : (i, v_i) \longrightarrow (j, T_{ij}^{\mathbf{T}}(v_i))$$
with cost
$$c_{\text{trans}}(e) := \|T_{ij}^{\mathbf{T}}(v_i)\|.$$

(Any consistent bounded positive edge-weighting based on the translator action would suffice; the specific choice is not essential for the subsequent theory.)

Path pseudo-metric on $X^{\mathbf{T}}$. A *finite edge-path* in $X^{\mathbf{T}}$ is a sequence $p = (e_1, \ldots, e_m)$ of composable edges. We define its cost

$$\mathrm{len}(p) := \sum_{r=1}^{m} c(e_r),$$

where $c(e)$ is $c_{\mathrm{intra}}(e)$ or $c_{\mathrm{trans}}(e)$ depending on the edge type. For any two points $x, y \in X^{\mathbf{T}}$, we define the *path pseudo-metric*

$$\tilde{d}_{\mathbf{T}}^{\mathrm{raw}}(x, y) := \inf\{\mathrm{len}(p) : p \text{ is a finite edge-path from } x \text{ to } y\}.$$

Standard arguments show that $\tilde{d}_{\mathbf{T}}^{\mathrm{raw}}$ is a pseudo-metric: it is symmetric, nonnegative, and satisfies the triangle inequality, but may assign zero distance to distinct points.

Definition 12 (Translator metric quotient)**.** *Define an equivalence relation $\sim_0$ on $X^{\mathbf{T}}$ by*

$$x \sim_0 y \quad \Longleftrightarrow \quad \tilde{d}_{\mathbf{T}}^{\mathrm{raw}}(x, y) = 0.$$

The translator metric quotient *of $\mathbf{T}$ is the metric space*

$$(\widetilde{\mathcal{M}}^{\mathbf{T}}, \tilde{d}_{\mathbf{T}})$$

obtained as the quotient $X^{\mathbf{T}}/\sim_0$ with metric

$$\tilde{d}_{\mathbf{T}}([x], [y]) := \tilde{d}_{\mathbf{T}}^{\mathrm{raw}}(x, y),$$

which is well-defined and satisfies the metric axioms.

For each modality $i \in \mathcal{I}$ there is a canonical quotient parameterization

$$\tilde{\psi}_i^{\mathbf{T}} : \Omega_i^{\mathrm{reg}}(\mathbf{T}) \to \widetilde{\mathcal{M}}^{\mathbf{T}}, \qquad \tilde{\psi}_i^{\mathbf{T}}(v_i) := [(i, v_i)].$$

Intuitively, $\widetilde{\mathcal{M}}^{\mathbf{T}}$ is the "patchwork quilt" you obtain by gluing together the in-regime feature charts along the translator edges, with $\tilde{d}_{\mathbf{T}}$ measuring the shortest-path cost induced by chart distances and translator hops. The construction is canonical (up to isometry) and depends only on the TSR $\mathbf{T}$.

4.2 Intrinsic smooth chartability

The previous construction associates a metric space $(\widetilde{\mathcal{M}}^{\mathbf{T}}, \tilde{d}_{\mathbf{T}})$ to any TSR $\mathbf{T}$. However, for an arbitrary $\mathbf{T}$ this space may be pathological: it need not resemble a manifold at any scale. We now isolate a set of purely *intrinsic* conditions on $\mathbf{T}$ under which the translator metric quotient behaves like a well-controlled, finite-dimensional geometric object.

Definition 13 (Intrinsic smooth chartability)**.** *A TSR $\mathbf{T}$ is* intrinsically smoothly chartable *if the following conditions hold on an in-regime subset Ω^{ideal} and the associated feature sets $\Omega_i^{\mathrm{reg}}(\mathbf{T})$:*

1. ***Translator regularity.*** *For each $(i, j) \in E_{\mathbf{T}}$, the translator $T_{ij}^{\mathbf{T}} : V_i^{\mathbf{T}} \to V_j^{\mathbf{T}}$ is C^2 on an open* *neighbourhood of $\Omega_i^{\mathrm{reg}}(\mathbf{T})$, with uniformly bounded Jacobian and inverse Jacobian:*

$$\sup_{v_i \in \Omega_i^{\mathrm{reg}}(\mathbf{T})} \|DT_{ij}^{\mathbf{T}}(v_i)\| \leq C_{\mathrm{Jac}}, \qquad \sup_{v_i \in \Omega_i^{\mathrm{reg}}(\mathbf{T})} \|DT_{ij}^{\mathbf{T}}(v_i)^{-1}\| \leq C_{\mathrm{Jac}}.$$

2. ***Round-trip asymmetry bound.*** *There exists $\delta_{\mathrm{asymm}}^{\mathbf{T}} < \infty$ such that for every $(i,j) \in E_{\mathbf{T}}$ and every $v_i \in \Omega_i^{\mathrm{reg}}(\mathbf{T})$ with $T_{ij}^{\mathbf{T}}(v_i) \in \Omega_j^{\mathrm{reg}}(\mathbf{T})$ and $T_{ji}^{\mathbf{T}}(T_{ij}^{\mathbf{T}}(v_i)) \in \Omega_i^{\mathrm{reg}}(\mathbf{T})$ we have*

$$\|T_{ji}^{\mathbf{T}}(T_{ij}^{\mathbf{T}}(v_i)) - v_i\| \leq \delta_{\mathrm{asymm}}^{\mathbf{T}}.$$

3. ***Cocycle bound.*** *There exists $\delta_{\mathrm{cocycle}}^{\mathbf{T}} < \infty$ such that for any triple (i,j,k) with $(i,j),(j,k),(i,k) \in E_{\mathbf{T}}$ and any $v_i \in \Omega_i^{\mathrm{reg}}(\mathbf{T})$ lying in the domain of all three compositions,*

$$\|T_{jk}^{\mathbf{T}}(T_{ij}^{\mathbf{T}}(v_i)) - T_{ik}^{\mathbf{T}}(v_i)\| \leq \delta_{\mathrm{cocycle}}^{\mathbf{T}}.$$

   This ensures that multi-hop translator paths are consistent, up to a controlled slack, with direct translations.

4. ***Latent coverage and drift control.*** *There is an operational horizon $L_\star > 0$ such that:*
   - *every benign latent path $\gamma \in \mathcal{P}_S(L_\star)$ projects to in-regime observations $\pi_i(\gamma(t))$ and hence to in-regime features $v_i(\gamma(t)) = \varphi_i^{\mathbf{T}}(\pi_i(\gamma(t)))$;*
   - *the drift profile $\varepsilon_{\mathrm{trans}}^{\mathbf{T}}(k)$ is finite for all k and satisfies the per-time-step bound of Definition **??** on these paths.*

Definition 14 (Well-behaved subcategory $\mathsf{SamTrans}^0(S)$)**.** *We define $\mathsf{SamTrans}^0(S)$ as the full subcategory of $\mathsf{SamTrans}(S)$ whose objects are intrinsically smoothly chartable TSRs in the sense of Definition **??**, and whose morphisms are exactly the morphisms of $\mathsf{SamTrans}(S)$ between such objects.*

By construction, membership in $\mathsf{SamTrans}^0(S)$ is a property of the TSR $\mathbf{T}$ and its translators alone. No manifold is mentioned in Definition **??**: it is a falsifiable hypothesis about the regularity, consistency, and coverage of the translators.

4.3 Riemannian realizability of the translator metric

The translator metric quotient $(\widetilde{\mathcal{M}}^{\mathbf{T}}, \tilde{d}_{\mathbf{T}})$ from Definition **??** is a canonical metric space built from $\mathbf{T}$ alone. Intrinsic chartability ensures that this space behaves like a finite-dimensional, locally well-behaved metric space: translators are smooth with bounded Jacobians, roundtrip and cocycle inconsistencies are small, and latent data populate a connected, bounded region.

To bridge from metric geometry to Riemannian geometry we isolate the following explicit assumption, which is the only place where we posit the existence of a smooth manifold "shadow."

Assumption 1 (Riemannian realizability of translator metric)**.** *For every $\mathbf{T} \in \mathsf{SamTrans}^0(S)$, the translator metric quotient $(\widetilde{\mathcal{M}}^{\mathbf{T}}, \tilde{d}_{\mathbf{T}})$ admits a Riemannian realization: there exists a smooth, finite-dimensional Riemannian manifold $(\mathcal{M}^{\mathbf{T}}, g^{\mathbf{T}})$ and a quasi-isometry*

$$q^{\mathbf{T}} : (\widetilde{\mathcal{M}}^{\mathbf{T}}, \tilde{d}_{\mathbf{T}}) \longrightarrow (\mathcal{M}^{\mathbf{T}}, d_{g^{\mathbf{T}}})$$

with distortion bounded by a function $\varepsilon^{\mathbf{T}}_{\text{realize}}(\cdot)$ that depends continuously and monotonically on the intrinsic slack parameters $\varepsilon^{\mathbf{T}}_{\text{trans}}(\cdot)$, $\delta^{\mathbf{T}}_{\text{asymm}}$ and $\delta^{\mathbf{T}}_{\text{cocycle}}$.

Moreover, we assume that $q^{\mathbf{T}}$ is coarse-biLipschitz on the in-regime region populated by Ω^{ideal}, so that for all x, y in this region

$$\Big|\frac{d_{g^{\mathbf{T}}}(q^{\mathbf{T}}(x), q^{\mathbf{T}}(y))}{\tilde{d}_{\mathbf{T}}(x,y)} - 1\Big| \leq \varepsilon^{\mathbf{T}}_{\text{realize}}(\tilde{d}_{\mathbf{T}}(x,y)).$$

We do *not* attempt to prove Assumption **??** in this work; it is our single, explicit geometric hypothesis. In practice, it is an empirical question whether a given TSR $\mathbf{T}$ satisfies this assumption, and Section **??** discusses audit procedures for probing Riemannian realizability.

4.4 Constructing F_S on objects

With the metric quotient and realizability assumption in hand, we can now define F_S on objects in a non-circular way.

Definition 15 (Functor F_S on objects)**.** *Let $\mathbf{T} \in \mathsf{SamTrans}^0(S)$ be an intrinsically smoothly chartable TSR. Consider its translator metric quotient $(\widetilde{\mathcal{M}}^{\mathbf{T}}, \tilde{d}_{\mathbf{T}})$ and a Riemannian realization $(\mathcal{M}^{\mathbf{T}}, g^{\mathbf{T}})$ with quasi-isometry $q^{\mathbf{T}}$ given by Assumption **??**. For each modality $i \in \mathcal{I}$, define the parameterization*

$$\psi^{\mathbf{T}}_i : V^{\mathbf{T}}_i \supseteq \Omega^{\text{reg}}_i(\mathbf{T}) \longrightarrow \mathcal{M}^{\mathbf{T}}, \qquad \psi^{\mathbf{T}}_i(v_i) := q^{\mathbf{T}}(\tilde{\psi}^{\mathbf{T}}_i(v_i)),$$

*where $\tilde{\psi}^{\mathbf{T}}_i$ is the quotient map from Definition **??**.*

We define

$$F_S(\mathbf{T}) := \mathbf{M}^{\mathbf{T}} := \Big(S,\ \{V^{\mathbf{T}}_i\}_{i\in\mathcal{I}},\ \{\varphi^{\mathbf{T}}_i\}_{i\in\mathcal{I}},\ (\mathcal{M}^{\mathbf{T}}, g^{\mathbf{T}}),\ \{\psi^{\mathbf{T}}_i\}_{i\in\mathcal{I}},\ \varepsilon^{\mathbf{T}}_{\text{dist}}(\cdot)\Big),$$

where the distortion profile $\varepsilon^{\mathbf{T}}_{\text{dist}}$ is defined by pulling back the quasi-isometry distortion:

$$\Big|\frac{\delta^{\mathbf{T}}_i(z_a, z_b)}{d_{g^{\mathbf{T}}}(z_a, z_b)} - 1\Big| \leq \varepsilon^{\mathbf{T}}_{\text{dist}}(d_{g^{\mathbf{T}}}(z_a, z_b)), \qquad z_a, z_b \in \psi^{\mathbf{T}}_i(\Omega^{\text{reg}}_i(\mathbf{T})),$$

with chart distances $\delta^{\mathbf{T}}_i(z_a, z_b) := \|\chi^{\mathbf{T}}_i(z_a) - \chi^{\mathbf{T}}_i(z_b)\|$ and chart maps $\chi^{\mathbf{T}}_i := (\psi^{\mathbf{T}}_i)^{-1}$ on their images.

*By construction, $\mathbf{M}^{\mathbf{T}}$ is an MSR of S in the sense of Definition **??**, and we have $F_S(\mathbf{T}) \in \mathsf{SamGeom}^0(S)$.*

Importantly, $F_S(\mathbf{T})$ is now determined (up to quasi-isometry) by $\mathbf{T}$ itself: the only freedom lies in the choice of Riemannian realization $(\mathcal{M}^{\mathbf{T}}, g^{\mathbf{T}})$ of the canonical metric quotient. Different choices of $q^{\mathbf{T}}$ give rise to MSRs that are close in the error gauge $\Delta_{\mathsf{SamGeom}}$ of Definition **??**, and Theorem **??** will show that this slack is first-order in the intrinsic errors of $\mathbf{T}$.

4.5 Constructing F_S on morphisms

We next define F_S on morphisms. Let $R : \mathbf{T} \to \mathbf{T}'$ be a morphism in $\mathsf{SamTrans}^0(S)$ as in Definition **??**, with reparameterizations $R_i : V^{\mathbf{T}}_i \to V^{\mathbf{T}'}_i$ satisfying the encoder- and translator-compatibility bounds on ideal data.

The TSR morphism R induces a map between the translator metric quotients,

$$\tilde{h}^R : \widetilde{\mathcal{M}}^{\mathbf{T}} \longrightarrow \widetilde{\mathcal{M}}^{\mathbf{T}'},$$

defined on representatives by

$$\tilde{h}^R(\tilde{\psi}_i^{\mathbf{T}}(v_i)) := \tilde{\psi}_i^{\mathbf{T}'}(R_i(v_i)),$$

and extended by continuity. The compatibility conditions in Definition **??** ensure that $\tilde{h}^R$ is well-defined up to the metric identification $\sim_0$ and is Lipschitz (with constants depending on the morphism error tolerances $(\eta_{\text{enc}}, \eta_{\text{trans}})$).

Using the quasi-isometries $q^{\mathbf{T}}$ and $q^{\mathbf{T}'}$ from Assumption **??**, we then define the manifold-level map

$$h^R : \mathcal{M}^{\mathbf{T}} \longrightarrow \mathcal{M}^{\mathbf{T}'}, \qquad h^R := q^{\mathbf{T}'} \circ \tilde{h}^R \circ (q^{\mathbf{T}})^{-1},$$

understanding $(q^{\mathbf{T}})^{-1}$ as a coarse inverse on the in-regime region.

Definition 16 (Functor F_S on morphisms)**.** *For a morphism $R : \mathbf{T} \to \mathbf{T}'$ in $\mathsf{SamTrans}^0(S)$, we define*

$$F_S(R) := H^R := (h^R,\ R) : \mathbf{M}^{\mathbf{T}} \longrightarrow \mathbf{M}^{\mathbf{T}'},$$

where $\mathbf{M}^{\mathbf{T}} = F_S(\mathbf{T})$ and $\mathbf{M}^{\mathbf{T}'} = F_S(\mathbf{T}')$, and h^R is the manifold map defined above.

The encoder-compatibility condition for H^R,

$$\|R_i(\varphi_i^{\mathbf{M}^{\mathbf{T}}}(\pi_i(u))) - \varphi_i^{\mathbf{M}^{\mathbf{T}'}}(\pi_i(u))\| \leq \eta'_{\text{enc}},$$

is inherited from R, while the chart-compatibility condition,

$$\|R_i(\varphi_i^{\mathbf{M}^{\mathbf{T}}}(\pi_i(u))) - \chi_i^{\mathbf{M}^{\mathbf{T}'}}(h^R(z^{\mathbf{T}}(u)))\| \leq \eta'_{\text{chart}},$$

with $z^{\mathbf{T}}(u)$ the latent state in $\mathcal{M}^{\mathbf{T}}$, follows from the construction of h^R and the quasi-isometry bounds. Thus H^R is a morphism in $\mathsf{SamGeom}(S)$ in the sense of Definition **??***.*

Proposition 1 (Functoriality and stability of F_S)**.** *The assignment F_S of Definitions* **??** *and* **??** *defines a functor*

$$F_S : \mathsf{SamTrans}^0(S) \longrightarrow \mathsf{SamGeom}^0(S)$$

satisfying:

1. *For every object $\mathbf{T} \in \mathsf{SamTrans}^0(S)$, $F_S(\mathrm{id}_{\mathbf{T}}) = \mathrm{id}_{F_S(\mathbf{T})}$.*
2. *For any composable morphisms $R : \mathbf{T} \to \mathbf{T}'$ and $R' : \mathbf{T}' \to \mathbf{T}''$ in $\mathsf{SamTrans}^0(S)$,*

$$F_S(R' \circ R) = F_S(R') \circ F_S(R).$$

3. *There exists a universal, increasing function C_F such that for any $\mathbf{T}, \mathbf{T}' \in \mathsf{SamTrans}^0(S)$ the error gauge satisfies*

$$\Delta_{\mathsf{SamGeom}}(F_S(\mathbf{T}), F_S(\mathbf{T}')) \leq C_F\Big(\Delta_{\mathsf{SamTrans}}(\mathbf{T}, \mathbf{T}'),\ \delta_{\text{chart}}(\mathbf{T}),\ \delta_{\text{chart}}(\mathbf{T}')\Big),$$

where $\delta_{\text{chart}}(\mathbf{T})$ is the chart-overlap slack of $\mathbf{T}$ from Definition **??***. Moreover, C_F is first-order* *in its arguments: it vanishes as* $\Delta_{\mathsf{SamTrans}} \to 0$ *and* $\delta_{\mathrm{chart}} \to 0$.

Proof sketch. Functoriality of F_S on objects and morphisms follows from the definitions of the translator metric quotient and the functoriality of the quasi-isometry construction in Assumption **??**. The stability bound in (iii) leverages the fact that the distortion profile $\varepsilon^{\mathbf{T}}_{\mathrm{dist}}$ is controlled by the intrinsic errors of $\mathbf{T}$ via the quasi-isometry, and that the error gauge $\Delta_{\mathsf{SamGeom}}$ is built from these distortion profiles and chart disagreements. Full details are deferred to Section 6, where we prove the stronger Error Budget Transfer Theorem. □

4.6 Discussion: what F_S really does

The functor F_S no longer hides a manifold inside the definition of smooth chartability. Instead, it proceeds in two conceptually clean stages:

- **From flat maps to a quilt.** For any intrinsically well-behaved TSR $\mathbf{T}$, we canonically construct the translator metric quotient $(\widetilde{\mathcal{M}}^{\mathbf{T}}, \tilde{d}_{\mathbf{T}})$ that encodes how features in different modalities can be transported and compared via translators. This is the "patchwork quilt" of all flat maps stitched together along their edges, with a well-defined metric.
- **From the quilt to the globe.** The single geometric hypothesis of Riemannian realizability asserts that this quilt is quasi-isometric to some smooth manifold $(\mathcal{M}^{\mathbf{T}}, g^{\mathbf{T}})$. The functor F_S then declares this manifold, together with the induced parameterizations $\psi_i^{\mathbf{T}}$, to be the MSR $F_S(\mathbf{T})$.

In this way, F_S is a genuine construction rather than a renaming: it builds a metric geometry from the translators, and only then—under an explicit, isolated assumption—passes to a smooth Riemannian geometry. The error gauge stability in Proposition **??** will feed directly into the Error Budget Transfer Theorem in Section 6, where we show that detection guarantees can be transported across this functor with only first-order slack.

5 From manifolds back to translators: the functor G_S

Section **??** constructed a functor $F_S : \mathsf{SamTrans}^0(S) \to \mathsf{SamGeom}^0(S)$ by starting from a smoothly chartable TSR $\mathbf{T}$, forming its translator metric quotient $(\widetilde{\mathcal{M}}^{\mathbf{T}}, \tilde{d}_{\mathbf{T}})$, and—under a Riemannian realizability assumption—obtaining a manifold-based realization $F_S(\mathbf{T})$ whose charts approximate the original translators. In this section we build the reverse functor

$$G_S : \mathsf{SamGeom}^0(S) \to \mathsf{SamTrans}^0(S),$$

which takes a well-parameterized manifold-based realization $\mathbf{M}$ and extracts a canonical translator-based realization $G_S(\mathbf{M})$. At the level of the Globe–Flat Maps analogy, F_S turns a consistent atlas of flat maps into a globe; G_S takes a globe with well-behaved charts and recovers a translator graph encoding the chart transitions.

Throughout this section we fix a semantic sameness structure S and work within the realization categories $\mathsf{SamGeom}(S)$ and $\mathsf{SamTrans}(S)$ from Section **??**. Norms $\|\cdot\|$ denote the natural inner-product norms on the relevant feature spaces, as in Section **??**.

5.1 Geometrically well-parameterized MSRs

We first specify which manifold-based realizations admit a stable translation view. Intuitively, an MSR is *geometrically well-parameterized* when its chart maps behave like well-conditioned coordinate systems on the region traced out by benign latent paths, with controlled distortion and overlap.

Definition 17 (Geometrically well-parameterized MSR)**.** *Let* $\mathbf{M}$ *be an MSR of* S *in the sense of Definition* **??***:*

$$\mathbf{M} = \Big(S,\ \{V_i^{\mathbf{M}}\}_{i\in\mathcal{I}},\ \{\varphi_i^{\mathbf{M}}\}_{i\in\mathcal{I}}, (\mathcal{M}^{\mathbf{M}}, g^{\mathbf{M}}),\ \{\psi_i^{\mathbf{M}}\}_{i\in\mathcal{I}}, \varepsilon_{\mathrm{dist}}^{\mathbf{M}}(\cdot)\Big).$$

Write $U_i^{\mathbf{M}} := \psi_i^{\mathbf{M}}(V_i^{\mathbf{M}})$ *and* $\chi_i^{\mathbf{M}} := (\psi_i^{\mathbf{M}})^{-1} : U_i^{\mathbf{M}} \to V_i^{\mathbf{M}}$*, and define* $z^{\mathbf{M}}(u)$ *as in Definition* **??***.*

We say that $\mathbf{M}$ *is* geometrically well-parameterized *if the following conditions hold.*

(i) ***Shared feature model.*** *For each modality* $i \in \mathcal{I}$ *the feature spaces and encoders agree (up to isometry) with those used in the TSR view: there exist inner-product-preserving isomorphisms* $Q_i : V_i^{\mathbf{M}} \to V_i$ *such that, after identifying* $V_i^{\mathbf{M}}$ *with* V_i *via* Q_i*, we may write* $\varphi_i^{\mathbf{M}} = \varphi_i$ *and use the common notation* V_i*,* φ_i*.*

(ii) ***Chart regularity and bounded distortion.*** *Each parameterization* $\psi_i : V_i \to \mathcal{M}^{\mathbf{M}}$ *is* C^2 *on an in-regime region* $V_i^{\mathrm{reg}} \subseteq V_i$ *that contains the image of the benign observation set under* $\varphi_i \circ \pi_i$*. The corresponding chart maps* $\chi_i : U_i \to V_i$ *with* $U_i := \psi_i(V_i^{\mathrm{reg}})$ *satisfy the distortion profile of Definition* **??***: for all* $z_a, z_b \in U_i$ *with* $d_{g^{\mathbf{M}}}(z_a, z_b) > 0$*,*

$$\left|\frac{\delta_i(z_a, z_b)}{d_{g^{\mathbf{M}}}(z_a, z_b)} - 1\right| \le \varepsilon_{\mathrm{dist}}^{\mathbf{M}}\big(d_{g^{\mathbf{M}}}(z_a, z_b)\big), \qquad \delta_i(z_a, z_b) := \big\|\chi_i(z_a) - \chi_i(z_b)\big\|_{V_i}. \tag{9}$$

(iii) ***Chart overlap control.*** *For any* i, j *with* $U_i \cap U_j \ne \varnothing$ *and any* $z \in U_i \cap U_j$*, the local chart transition* $\chi_j \circ \psi_i : V_i^{\mathrm{reg}} \to V_j$ *is well-defined on a neighborhood of* $\chi_i(z)$ *and has uniformly bounded Jacobian and inverse Jacobian there. Equivalently, the family* $\{\chi_i\}$ *forms a uniformly well-conditioned atlas on the in-regime region of* $\mathcal{M}^{\mathbf{M}}$*.*

(iv) ***Benign path coverage.*** *There exists an operational horizon* $L_\star$ *such that for every benign latent path* $\gamma \in \mathcal{P}_S(L_\star)$*, the induced manifold path* $z^{\mathbf{M}}(\cdot)$ *remains in* $\bigcup_i U_i$ *and crosses only finitely many chart overlaps. The semantic drift profile* $\sigma_S^{\mathbf{M}}(L_\star)$ *from Definition* **??** *is finite.*

Condition (i) ensures that the manifold view and the translator view share the same feature model; (ii)–(iii) guarantee that chart coordinates are well-behaved and that chart transitions are controlled; and (iv) ensures that the benign path family never leaves the region where charts are valid. Taken together, these conditions are the geometric mirror of the smooth chartability conditions imposed on TSRs in Section **??**.

Definition 18 (Well-behaved geometric subcategory)**.** *Let* $\mathsf{SamGeom}^0(S)$ *denote the full subcategory of* $\mathsf{SamGeom}(S)$ *whose objects are geometrically well-parameterized MSRs in the sense of Definition* **??***, and whose morphisms are those of* $\mathsf{SamGeom}(S)$ *(Definition* **??***) between such objects.*

5.2 Constructing G_S on objects

We now define G_S on objects by "unfolding" a well-parameterized manifold back into a TSR whose translators are given by chart transitions.

Definition 19 (Functor G_S on objects)**.** *Let* $\mathbf{M} \in \mathsf{SamGeom}^0(S)$ *with data as in Definition* **??**. *We define the TSR* $G_S(\mathbf{M}) =: \mathbf{T}^{\mathbf{M}}$ *by*

$$\mathbf{T}^{\mathbf{M}} = \Big(S,\ \{V_i^{\mathbf{M}}\}_{i\in\mathcal{I}},\ \{\varphi_i^{\mathbf{M}}\}_{i\in\mathcal{I}}, \{T_{ij}^{\mathbf{M}}\}_{(i,j)\in E^{\mathbf{M}}},\ G_{\mathbf{T}^{\mathbf{M}}}, \varepsilon_{\text{trans}}^{\mathbf{M}}(\cdot)\Big),$$

where:

(i) ***Feature spaces and encoders.*** *We reuse the feature spaces and encoders of* $\mathbf{M}$*:* $V_i^{\mathbf{M}}$ *and* $\varphi_i^{\mathbf{M}}$ *become the feature spaces and encoders of* $\mathbf{T}^{\mathbf{M}}$.

(ii) ***Translator graph.*** *The vertex set of* $G_{\mathbf{T}^{\mathbf{M}}}$ *is* $\mathcal{I}$*, and we include a directed edge* (i, j) *whenever* $U_i \cap U_j$ *is non-empty on the in-regime region traced by benign paths. Thus translators exist exactly where chart overlaps do.*

(iii) ***Chart-transition translators.*** *For each edge* $(i, j) \in E^{\mathbf{M}}$ *we define the translator* $T_{ij}^{\mathbf{M}} : V_i^{\mathbf{M}} \to V_j^{\mathbf{M}}$ *on the in-regime domain* V_i^{reg} *by*

$$T_{ij}^{\mathbf{M}}(v_i) := \chi_j^{\mathbf{M}}(\psi_i^{\mathbf{M}}(v_i)) \qquad \textit{whenever } \psi_i^{\mathbf{M}}(v_i) \in U_i \cap U_j. \tag{10}$$

On points not in V_i^{reg} *or not mapping into* $U_i \cap U_j$ *we leave* $T_{ij}^{\mathbf{M}}$ *undefined; the TSR semantics restrict attention to in-regime usage as in Definition* **??**.

(iv) ***Drift profile.*** *We define the drift profile* $\varepsilon_{\text{trans}}^{\mathbf{M}}(k)$ *for* k*-hop translator paths by composing the one-hop distortions induced by the chart maps. Informally, for* $k = 1$ *and an in-regime latent entity* u *with* $z^{\mathbf{M}}(u) \in U_i \cap U_j$ *and* $v_i := \chi_i^{\mathbf{M}}(z^{\mathbf{M}}(u))$,

$$\|T_{ij}^{\mathbf{M}}(v_i) - \chi_j^{\mathbf{M}}(z^{\mathbf{M}}(u))\| = 0,$$

and departures from this ideal arise only through observational and encoder noise (captured in E_{link}*) and chart distortion (Equation (***??***)). For general* k*, we define* $\varepsilon_{\text{trans}}^{\mathbf{M}}(k)$ *as a non-decreasing function of* k *whose leading-order term is a linear accumulation of the one-hop distortion contributions; Section 6 makes this dependence explicit and shows that under the bounded-distortion regime of Definition* **??**, $\varepsilon_{\text{trans}}^{\mathbf{M}}(k)$ *is first-order in the distortion profile* $\varepsilon_{\text{dist}}^{\mathbf{M}}$.

By construction, $G_S(\mathbf{M})$ is a TSR of S in the sense of Definition **??**. Intuitively, G_S forgets the manifold and keeps only the chart-transition structure, turning the Globe back into a translation graph whose edges represent coordinate changes between overlapping charts.

5.3 Constructing G_S on morphisms

We now define G_S on morphisms by discarding the manifold map and retaining only the feature-space reparameterizations.

Definition 20 (Functor G_S on morphisms)**.** *Let $H : \mathbf{M} \to \mathbf{M}'$ be a morphism in* $\mathsf{SamGeom}^0(S)$ *in the sense of Definition* **??**, *with*

$$H = (h,\ R), \qquad h : \mathcal{M}^{\mathbf{M}} \to \mathcal{M}^{\mathbf{M}'}, \quad R = \{R_i : V_i^{\mathbf{M}} \to V_i^{\mathbf{M}'}\}_{i \in \mathcal{I}}.$$

We define

$$G_S(H) := R : G_S(\mathbf{M}) \to G_S(\mathbf{M}'),$$

with tolerances $(\eta_{\text{enc}}, \eta_{\text{trans}})$ to be specified below. That is, the morphism $G_S(H)$ is simply the family of feature-space reparameterizations from H, now viewed as a TSR morphism in the sense of Definition **??**.

The following lemma records that this definition is well-posed.

Lemma 1 (Well-definedness of G_S on morphisms)**.** *Let $H = (h, R) : \mathbf{M} \to \mathbf{M}'$ be a morphism in* $\mathsf{SamGeom}^0(S)$ *with encoder tolerance η_{enc} and chart tolerance η_{chart} as in Equations* (**??**) *and* (**??**)*. Then $G_S(H) = R$ is a morphism in* $\mathsf{SamTrans}^0(S)$ *from $\mathbf{T}^{\mathbf{M}} = G_S(\mathbf{M})$ to $\mathbf{T}^{\mathbf{M}'} = G_S(\mathbf{M}')$ with encoder tolerance η_{enc} and translator tolerance η_{trans} bounded by a first-order function of*

$$\eta_{\text{chart}}, \quad \varepsilon_{\text{dist}}^{\mathbf{M}}(\cdot), \quad \varepsilon_{\text{dist}}^{\mathbf{M}'}(\cdot).$$

Proof sketch. *Encoder compatibility for $G_S(H)$ is exactly Equation* (**??**)*, so the encoder tolerance is preserved.*

For translators, recall that $T_{ij}^{\mathbf{M}} = \chi_j^{\mathbf{M}} \circ \psi_i^{\mathbf{M}}$ and $T_{ij}^{\mathbf{M}'} = \chi_j^{\mathbf{M}'} \circ \psi_i^{\mathbf{M}'}$ on in-regime domains. Given an in-regime latent entity u and the corresponding chart coordinates $v_i := \chi_i^{\mathbf{M}}(z^{\mathbf{M}}(u))$ and $v_i' := \chi_i^{\mathbf{M}'}(z^{\mathbf{M}'}(u))$, the chart compatibility condition for H can be written as

$$\|R_i(v_i) - R_h^i(v_i)\| \le \eta_{\text{chart}}, \qquad R_h^i := \chi_i^{\mathbf{M}'} \circ h \circ \psi_i^{\mathbf{M}}.$$

A similar relation holds for R_j. The translator compatibility condition we must check for $G_S(H)$ is

$$\|R_j(T_{ij}^{\mathbf{M}}(v_i)) - T_{ij}^{\mathbf{M}'}(R_i(v_i))\| \le \eta_{\text{trans}}.$$

By inserting and subtracting the h-induced reparameterizations R_h^i and R_h^j, and then using the distortion bounds (**??**) *for both $\mathbf{M}$ and $\mathbf{M}'$, one can show that η_{trans} is controlled by a first-order function of η_{chart} and the distortion profiles $\varepsilon_{\text{dist}}^{\mathbf{M}}$, $\varepsilon_{\text{dist}}^{\mathbf{M}'}$. A detailed proof follows from the distortion bounds and chart-compatibility conditions defined above.*

5.4 Functoriality and gauge stability

We now state the functoriality of G_S and its stability with respect to the error gauges of Definition **??**.

Proposition 2 (Functoriality and gauge stability of G_S)**.** *The assignment G_S defines a functor G_S :* $\mathsf{SamGeom}^0(S) \to \mathsf{SamTrans}^0(S)$ *with the following properties.*

(i) ***Functoriality.*** *For every $\mathbf{M} \in$* $\mathsf{SamGeom}^0(S)$*, $G_S(\mathrm{id}_{\mathbf{M}})$ is the identity morphism on $G_S(\mathbf{M})$, and for any composable morphisms $H_1 : \mathbf{M}_0 \to \mathbf{M}_1$ and $H_2 : \mathbf{M}_1 \to \mathbf{M}_2$ in* $\mathsf{SamGeom}^0(S)$*,*

$$G_S(H_2 \circ H_1) = G_S(H_2) \circ G_S(H_1)$$

as morphisms in $\mathsf{SamTrans}^0(S)$*.*

(ii) ***Gauge stability.*** *There exists a first-order slack function*

$$C_G(\cdot,\cdot,\cdot) : \mathbb{R}_{\geq 0} \times \mathcal{E} \times \mathcal{E} \to \mathbb{R}_{\geq 0}$$

(where $\mathcal{E}$ *denotes the space of distortion profiles) such that for any* $\mathbf{M}, \mathbf{M}' \in \mathsf{SamGeom}^0(S)$*,*

$$\Delta_{\mathsf{SamTrans}}(G_S(\mathbf{M}), G_S(\mathbf{M}')) \leq C_G\Big(\Delta_{\mathsf{SamGeom}}(\mathbf{M}, \mathbf{M}'), \varepsilon^{\mathbf{M}}_{\mathrm{dist}}(\cdot), \varepsilon^{\mathbf{M}'}_{\mathrm{dist}}(\cdot)\Big), \tag{11}$$

with $C_G(0, \varepsilon, \varepsilon)$ *first-order in* ε*.*

Proof sketch. Functoriality in (i) is immediate from the definition $G_S(H) = R$: the identity morphism in $\mathsf{SamGeom}^0(S)$ has $h = \mathrm{id}$ and $R_i = \mathrm{id}$ for all i, so $G_S(\mathrm{id}_{\mathbf{M}})$ is the identity morphism on $G_S(\mathbf{M})$. Composition is preserved because $(h_2 \circ h_1, \{R_i^{(2)} \circ R_i^{(1)}\})$ maps under G_S to $\{R_i^{(2)} \circ R_i^{(1)}\}$, which is exactly the composite of $G_S(H_1)$ and $G_S(H_2)$ in $\mathsf{SamTrans}^0(S)$.

For (ii), consider two MSRs $\mathbf{M}, \mathbf{M}' \in \mathsf{SamGeom}^0(S)$ and the corresponding TSRs $\mathbf{T}^{\mathbf{M}}$ and $\mathbf{T}^{\mathbf{M}'}$. The gauge $\Delta_{\mathsf{SamTrans}}(\mathbf{T}^{\mathbf{M}}, \mathbf{T}^{\mathbf{M}'})$ from Definition **??** has two terms: the difference between the drift profiles $\varepsilon^{\mathbf{M}}_{\mathrm{trans}}$ and $\varepsilon^{\mathbf{M}'}_{\mathrm{trans}}$, and the supremum of translator disagreements on ideal latent entities. Both are controlled by the chart structure and the geodesic distances $d_{g^{\mathbf{M}}}$, $d_{g^{\mathbf{M}'}}$ via Equation (**??**). Using the definition of $\Delta_{\mathsf{SamGeom}}$, a chaining argument over benign paths, and the bounded-distortion regime, one can show that the right-hand side of (**??**) holds with C_G first-order in each of its arguments. Full details are deferred to Section 6.

5.5 Interpretation and limitations

The functor G_S provides a canonical way to view a well-parameterized MSR $\mathbf{M}$ as a TSR $\mathbf{T}^{\mathbf{M}}$ whose translators implement chart transitions. This has two conceptual implications.

First, G_S shows that the translator view and the manifold view are not competing descriptions but dual coordinatizations of the same semantic sameness structure S. A practitioner who starts from a geometry-first design (an MSR) can always "unfold" it into a translator graph that lies in $\mathsf{SamTrans}^0(S)$, and error gauges are preserved up to first-order slack in the distortion profiles.

Second, G_S makes explicit that the stability of the translator view depends not only on the difference between manifolds $\Delta_{\mathsf{SamGeom}}(\mathbf{M}, \mathbf{M}')$ but also on the intrinsic distortion of each manifold, as reflected in the dependence of C_G on $\varepsilon^{\mathbf{M}}_{\mathrm{dist}}$ and $\varepsilon^{\mathbf{M}'}_{\mathrm{dist}}$ in (**??**). This is the geometric analogue of the translator asymmetry and chart slack terms that appear on the TSR side.

Together with the construction of F_S in Section **??**, the functor G_S completes the bridge between the translator and manifold design spaces. Section 6 uses these functors to prove the ε-equivalence theorem (Theorem **??**) and, ultimately, the Error Budget Transfer Theorem, where manifold geometry and error budgets are kept separate but tightly coupled.

6 Error-budget transfer and detector-level equivalence

Sections **??** and **??** constructed functors

$$F_S : \mathsf{SamTrans}^0(S) \to \mathsf{SamGeom}^0(S), \qquad G_S : \mathsf{SamGeom}^0(S) \to \mathsf{SamTrans}^0(S),$$

between the well-behaved subcategories of translation-based and manifold-based realizations of a fixed semantic sameness structure S. The structural Theorem **??** showed that these functors form an ε-equivalence of categories in the sense of Definition **??**: they preserve realizations up to first-order slack in the error gauges $\Delta_{\mathsf{SamTrans}}$ and $\Delta_{\mathsf{SamGeom}}$.

In this section we lift that structural equivalence to the detector level. We assume that realizations are equipped with sameness detectors in the sense of Definition **??**, and we show that *error budgets and Davis-style correctness bounds are preserved in both directions up to first-order slack.* The key notion is that of *aligned detectors* on structurally linked realizations.

6.1 Aligned detectors across realizations

Fix a semantic sameness structure S and an in-regime latent subset $\Omega^{\text{ideal}} \subseteq I$ as in Definition **??**. Let $\mathbf{T}$ be a TSR and $\mathbf{M}$ an MSR of S that are linked either by chartability ($\mathbf{M} = F_S(\mathbf{T})$) or by geometric realization ($\mathbf{T} = G_S(\mathbf{M})$). Recall that for each $u \in \Omega^{\text{ideal}}$ and any modality i with $\pi_i(u)$ defined we write

$$v_i^{\mathbf{T}}(u) := \varphi_i^{\mathbf{T}}(\pi_i(u)) \in V_i^{\mathbf{T}}, \qquad z^{\mathbf{M}}(u) := \psi_i^{\mathbf{M}}(v_i^{\mathbf{T}}(u)) \in \mathcal{M}^{\mathbf{M}},$$

which is independent of i on the in-regime set by latent consistency.

Definition 21 (Aligned detectors)**.** *Let* $\mathbf{T}$ *and* $\mathbf{M}$ *be as above and let*

$$D^{\mathbf{T}} = (\mathbf{T}, \Sigma_{\mathbf{T}}, A_{\mathbf{T}}, E^{\mathbf{T}}_{\text{geom}}, E^{\mathbf{T}}_{\text{link}}, \xi^{\mathbf{T}}, \zeta^{\mathbf{T}}, \delta^{\mathbf{T}}_{\text{indep}}, \tau^{\mathbf{T}}_{\text{vac}})$$

and

$$D^{\mathbf{M}} = (\mathbf{M}, \Sigma_{\mathbf{M}}, A_{\mathbf{M}}, E^{\mathbf{M}}_{\text{geom}}, E^{\mathbf{M}}_{\text{link}}, \xi^{\mathbf{M}}, \zeta^{\mathbf{M}}, \delta^{\mathbf{M}}_{\text{indep}}, \tau^{\mathbf{M}}_{\text{vac}})$$

be sameness detectors on $\mathbf{T}$ *and* $\mathbf{M}$ *in the sense of Definition* **??**. *We say that* $(D^{\mathbf{T}}, D^{\mathbf{M}})$ *are* aligned detectors *for* S *if:*

(1) ***Statistic-level alignment.*** *There exists* $\eta_{\text{stat}} \geq 0$ *such that for all* $u \in \Omega^{\text{ideal}}$ *and any modality* i *with* $\pi_i(u)$ *defined,*

$$\left|\Sigma_{\mathbf{T}}(v_i^{\mathbf{T}}(u)) - \Sigma_{\mathbf{M}}(z^{\mathbf{M}}(u))\right| \leq \eta_{\text{stat}}. \tag{12}$$

(2) ***Decision-region compatibility.*** *The decision rules* $A_{\mathbf{T}}$ *and* $A_{\mathbf{M}}$ *use compatible high-risk thresholds and abstention regions in the statistic coordinate: there exists* $\eta_{\text{thresh}} \geq 0$ *such that the high-risk events* $H^{\mathbf{T}}, H^{\mathbf{M}}$ *obey*

$$H^{\mathbf{T}} \triangle H^{\mathbf{M}} \subseteq \{x : |\Sigma_{\mathbf{T}}(x) - \Sigma_{\mathbf{M}}(x)| \leq \eta_{\text{thresh}}\}$$

on non-abstaining, in-regime inputs, where $\triangle$ *denotes symmetric difference.*

(3) ***Shared calibration and abstention.*** *The calibration and abstention components are either shared or differ only by small slacks: there exist* $\eta_{\text{cal}}, \eta_{\text{abst}} \geq 0$ *such that*

$$|\xi^{\mathbf{T}} - \xi^{\mathbf{M}}| \leq \eta_{\text{cal}}, \qquad |\zeta^{\mathbf{T}} - \zeta^{\mathbf{M}}| \leq \eta_{\text{abst}}.$$

We collect the alignment slacks into

$$\alpha_{\text{align}} := (\eta_{\text{stat}}, \eta_{\text{thresh}}, \eta_{\text{cal}}, \eta_{\text{abst}}).$$

In practice, aligned detectors arise by design: given a TSR detector $D^{\mathbf{T}}$ built from features on $\{V_i^{\mathbf{T}}\}$, the associated manifold detector $D^{\mathbf{M}}$ typically uses the same parameterization of the statistic on $\mathcal{M}^{\mathbf{M}}$ via $z^{\mathbf{M}}(u)$ and reuses the same calibration and abstention policies. In such cases the alignment slacks are dominated by the structural chartability parameters and by approximation error in the statistics.

6.2 Error-budget equivalence along F_S

We first analyze the detector-level effect of the forward functor $F_S : \mathsf{SamTrans}^0(S) \to \mathsf{SamGeom}^0(S)$. Recall that for a smoothly chartable TSR $\mathbf{T}$ the functor constructs a translator metric quotient and, under the Riemannian realizability assumption, a Davis-style MSR $\mathbf{M} = F_S(\mathbf{T})$ whose chart structure is compatible with the translators. The chartability parameters

$$\varepsilon^{\mathbf{T}}_{\text{trans}}, \quad \varepsilon^{\mathbf{M}}_{\text{dist}}, \quad \delta_{\text{chart}}(\mathbf{T}, \mathbf{M}), \quad \delta_{\text{asymm}}(\mathbf{T})$$

control translator drift, metric distortion, chart overlap slack, and translator asymmetry, respectively.

We now show that for aligned detectors $(D^{\mathbf{T}}, D^{\mathbf{M}})$ on $(\mathbf{T}, \mathbf{M})$ the error budgets and statistic separation are equivalent up to first-order slack in these quantities and in the alignment slacks α_{align}.

Theorem 2 (Error-budget transfer along F_S)**.** *Fix a sameness structure S and a smoothly chartable TSR $\mathbf{T} \in \mathsf{SamTrans}^0(S)$ with associated MSR $\mathbf{M} = F_S(\mathbf{T})$. Let $D^{\mathbf{T}}$ and $D^{\mathbf{M}}$ be aligned detectors in the sense of Definition* **??**. *Write $E^{\mathbf{R}}_{\text{geom}}$, $E^{\mathbf{R}}_{\text{link}}$, $\xi^{\mathbf{R}}$, $\zeta^{\mathbf{R}}$, $\delta^{\mathbf{R}}_{\text{indep}}$, $\Delta^{\mathbf{R}}_S$ for the geometric error, linkage error, calibration error, abstention failure, independence slack, and statistic separation of $D^{\mathbf{R}}$ for $\mathbf{R} \in \{\mathbf{T}, \mathbf{M}\}$, with $\Delta^{\mathbf{R}}_S$ defined as in the Davis framework for the induced statistic.:contentReference[oaicite:0]index=0*

Then there exist continuous first-order *slack functions*

$$\Gamma^F_{\text{geom}},\ \Gamma^F_{\text{link}},\ \Gamma^F_{\text{sep}},\ \Gamma^F_{\text{det}}$$

such that:

$$|E^{\mathbf{T}}_{\text{geom}} - E^{\mathbf{M}}_{\text{geom}}| \le \Gamma^F_{\text{geom}}(\varepsilon^{\mathbf{T}}_{\text{trans}}, \varepsilon^{\mathbf{M}}_{\text{dist}}, \delta_{\text{chart}}(\mathbf{T}, \mathbf{M}), \delta_{\text{asymm}}(\mathbf{T}), \alpha_{\text{align}}), \tag{13}$$

$$|E^{\mathbf{T}}_{\text{link}} - E^{\mathbf{M}}_{\text{link}}| \le \Gamma^F_{\text{link}}(\varepsilon^{\mathbf{T}}_{\text{trans}}, \varepsilon^{\mathbf{M}}_{\text{dist}}, \delta_{\text{chart}}(\mathbf{T}, \mathbf{M}), \delta_{\text{asymm}}(\mathbf{T}), \alpha_{\text{align}}), \tag{14}$$

$$|\Delta^{\mathbf{T}}_S - \Delta^{\mathbf{M}}_S| \le \Gamma^F_{\text{sep}}(\varepsilon^{\mathbf{T}}_{\text{trans}}, \varepsilon^{\mathbf{M}}_{\text{dist}}, \delta_{\text{chart}}(\mathbf{T}, \mathbf{M}), \delta_{\text{asymm}}(\mathbf{T}), \alpha_{\text{align}}), \tag{15}$$

and $\Gamma^F_{\text{geom}}(0) = \Gamma^F_{\text{link}}(0) = \Gamma^F_{\text{sep}}(0) = 0$ with linear (or sublinear) behavior near 0.

Moreover, let $LB(\cdot)$ denote any Davis-style lower bound on posterior correctness that is monotone non-decreasing in Δ_S and monotone non-increasing in each of $E_{\text{geom}}, E_{\text{link}}, \xi, \zeta,$ and δ_{indep}; for example, the Cantelli-based bound of Theorem 2 in the Davis-manifold paper.:contentReference[oaicite:1]index=1 Then

$$|LB(D^{\mathbf{T}}) - LB(D^{\mathbf{M}})| \le \Gamma^F_{\text{det}}(\varepsilon^{\mathbf{T}}_{\text{trans}}, \varepsilon^{\mathbf{M}}_{\text{dist}}, \delta_{\text{chart}}(\mathbf{T}, \mathbf{M}), \delta_{\text{asymm}}(\mathbf{T}), \alpha_{\text{align}}), \tag{16}$$

with Γ^F_{det} first-order in its arguments.

Proof sketch. The structural ε-equivalence of F_S and the chartability conditions imply that for in-regime latent entities u the translator graph and manifold geometry induce path families and distances that agree up to first-order slack in the parameters $(\varepsilon^{\mathbf{T}}_{\text{trans}}, \varepsilon^{\mathbf{M}}_{\text{dist}}, \delta_{\text{chart}}, \delta_{\text{asymm}})$. Aligned statistics then ensure that the detectors see essentially the same underlying semantic separation along corresponding paths, and that decision regions differ by at most α_{align} in the statistic coordinate.

By definition, $E^{\mathbf{R}}_{\text{geom}}$ and $E^{\mathbf{R}}_{\text{link}}$ are probabilities of structural and linkage failures on ideal data. The chartability bounds give an inclusion between the corresponding failure events for $\mathbf{T}$ and $\mathbf{M}$, up to sets of probability controlled by the small parameters above, yielding the absolute-difference bounds (**??**)–(**??**). A similar argument applied to the Davis-style separation parameter Δ_S gives (**??**). Finally, monotonicity of $LB(\cdot)$ in its arguments implies that the induced change in the lower bound is at most first-order in the changes of $(E_{\text{geom}}, E_{\text{link}}, \xi, \zeta, \delta_{\text{indep}}, \Delta_S)$, giving (**??**). Full details, including explicit expressions for the slack functions, are deferred to Section **??**.

6.3 Error-budget equivalence along G_S

We now prove the dual result for the geometric-to-translator functor $G_S : \mathsf{SamGeom}^0(S) \to \mathsf{SamTrans}^0(S)$. Here the roles of translator drift and metric distortion are swapped: $\varepsilon_{\text{dist}}$ is given as part of the MSR, while $\varepsilon_{\text{trans}}$ and δ_{asymm} arise from the induced translators $T^{\mathbf{M}}_{ij} := \chi^{\mathbf{M}}_j \circ \psi^{\mathbf{M}}_i$.

Theorem 3 (Error-budget transfer along G_S)**.** *Fix a sameness structure S and a geometrically well-parameterized MSR $\mathbf{M} \in \mathsf{SamGeom}^0(S)$ in the sense of Definition* **??**. *Let $\mathbf{T} := G_S(\mathbf{M})$ be the induced TSR, and let $D^{\mathbf{M}}$ and $D^{\mathbf{T}}$ be aligned detectors as in Definition* **??**. *Then there exist first-order slack functions*

$$\Gamma^G_{\text{geom}},\ \Gamma^G_{\text{link}},\ \Gamma^G_{\text{sep}},\ \Gamma^G_{\text{det}}$$

such that:

$$|E^{\mathbf{M}}_{\text{geom}} - E^{\mathbf{T}}_{\text{geom}}| \le \Gamma^G_{\text{geom}}(\varepsilon^{\mathbf{M}}_{\text{dist}}, \varepsilon^{\mathbf{T}}_{\text{trans}}, \delta_{\text{chart}}(\mathbf{T},\mathbf{M}), \delta_{\text{asymm}}(\mathbf{T}), \alpha_{\text{align}}), \tag{17}$$

$$|E^{\mathbf{M}}_{\text{link}} - E^{\mathbf{T}}_{\text{link}}| \le \Gamma^G_{\text{link}}(\varepsilon^{\mathbf{M}}_{\text{dist}}, \varepsilon^{\mathbf{T}}_{\text{trans}}, \delta_{\text{chart}}(\mathbf{T},\mathbf{M}), \delta_{\text{asymm}}(\mathbf{T}), \alpha_{\text{align}}), \tag{18}$$

$$|\Delta^{\mathbf{M}}_S - \Delta^{\mathbf{T}}_S| \le \Gamma^G_{\text{sep}}(\varepsilon^{\mathbf{M}}_{\text{dist}}, \varepsilon^{\mathbf{T}}_{\text{trans}}, \delta_{\text{chart}}(\mathbf{T},\mathbf{M}), \delta_{\text{asymm}}(\mathbf{T}), \alpha_{\text{align}}), \tag{19}$$

and, for any Davis-style lower bound $LB(\cdot)$ as in Theorem **??**,

$$|LB(D^{\mathbf{M}}) - LB(D^{\mathbf{T}})| \le \Gamma^G_{\text{det}}(\varepsilon^{\mathbf{M}}_{\text{dist}}, \varepsilon^{\mathbf{T}}_{\text{trans}}, \delta_{\text{chart}}(\mathbf{T},\mathbf{M}), \delta_{\text{asymm}}(\mathbf{T}), \alpha_{\text{align}}), \tag{20}$$

with all Γ^G first-order at the origin.

Proof sketch. The proof mirrors that of Theorem **??**. The construction of G_S expresses each translator $T^{\mathbf{M}}_{ij} = \chi^{\mathbf{M}}_j \circ \psi^{\mathbf{M}}_i$ as a chart-transition map on $\mathcal{M}^{\mathbf{M}}$, so that translator drift and asymmetry on $\mathbf{T}$ are controlled by the metric distortion and chart geometry of $\mathbf{M}$. Geometric well-parameterization ensures that these quantities are small on the in-regime subset. Aligned detectors then again see (up to α_{align}) the same underlying semantic separation, and the same argument as before bounds differences in E_{geom}, E_{link}, Δ_S, and $LB(\cdot)$ by first-order functions of the structural and alignment parameters. Details are deferred to Section **??**.

6.4 Detector-level ε-equivalence

Combining the structural ε-equivalence of Theorem **??** with the detector-level transfer results above yields an ε-equivalence *of detectors* on $\mathsf{SamTrans}^0(S)$ and $\mathsf{SamGeom}^0(S)$.

Corollary 1 (Round-trip detector-level equivalence)**.** *Fix S and consider a smoothly chartable TSR $\mathbf{T} \in \mathsf{SamTrans}^0(S)$ with $\mathbf{M} = F_S(\mathbf{T})$, and a geometrically well-parameterized MSR $\mathbf{M}' \in \mathsf{SamGeom}^0(S)$ with $\mathbf{T}' = G_S(\mathbf{M}')$. Let $D^{\mathbf{T}}$, $D^{\mathbf{M}}$, $D^{\mathbf{M}'}$, $D^{\mathbf{T}'}$ be aligned detectors on these realizations.*

Then there exist first-order slack functions C^F_{det} and C^G_{det} such that

$$|LB(D^{\mathbf{T}}) - LB(D^{F_S(\mathbf{T})})| \le C^F_{\mathrm{det}}(\varepsilon^{\mathbf{T}}_{\mathrm{trans}}, \varepsilon^{F_S(\mathbf{T})}_{\mathrm{dist}}, \delta_{\mathrm{chart}}(\mathbf{T}, F_S(\mathbf{T})), \delta_{\mathrm{asymm}}(\mathbf{T})), \tag{21}$$

$$|LB(D^{\mathbf{M}'}) - LB(D^{G_S(\mathbf{M}')})| \le C^G_{\mathrm{det}}(\varepsilon^{\mathbf{M}'}_{\mathrm{dist}}, \varepsilon^{G_S(\mathbf{M}')}_{\mathrm{trans}}, \delta_{\mathrm{chart}}(G_S(\mathbf{M}'), \mathbf{M}'), \delta_{\mathrm{asymm}}(G_S(\mathbf{M}'))), \tag{22}$$

with $C^F_{\mathrm{det}}(0) = C^G_{\mathrm{det}}(0) = 0$ and first- order dependence near the origin.

In particular, up to first-order slack in the structural parameters $(\varepsilon_{\mathrm{trans}}, \varepsilon_{\mathrm{dist}}, \delta_{\mathrm{chart}}, \delta_{\mathrm{asymm}})$, Davis-style correctness bounds that hold for a manifold-based detector can be transported to a translator-based detector realizing the same semantic sameness structure, and conversely.

Conceptually, this corollary is the detector-level analogue of the categorical ε-equivalence in Theorem **??**: *within the smoothly chartable / geometrically well-parameterized regime, working on the "globe"* $\mathbf{M}$ *or on a compatible "atlas of flat maps"* $\mathbf{T}$ *leads to equivalent guarantees up to first-order slack in the geometry and chartability parameters.* This justifies deriving Davis-style theorems on manifolds and then instantiating them on real systems implemented as translator graphs, as long as chartability audits support the small-parameter regime.

7 Detailed proofs and structural slack bounds

This section collects explicit formulas for the slack functions Γ^F_{real}, Γ^G_{real}, C_F, and C_G, and proves the structural results announced in Sections **??**–**??**. The central statement is a gauge-stability theorem for the functors F_S and G_S; the error-budget transfer results then appear as corollaries.

Throughout we fix a sameness structure S and work inside the well–behaved subcategories $\mathsf{SamTrans}^0(S) \subset \mathsf{SamTrans}(S)$ and $\mathsf{SamGeom}^0(S) \subset \mathsf{SamGeom}(S)$ defined in Sections **??**–**??**. We assume the realizability hypothesis for translator metric quotients from Section **??**.

7.1 Slack functions and realizability

We first make explicit the "realizability" slacks that measure how far a translator-based or manifold-based realization is from an ideal, infinitesimally exact geometry, and then define the gauge-stability slack functions C_F and C_G.

Definition 22 (Realizability slacks)**.** *Let $\mathbf{T} \in \mathsf{SamTrans}^0(S)$ be a smoothly chartable TSR with drift profile $\varepsilon^{\mathbf{T}}_{\mathrm{trans}}(k)$ and asymmetry parameter $\delta_{\mathrm{asymm}}(\mathbf{T})$ (Definition **??**). Let*

$$\tilde{\mathcal{M}}^{\mathbf{T}},\ \tilde{d}_{\mathbf{T}}$$

*be its translator metric quotient as in Definition **??**, and let $(\mathcal{M}^{\mathbf{T}}, g^{\mathbf{T}})$ be a Riemannian manifold and quasi-isometry $q^{\mathbf{T}} : \tilde{\mathcal{M}}^{\mathbf{T}} \to \mathcal{M}^{\mathbf{T}}$ given by the Riemannian realizability assumption.*

We define the realizability slack *of* $\mathbf{T}$ *by*

$$\Gamma^F_{\mathrm{real}}(\mathbf{T}) := \sup_{[v],[w]\in\tilde{\mathcal{M}}^{\mathbf{T}}} \left| \frac{\tilde{d}_{\mathbf{T}}([v],[w])}{d_{g^{\mathbf{T}}}(q^{\mathbf{T}}([v]), q^{\mathbf{T}}([w]))} - 1 \right|. \tag{23}$$

By realizability, $\Gamma^F_{\mathrm{real}}(\mathbf{T}) < \infty$ *and, for chartable TSRs, it is controlled by the intrinsic translator quantities:*

$$\Gamma^F_{\mathrm{real}}(\mathbf{T}) \lesssim a_1 \sup_{k\geq 1} \varepsilon^{\mathbf{T}}_{\mathrm{trans}}(k) + a_2\, \delta_{\mathrm{asymm}}(\mathbf{T}),$$

for constants a_1, a_2 *depending only on the ambient architecture and the chosen path horizon.*

Dually, for $\mathbf{M} \in \mathsf{SamGeom}^0(S)$ *with chart family* $\{\chi^{\mathbf{M}}_i\}$ *and distortion profile* $\varepsilon^{\mathbf{M}}_{\mathrm{dist}}(\cdot)$*, we define*

$$\Gamma^G_{\mathrm{real}}(\mathbf{M}) := \sup_{z,z'\in\mathcal{M}^{\mathbf{M}}} \left| \frac{\delta_i(z,z')}{d_{g^{\mathbf{M}}}(z,z')} - 1 \right|, \tag{24}$$

where $\delta_i(z,z') := \|\chi^{\mathbf{M}}_i(z) - \chi^{\mathbf{M}}_i(z')\|_{V_i}$ *on any chart* U_i *containing* z, z'*. By Definition* **??***, this is bounded by the distortion profile:*

$$\Gamma^G_{\mathrm{real}}(\mathbf{M}) \lesssim \sup_{\ell>0} \varepsilon^{\mathbf{M}}_{\mathrm{dist}}(\ell).$$

Both slacks are first-order*: whenever the intrinsic translator errors (for* $\mathbf{T}$*) or chart distortions (for* $\mathbf{M}$*) tend to zero, the corresponding* Γ^F_{real} *or* Γ^G_{real} *tends to zero.*

We can now define the gauge-level slack functions that appear in the functorial stability theorem.

Definition 23 (Gauge-stability slack functions)**.** *Let* $\Delta_{\mathsf{SamTrans}}$ *and* $\Delta_{\mathsf{SamGeom}}$ *be the error gauges on* $\mathsf{SamTrans}^0(S)$ *and* $\mathsf{SamGeom}^0(S)$ *from Definition* **??***. We define:*

$$C_F(\mathbf{T},\mathbf{T}') := \alpha_F\, \Delta_{\mathsf{SamTrans}}(\mathbf{T},\mathbf{T}') + \beta_F(\Gamma^F_{\mathrm{real}}(\mathbf{T}) + \Gamma^F_{\mathrm{real}}(\mathbf{T}')), \tag{25}$$

$$C_G(\mathbf{M},\mathbf{M}') := \alpha_G\, \Delta_{\mathsf{SamGeom}}(\mathbf{M},\mathbf{M}') + \beta_G(\Gamma^G_{\mathrm{real}}(\mathbf{M}) + \Gamma^G_{\mathrm{real}}(\mathbf{M}')), \tag{26}$$

for fixed constants $\alpha_F, \beta_F, \alpha_G, \beta_G > 0$ *that depend only on the ambient architectures and on the choice of path horizon and operational region, not on particular realizations* S*,* $\mathbf{T}$*, or* $\mathbf{M}$*.*

By construction these are first-order *in the gauges and slacks: if*

$$\Delta_{\mathsf{SamTrans}}(\mathbf{T},\mathbf{T}') \to 0, \quad \Gamma^F_{\mathrm{real}}(\mathbf{T}), \Gamma^F_{\mathrm{real}}(\mathbf{T}') \to 0,$$

then $C_F(\mathbf{T},\mathbf{T}') \to 0$*, and likewise for* C_G *with* $\Delta_{\mathsf{SamGeom}}$ *and* Γ^G_{real}*.*

7.2 Main structural theorem: gauge stability of F_S and G_S

We now prove the main structural result: the functors F_S and G_S are Lipschitz with respect to the error gauges, with slack functions C_F and C_G from Definition **??**. This theorem is the "machine" that all error-budget transfer statements in Section **??** rest on.

Theorem 4 (Gauge stability of F_S and G_S)**.** *Fix a sameness structure* S *and realizability constants as above.*

1. ***Functorial stability of*** F_S***.*** *For any* $\mathbf{T}, \mathbf{T}' \in \mathsf{SamTrans}^0(S)$*, let* $\mathbf{M} := F_S(\mathbf{T})$ *and* $\mathbf{M}' := F_S(\mathbf{T}')$ *be the manifold-based realizations constructed via the translator metric quotient and Riemannian realizability. Then*

$$\Delta_{\mathsf{SamGeom}}(\mathbf{M}, \mathbf{M}') \ \leq \ C_F(\mathbf{T}, \mathbf{T}'), \tag{27}$$

with C_F *as in* (**??**)*. In particular, if* $\Delta_{\mathsf{SamTrans}}(\mathbf{T}, \mathbf{T}') \to 0$ *and the realizability slacks of* $\mathbf{T}, \mathbf{T}'$ *tend to zero, then* $\Delta_{\mathsf{SamGeom}}(\mathbf{M}, \mathbf{M}') \to 0$*.*

2. ***Functorial stability of*** G_S***.*** *For any* $\mathbf{M}, \mathbf{M}' \in \mathsf{SamGeom}^0(S)$*, let* $\mathbf{T} := G_S(\mathbf{M})$ *and* $\mathbf{T}' := G_S(\mathbf{M}')$ *be the translator-based realizations constructed in Section* **??***. Then*

$$\Delta_{\mathsf{SamTrans}}(\mathbf{T}, \mathbf{T}') \ \leq \ C_G(\mathbf{M}, \mathbf{M}'), \tag{28}$$

with C_G *as in* (**??**)*. In particular, if* $\Delta_{\mathsf{SamGeom}}(\mathbf{M}, \mathbf{M}') \to 0$ *and the realizability slacks of* $\mathbf{M}, \mathbf{M}'$ *tend to zero, then* $\Delta_{\mathsf{SamTrans}}(\mathbf{T}, \mathbf{T}') \to 0$*.*

Proof sketch. We sketch the argument for F_S; the case of G_S is dual with translator/metric roles switched.

Let $\mathbf{T}, \mathbf{T}' \in \mathsf{SamTrans}^0(S)$, and let $\mathbf{M} = F_S(\mathbf{T})$, $\mathbf{M}' = F_S(\mathbf{T}')$ be obtained by: (i) forming translator metric quotients $(\tilde{\mathcal{M}}^{\mathbf{T}}, \tilde{d}_{\mathbf{T}})$ and $(\tilde{\mathcal{M}}^{\mathbf{T}'}, \tilde{d}_{\mathbf{T}'})$ as in Definition **??**; (ii) invoking realizability to obtain quasi-isometries $q^{\mathbf{T}} : \tilde{\mathcal{M}}^{\mathbf{T}} \to \mathcal{M}^{\mathbf{T}}$, $q^{\mathbf{T}'} : \tilde{\mathcal{M}}^{\mathbf{T}'} \to \mathcal{M}^{\mathbf{T}'}$; and (iii) defining charts $\psi_i^{\mathbf{T}}, \psi_i^{\mathbf{T}'}$ by composing the canonical quotient maps with $q^{\mathbf{T}}, q^{\mathbf{T}'}$.

Profile part. The distortion profiles $\varepsilon_{\text{dist}}^{\mathbf{M}}$ and $\varepsilon_{\text{dist}}^{\mathbf{M}'}$ compare chart distances δ_i to geodesic distances $d_{g^{\mathbf{T}}}$ and $d_{g^{\mathbf{T}'}}$. For any pair of latent entities u, u' in the in-regime set, we compare the corresponding manifold distances via the chain

$$\left| d_{g^{\mathbf{T}}}(z(u), z(u')) - d_{g^{\mathbf{T}'}}(z'(u), z'(u')) \right| \leq A + B + C,$$

where

$$\begin{aligned} A &= \left| d_{g^{\mathbf{T}}}(z, z') - \tilde{d}_{\mathbf{T}}([v], [v']) \right|, \\ B &= \left| \tilde{d}_{\mathbf{T}}([v], [v']) - \tilde{d}_{\mathbf{T}'}([v'], [v'']) \right|, \\ C &= \left| \tilde{d}_{\mathbf{T}'}([v'], [v'']) - d_{g^{\mathbf{T}'}}(z', z'') \right|. \end{aligned}$$

Here $[v], [v'], [v'']$ are the equivalence classes of feature vectors encoding u, u' in the translator quotients; the exact indexing is not important.

Terms A and C are controlled by the realizability slacks: by definition of Γ_{real}^F,

$$A \ \lesssim \ \Gamma_{\text{real}}^F(\mathbf{T}), \qquad C \ \lesssim \ \Gamma_{\text{real}}^F(\mathbf{T}').$$

Term B is controlled by the translator gauge $\Delta_{\mathsf{SamTrans}}(\mathbf{T}, \mathbf{T}')$: the quotient metric $\tilde{d}_{\mathbf{T}}$ is defined as an infimum over path costs built from the translators $T_{ij}^{\mathbf{T}}$, and likewise for $\tilde{d}_{\mathbf{T}'}$ and $T_{ij}^{\mathbf{T}'}$. By construction of $\Delta_{\mathsf{SamTrans}}$, the per-hop deviations $\|T_{ij}^{\mathbf{T}}(v) - T_{ij}^{\mathbf{T}'}(v')\|$ are bounded by $\Delta_{\mathsf{SamTrans}}(\mathbf{T}, \mathbf{T}')$ on in-regime latent points, and a standard path-comparison argument in metric geometry then yields

$$B \ \lesssim \ \Delta_{\mathsf{SamTrans}}(\mathbf{T}, \mathbf{T}').$$

Combining the three terms and taking suprema over in-regime latent pairs shows that the profile component of $\Delta_{\mathsf{SamGeom}}(\mathbf{M},\mathbf{M}')$ is bounded by a constant multiple of $C_F(\mathbf{T},\mathbf{T}')$.

Operational part. The operational component of $\Delta_{\mathsf{SamGeom}}$ compares the induced manifold distances and chart coordinates along benign paths, again over in-regime latent entities u. The same triangle-inequality decomposition

$$|d_{g^{\mathbf{T}}} - d_{g^{\mathbf{T}'}}| \;\leq\; |d_{g^{\mathbf{T}}} - \tilde{d}_{\mathbf{T}}| + |\tilde{d}_{\mathbf{T}} - \tilde{d}_{\mathbf{T}'}| + |\tilde{d}_{\mathbf{T}'} - d_{g^{\mathbf{T}'}}|$$

controls these discrepancies uniformly over the operational region, with the first and last terms again bounded by Γ^F_{real} and the middle term by $\Delta_{\mathsf{SamTrans}}$. This yields the desired bound for the operational piece and thus for $\Delta_{\mathsf{SamGeom}}(\mathbf{M},\mathbf{M}')$ as a whole.

Collecting constants into α_F, β_F produces (**??**). The argument for G_S follows the same template with chart distances and manifold distances swapped and Γ^G_{real} in place of Γ^F_{real}, yielding (**??**). □

7.3 Error-budget transfer as corollaries

We now explain how the main error-budget transfer results in Section **??** (Theorems **??** and **??**) follow from the gauge-stability theorem **??** together with the Lipschitz dependence of detector error budgets on the gauges.

Recall that Section **??** introduces aligned detectors $D^{\mathbf{R}}$ for $\mathbf{R} \in \mathsf{SamTrans}^0(S)$ or $\mathsf{SamGeom}^0(S)$, each with an error budget

$$(E^{\mathbf{R}}_{\mathrm{geom}}, E^{\mathbf{R}}_{\mathrm{link}}, \xi^{\mathbf{R}}, \zeta^{\mathbf{R}}, \delta^{\mathbf{R}}_{\mathrm{indep}}),$$

and an induced lower bound on correctness for high-risk alerts $\mathrm{LB}(\mathbf{R})$ of the Davis form.

The key analytic input from Section **??** is that these budgets are *Lipschitz in the gauges*: there exist functions $\Lambda_{\mathsf{SamTrans}}$ and $\Lambda_{\mathsf{SamGeom}}$ such that, for any pair of realizations $\mathbf{T},\mathbf{T}'$ and $\mathbf{M},\mathbf{M}'$,

$$|E^{\mathbf{T}}_{\mathrm{geom}} - E^{\mathbf{T}'}_{\mathrm{geom}}| \leq \Lambda_{\mathsf{SamTrans}}(\Delta_{\mathsf{SamTrans}}(\mathbf{T},\mathbf{T}')), \quad \text{and likewise for } E_{\mathrm{link}}, \xi, \zeta, \delta_{\mathrm{indep}}, \tag{29}$$

$$|E^{\mathbf{M}}_{\mathrm{geom}} - E^{\mathbf{M}'}_{\mathrm{geom}}| \leq \Lambda_{\mathsf{SamGeom}}(\Delta_{\mathsf{SamGeom}}(\mathbf{M},\mathbf{M}')), \tag{30}$$

with $\Lambda_{\mathsf{SamTrans}}, \Lambda_{\mathsf{SamGeom}}$ first-order at the origin.

We now combine these Lipschitz properties with Theorem **??**.

Corollary 2 (Error-budget equivalence along F_S)**.** *Let $\mathbf{T},\mathbf{T}' \in \mathsf{SamTrans}^0(S)$ and $\mathbf{M} = F_S(\mathbf{T})$, $\mathbf{M}' = F_S(\mathbf{T}')$. Let $D^{\mathbf{T}}$ and $D^{\mathbf{T}'}$ be aligned detectors on $\mathbf{T}$ and $\mathbf{T}'$ with Davis-style error budgets and lower bounds $\mathrm{LB}(\mathbf{T})$, $\mathrm{LB}(\mathbf{T}')$, and let $D^{\mathbf{M}}$, $D^{\mathbf{M}'}$ be their aligned MSR detectors with bounds $\mathrm{LB}(\mathbf{M})$, $\mathrm{LB}(\mathbf{M}')$.*

*Then there exists a first-order function $\Gamma^{\det}_F$ (depending only on the Lipschitz moduli in (**??**)–(**??**) and on the constants in C_F) such that*

$$\big|\mathrm{LB}(\mathbf{M}) - \mathrm{LB}(\mathbf{M}')\big| \;\leq\; \Gamma^{\det}_F(\Delta_{\mathsf{SamTrans}}(\mathbf{T},\mathbf{T}'), \Gamma^F_{\mathrm{real}}(\mathbf{T}), \Gamma^F_{\mathrm{real}}(\mathbf{T}')),$$

and similarly with $(\mathbf{M},\mathbf{M}')$ and $(\mathbf{T},\mathbf{T}')$ reversed:

$$\big|\mathrm{LB}(\mathbf{T}) - \mathrm{LB}(\mathbf{T}')\big| \;\leq\; \Gamma^{\det}_F(\Delta_{\mathsf{SamTrans}}(\mathbf{T},\mathbf{T}'), \Gamma^F_{\mathrm{real}}(\mathbf{T}), \Gamma^F_{\mathrm{real}}(\mathbf{T}')).$$

In particular, if $\Delta_{\mathsf{SamTrans}}(\mathbf{T},\mathbf{T}')$ and the realizability slacks of $\mathbf{T},\mathbf{T}'$ tend to zero, then all four *lower bounds* $\mathrm{LB}(\mathbf{T}), \mathrm{LB}(\mathbf{T}'), \mathrm{LB}(\mathbf{M}), \mathrm{LB}(\mathbf{M}')$ *converge together.*

Proof. Apply the Lipschitz property (**??**) to the pair $(\mathbf{M}, \mathbf{M}')$ and then use the gauge-stability bound (**??**):

$$|\mathrm{LB}(\mathbf{M}) - \mathrm{LB}(\mathbf{M}')| \lesssim \Lambda_{\mathsf{SamGeom}}(\Delta_{\mathsf{SamGeom}}(\mathbf{M}, \mathbf{M}')) \leq \Lambda_{\mathsf{SamGeom}}(C_F(\mathbf{T}, \mathbf{T}')).$$

Since C_F is first-order in $\Delta_{\mathsf{SamTrans}}$ and Γ^F_{real}, and $\Lambda_{\mathsf{SamGeom}}$ is first-order at the origin, the composite $\Gamma^{\mathrm{det}}_F := \Lambda_{\mathsf{SamGeom}} \circ C_F$ has the claimed first-order dependence.

The same reasoning applied directly to $(\mathbf{T}, \mathbf{T}')$ via (**??**) bounds $|\mathrm{LB}(\mathbf{T}) - \mathrm{LB}(\mathbf{T}')|$ by a (possibly different but comparable) first-order function of $\Delta_{\mathsf{SamTrans}}$ and Γ^F_{real}; we absorb both into the single Γ^{det}_F. □

Corollary 3 (Error-budget equivalence along G_S)**.** *Let* $\mathbf{M}, \mathbf{M}' \in \mathsf{SamGeom}^0(S)$ *and* $\mathbf{T} = G_S(\mathbf{M})$, $\mathbf{T}' = G_S(\mathbf{M}')$. *With aligned detectors and Davis-style error budgets as in Corollary* **??***, there exists a first-order function* Γ^{det}_G *such that*

$$|\mathrm{LB}(\mathbf{M}) - \mathrm{LB}(\mathbf{M}')|, \quad |\mathrm{LB}(\mathbf{T}) - \mathrm{LB}(\mathbf{T}')| \leq \Gamma^{\mathrm{det}}_G(\Delta_{\mathsf{SamGeom}}(\mathbf{M}, \mathbf{M}'), \Gamma^G_{\mathrm{real}}(\mathbf{M}), \Gamma^G_{\mathrm{real}}(\mathbf{M}')).$$

In particular, if $\Delta_{\mathsf{SamGeom}}(\mathbf{M}, \mathbf{M}')$ *and the realizability slacks of* $\mathbf{M}, \mathbf{M}'$ *tend to zero, all four lower bounds converge together.*

Proof. Identical in structure to Corollary **??**, using (**??**) and the Lipschitz property (**??**) for TSR error budgets. □

Together, Theorem **??** and Corollaries **??**–**??** formalize the idea that the TSR and MSR descriptions of a given sameness structure S are *equivalent as detection substrates*: small gauge distance between two TSRs implies small gauge distance between their manifold realizations, and vice versa, and aligned detectors on either side inherit this equivalence at the level of error budgets and lower correctness bounds.

8 Operational audits and evaluation protocols

This section specifies empirical audits and evaluation protocols for Davis systems. It is deliberately *procedural* and does not report numerical results, case studies, or deployment performance. The goal is to make the theoretical assumptions of Sections 3–5 falsifiable in practice: each audit targets a specific part of the geometry, separation, or error budget, and is designed to be run on held-out data or prospective streams.

Throughout, we fix a Davis system

$$\mathcal{D} = ((M, g, \delta), P(L), c, h, S, A, \text{error budget})$$

in some domain, with distortion profile $\varepsilon(L)$, configuration margins $(\kappa_{\mathrm{hard}}, \kappa_{\mathrm{soft}})$, and error-budget components $(E_{\mathrm{geom}}, E_{\mathrm{link}}, \xi, \zeta, \delta_{\mathrm{indep}})$ as defined earlier. All protocols assume strict train/validation splits and time-asymmetric estimation where applicable (no look-ahead).

8.1 Evaluation principles and pre-registration

Before running any audit or reporting any number, we recommend pre-registering:

(i) **Primary endpoints.** Which quantities will be assessed? Examples: (a) distortion tails $\varepsilon(L_\star)$ at the proposed operational horizon, (b) effective separation gap $\kappa_{\text{soft}} - 2R\varepsilon(L_\star)$, (c) empirical Cantelli bound tightness, (d) calibration error ξ in a high-risk band, (e) abstention coverage and residual failure ζ.

(ii) **Path and slice selection.** How will benign paths in $P(L)$ and evaluation slices be sampled (e.g., by time, region, demographic group, capture condition)? Which slices count as in-distribution R_{valid}?

(iii) **Thresholds and vacuity criteria.** Pre-specify acceptable ranges for distortion (e.g., $\varepsilon(L_\star) \le \varepsilon^\star$), error-budget sums $E_{\text{geom}} + E_{\text{link}} + \xi + \zeta \le \tau_{\text{vac}}$, and independence slack δ_{indep}.

(iv) **Baselines and ablations.** Define baselines (e.g., flat Euclidean models, models without smoothness regularization, non-geometric detectors), and which components will be ablated (e.g., remove path features, remove auxiliary detectors, remove abstention).

(v) **Data governance and leakage.** Fix time windows for training vs evaluation, forbid re-using evaluation data for tuning, and record configuration hashes for all runs to support auditability.

All audits below should be run on held-out validation data (retrospective or prospective) that is not used to train f_θ, fit S, or tune the abstention policy A.

8.2 Distortion audits: verifying the bounded-distortion regime

These audits target the geometric assumptions behind the distortion profile $\varepsilon(L)$ and the path-horizon choice $L_\star$ (Theorem 1 and the non-vacuity condition $\kappa_{\text{soft}} - 2R\varepsilon(L_\star) > 0$).

Protocol 1: geodesic–Euclidean distortion vs. path length.

(1) **Sample benign paths.** On validation data, construct a set of identity-preserving paths $\{\gamma_m\}_{m=1}^M \subset P(L_{\max})$ via natural temporal sequences (e.g., time windows, short trajectories) or controlled perturbations. Record their lengths $L(\gamma_m)$.

(2) **Subdivide into segments.** For each γ_m, sample pairs of times $0 \le s < t \le 1$ with arc-length distance $d_g\big(z(\gamma_m(s)), z(\gamma_m(t))\big)$ spanning a grid of target lengths $\ell \in \{\ell_1, \dots, \ell_K\}$ up to $L_{\max}$.

(3) **Approximate geodesic distance.** For each pair (s,t), approximate d_g using either:

- the path length of the interpolated trajectory in the embedding (if g is induced by a pullback metric), or
- a discrete geodesic approximation (e.g., graph-based shortest paths) constrained to lie within a local neighborhood.

(4) **Compute distortion ratios.** Let δ be the ambient chord metric (Euclidean distance in the embedding). For each pair, compute the distortion ratio

$$\rho(s,t) = \frac{\delta\big(z(\gamma_m(s)), z(\gamma_m(t))\big)}{d_g\big(z(\gamma_m(s)), z(\gamma_m(t))\big)}.$$

Aggregate ρ into bins by geodesic length ℓ.

(5) **Summarize distortion.** For each length bin ℓ, estimate $\mathbb{E}[|\rho - 1| \mid d_g \approx \ell]$ and upper quantiles (e.g., 95th/99th percentile). Plot or tabulate these as a "distortion curve" $\widehat{\varepsilon}_{\text{geom}}(\ell)$.

(6) **Choose $L_\star$ and validate.** Select an operational horizon $L_\star$ such that $\widehat{\varepsilon}_{\text{geom}}(\ell) \le \varepsilon^\star$ for all $\ell \le L_\star$, with $\varepsilon^\star$ chosen so that the effective soft margin $\kappa_{\text{soft}} - 2R\varepsilon^\star$ remains positive. Document any slices (e.g., specific conditions or subpopulations) where distortion tails are heavy; these may be excluded from R_{valid} or treated as out-of-distribution.

Protocol 2: curvature and smoothness along paths. To empirically test the smoothness assumptions used in Section 3:

(1) Approximate discrete curvature along each path by the angles between successive segments in the embedding, or by second differences of $z(t)$ in arc-length parameterization.

(2) Confirm that curvature statistics (mean, 95th percentile) remain below a pre-registered threshold for paths with length $\le L_\star$.

(3) Flag regimes where curvature spikes (e.g., hard scene cuts, recombination events, abrupt control changes) as candidates for abstention or separate modeling.

These two audits jointly support (T1) in Section 6.2: they test whether the bounded-distortion regime is empirically valid on the intended deployment distribution.

8.3 Configuration and margin audits

These audits target the configuration map c, coarse labels h, and the separation margins $(\kappa_{\text{hard}}, \kappa_{\text{soft}})$ used in the Cantelli bound and non-vacuity conditions.

Protocol 3: separation in geodesic and ambient distance.

(1) **Collect labeled pairs.** On validation data, assemble a set of endpoint pairs (x_a, x_b) with coarse labels $h\big(c(z(x_a))\big)$ and $h\big(c(z(x_b))\big)$, distinguishing "similar" vs. "changed" vs. "ambiguous" configurations.

(2) **Measure distances.** For each pair, compute both the approximate geodesic distance $d_g\big(z(x_a), z(x_b)\big)$ and ambient distance $\delta\big(z(x_a), z(x_b)\big)$.

(3) **Estimate margins.** For each label pair (e.g., similar–similar, similar–changed, changed–changed), estimate the empirical distribution of distances. From these, estimate:

$$\widehat{\kappa}_{\text{hard}} \approx \inf_{(a,b) \text{ cross-config}} d_g(z_a, z_b), \qquad \widehat{\kappa}_{\text{soft}} \approx \text{quantile}_q\big(d_g(z_a, z_b)\big)$$

for a pre-registered quantile q (e.g., 5th percentile of cross-config pairs).

(4) **Check non-vacuity under distortion.** Using the distortion estimates from the previous subsection, verify that $\widehat{\kappa}_{\text{soft}} - 2R\widehat{\varepsilon}_{\text{geom}}(L_\star) > 0$ on the slices of interest. If this inequality fails or is marginal, treat the corresponding slice as ambiguous and rely on abstention or auxiliary detectors.

Protocol 4: stability across slices. Repeat Protocol 3 across pre-registered slices (e.g., demographics, capture conditions, time periods) and compare the estimated margins; large variation may indicate bias in f_θ or in the configuration definition. Where gaps appear, consider slice-specific abstention policies or separate retraining.

8.4 Feature–geometry linkage audits

These audits target the monotone link $\Delta S \geq g(\Delta g)$ between configuration movement in geometry and changes in the scalar statistic S.

Protocol 5: monotone linkage in bins of geodesic displacement.

(1) **Define path segments and labels.** For each benign path $\gamma \in P(L_\star)$ and time pair (s, t), compute the geodesic displacement $\Delta g = d_g(z(\gamma(s)), z(\gamma(t)))$ and the corresponding feature-based change in the statistic, e.g., $\Delta S = S_t - S_s$ or S_t alone if S is cumulative.

(2) **Bin by geodesic distance.** Partition the range of Δg into bins $B_1, \ldots, B_K$ on an operational interval $[0, g_{\max}]$. For each bin, compute:

$$\overline{\Delta S}_k = \mathbb{E}[\Delta S \mid \Delta g \in B_k], \qquad \widehat{\mathrm{Var}}(\Delta S \mid \Delta g \in B_k).$$

(3) **Test monotonicity.** Fit a monotone regression $\widehat{g}$ of $\overline{\Delta S}_k$ on the midpoints of B_k, and measure deviations from monotonicity (e.g., number and magnitude of violations). Pre-register an acceptable tolerance for such violations.

(4) **Report linkage slack.** Summarize a linkage slack term $\widehat{\eta}_{\text{link}}$ capturing how far the empirical relation deviates from an ideal monotone link; this feeds into the linkage component of the error budget E_{link}.

Protocol 6: variance and finite-moment checks. To justify finite-variance Cantelli bounds, estimate second moments of S (or ΔS) conditional on configuration changes. If heavy tails are observed, tighten the operational range or incorporate robust statistics (e.g., winsorized S) and re-audit.

8.5 Calibration, abstention, and error-budget estimation

These audits estimate the non-geometric terms (ξ, ζ) and link theoretical error budgets to empirical behavior.

Protocol 7: calibration in the high-risk region.

(1) **Fit and freeze calibration.** Using historical training data, fit a monotone calibration map $\widehat{\pi} = g_{\text{cal}}(S)$ from the statistic to an interpretable risk quantity (e.g., probability of class change within horizon T). Freeze g_{cal} before calibration auditing.

(2) **Evaluate on validation.** On held-out or prospective data, compute calibration metrics (Brier score, expected calibration error, reliability curves) on:

- the full support of S, and

- a high-risk band (e.g., S above a pre-registered threshold where actions are taken).

(3) **Estimate ξ.** Define $\widehat{\xi}$ as the maximum deviation between empirical and nominal probabilities in the high-risk band, or as a pre-registered functional of the calibration curves. This quantity enters the error budget as the calibration term.

Protocol 8: abstention behavior and ζ.

(1) **Record abstentions.** Under a candidate abstention policy A, log for each evaluation sample whether the system abstained and the rationale (e.g., large distortion, high OOD score, large predictive interval).

(2) **Estimate coverage.** For each slice, estimate coverage $1 - A$ (fraction of non-abstained cases) and the conditional error rates among non-abstained decisions (e.g., misclassification rate or violation of a domain-specific safety criterion).

(3) **Estimate ζ.** Define $\widehat{\zeta}$ as the frequency with which the system *fails* to abstain in regimes where distortion or OOD scores exceed pre-registered safety thresholds. This captures abstention failures in the error budget.

Protocol 9: empirical error-budget decomposition. On validation data where ground-truth outcomes are available, label each failure with its dominant source (geometry drift, linkage failure, calibration error, abstention failure, or uncontrollable noise). Estimate empirical frequencies of each error type and compare to the theoretical components $(E_{\text{geom}}, E_{\text{link}}, \xi, \zeta)$; large discrepancies signal either model misspecification or unmodeled dependencies.

8.6 Independence slack and compositional bounds

The compositional dominance bound in Theorem 2 assumes approximate conditional independence of the error sources; when this fails, the theory recommends a conservative union bound. This subsection sketches how to empirically assess the independence slack δ_{indep}.

Protocol 10: independence diagnostics.

(1) **Error indicators.** For each sample, define Bernoulli indicators $E_{\text{geom}}, E_{\text{link}}, E_{\text{cal}}, E_{\text{abst}}$ for the occurrence of each error type, based on the audits above and domain-specific thresholds.

(2) **Estimate joint vs. product.** For selected pairs or triples of error types (e.g., $(E_{\text{geom}}, E_{\text{link}})$, $(E_{\text{link}}, E_{\text{cal}})$), estimate:

$$p_{12} = \mathbb{P}(E_1 \wedge E_2), \qquad p_1 = \mathbb{P}(E_1), \quad p_2 = \mathbb{P}(E_2),$$

and the deviation from independence $\Delta_{12} = p_{12} - p_1 p_2$.

(3) **Define δ_{indep}.** Aggregate the magnitudes $|\Delta_{12}|$ across pairs into an independence slack $\widehat{\delta}_{\text{indep}}$ (e.g., maximum or suitable norm), and compare the compositional bound using products to the conservative union bound using sums. If $\widehat{\delta}_{\text{indep}}$ is large, prefer the union bound in practice.

8.7 Slicing, OOD monitoring, and lifecycle audits

Finally, we outline cross-cutting audits that track geometry and performance over time and across slices.

Protocol 11: slice-wise Davis audits. For each pre-registered slice (e.g., region, demographic group, capture condition, hardware profile):

(1) Re-run Protocols 1–9 restricted to that slice.

(2) Compare distortion curves, margins, linkage slack, calibration error, and abstention behavior across slices.

(3) Flag slices where any component of the error budget becomes vacuous (e.g., $E_{\mathrm{geom}}+E_{\mathrm{link}}+\xi+\zeta > \tau_{\mathrm{vac}}$), and treat them as out-of-regime for the current system.

Protocol 12: OOD detection and retraining triggers. Monitor distributions of:

- distortion ratios $\rho(s,t)$,
- path curvature statistics,
- feature distributions used in S,
- OOD or density scores from auxiliary detectors.

If these drift beyond pre-registered control limits, increase abstention, freeze thresholds, and consider retraining f_θ or redefining $P(L)$ and $L_\star$ before resuming full operation.

Section scope note. This section specifies *audit protocols and reporting templates* for future empirical work with Davis systems. By design, it contains no numerical results, plots, or case studies. In practical deployments, we recommend that these protocols be instantiated as part of a comprehensive system dossier, with results refreshed on a fixed cadence and after major system changes.

9 Case study: a geometry-first multi-modal threat detector

The abstract framework of Sections **??**–**??** was developed independently of any particular domain. In parallel, the author built a deployed system—codenamed *KRAKEN* in a separate, application-focused manuscript—for real-time, multi-modal maritime threat detection.[3] In hindsight, KRAKEN can be understood as a concrete instance of the TSR–MSR equivalence and error-budget decomposition developed in this paper.

This section sketches that correspondence. It is not required for the theoretical results, but serves as a sanity check: a complex, safety-critical system built before the formalism was articulated nonetheless fits naturally into the semantic-sameness, realization, and functorial picture.

[3]The present section intentionally omits operational details (sensor layouts, campaign locations, adversary tactics) and focuses only on the structural aspects relevant to the TSR–MSR theory: semantic sameness, translator graphs, geometric realizations, and error budgets.

9.1 Operational setting (informal)

At a high level, the system fuses heterogeneous sensor streams to detect and track subsurface threats and undersea infrastructure hazards in coastal waters. Representative modalities include:

- **Kinetic** channels: distributed acoustic sensing (fiber strain), passive SONAR spectrograms, and related vibration measurements;
- **Electromagnetic** channels: synthetic aperture radar and optical/IR imagery of the sea surface, plus RF emissions from platforms and shore;
- **Contextual** channels: shipping lanes, bathymetry, environmental metadata, and operator annotations.

Each modality has its own units, SNR regime, failure modes, and coverage gaps. Operators care about *entities*: specific submarines, vessels, and patterns of behavior. The core question is semantic: when is a faint acoustic trace, a marginal image artifact, and a weak RF transient evidence of the *same* latent threat, and when are they unrelated clutter?

9.2 Semantic sameness structure for maritime threats

We briefly sketch a semantic sameness structure $S_{\text{mar}} = (I, \{X_i\}_{i\in\mathcal{I}}, \{\pi_i\}_{i\in\mathcal{I}}, \approx, \mathcal{P}_S(L))$ in the sense of Section **??** that captures this setting.

- **Latent space I.** Elements $u \in I$ represent latent maritime *threat states* at an instant: a tuple (C, E, Q) of threat class (platform or behavior type), environment (sea state, sound-speed profile, background traffic), and configuration (speed, depth, machinery state, loadout).
- **Modalities $\{X_i\}$.** Each $i \in \mathcal{I}$ indexes a sensor family (e.g., DAS fiber segments, SONAR beams, satellite tiles, RF channels). X_i is the space of idealized observations from that family over a short time window.
- **Rendering maps $\pi_i : I \to X_i$.** Given a latent state $u = (C, E, Q)$, the map $\pi_i(u)$ encodes what an ideal, noise-free sensor of modality i *would* observe: an acoustic spectrum induced by propeller cavitation, a surface wake pattern in SAR, a strain pattern along a buried fiber, an RF emission mask, and so on. In practice, real observations include noise, occlusion, and quantization; as in Section **??**, we blur notation and treat $x_i \in X_i$ as an observation "of" u when it lies within an acceptable noise ball around $\pi_i(u)$.
- **Sameness relation $\approx$.** Two observations $x \in X_i$, $x' \in X_j$ are semantically equivalent, $x \approx x'$, if they are rendered from the *same* latent entity within a short temporal tolerance: there exists a path $\gamma \in \mathcal{P}_S(L)$ and time t such that $x \approx \pi_i(\gamma(t))$, $x' \approx \pi_j(\gamma(t))$. Intuitively, they are evidence about the same contact rather than coincident clutter.
- **Benign path family $\mathcal{P}_S(L)$.** Paths $\gamma : [0, 1] \to I$ describe physically plausible threat trajectories over an operational horizon L: slow drifts in position and depth, machinery state changes, and routine environmental variation. This family encodes both how fast threats can move in latent space and which directions are benign.

KRAKEN was designed without this formal vocabulary, but its informal threat models and scenario design align closely with S_{mar}.

9.3 TSR instantiation: translator graph over sensor embeddings

On the translator side, KRAKEN provides a TSR $\mathbf{T}_{\text{mar}} \in \mathsf{SamTrans}^0(S_{\text{mar}})$ in the sense of Definition **??**.

- **Feature spaces V_i and encoders φ_i.** Each modality has an encoder $\varphi_i : X_i \to V_i$ (CNNs for spectrograms, transformers for sequences, etc.) producing modality-specific embeddings. These are the V_i in $\mathbf{T}_{\text{mar}}$.
- **Translator graph.** A sparse, directed graph G_T connects modalities via learned translators $T_{ij} : V_i \to V_j$. Edges exist where physics and coverage justify translation (e.g., DAS→SONAR, SONAR→DAS, DAS→SAT, RF→SAT). Translators are trained and calibrated to satisfy per-edge error bounds
$$\|T_{ij}(v_i) - v_j\| \le \varepsilon_{ij}$$
on in-regime data, yielding a drift profile $\varepsilon_{\text{trans}}(k)$ for k-hop paths as in Definition **??**. Spectral routing on G_T avoids paths whose accumulated $\varepsilon_{\text{trans}}$ exceeds a configured budget.
- **Chartability side conditions.** Empirically, the learned translators on the operational regime satisfy:
  - approximate local bi-Lipschitz behavior (bounded Jacobians);
  - small cocycle residual $\|T_{jk}(T_{ij}(v)) - T_{ik}(v)\|$ on triple overlaps;
  - bounded asymmetry $\|T_{ji}(T_{ij}(v)) - v\|$ on round-trips.

  These are precisely the smooth-chartability conditions from Section **??**, verified by held-out audits rather than assumed.

From the standpoint of this paper, KRAKEN's representation layer is a richly engineered TSR with a calibrated error budget and empirical evidence of smooth chartability.

9.4 MSR instantiation: a Davis-style threat manifold

On the geometric side, KRAKEN instantiates an MSR $\mathbf{M}_{\text{mar}} \in \mathsf{SamGeom}^0(S_{\text{mar}})$ in the spirit of Definition **??** and the Davis-manifold construction.

- **Threat manifold $(\mathcal{M}, g)$.** Multi-modal evidence fused through the translator graph is encoded into a low-dimensional Riemannian manifold $(\mathcal{M}, g)$ using a HERALD-style encoder trained with a contrastive objective and a smoothness regularizer. Geodesic distance d_g aligns with threat semantics: quiet and loud variants of the same platform lie close; distinct classes separate with a soft margin.
- **Charts and chart distances.** For each modality (or fused representation) there is a chart map $\chi_i : U_i \subset \mathcal{M} \to V_i$ with bounded distortion in the sense of Definition **??**. Distortion audits along benign paths in $\mathcal{P}_S(L)$ empirically validate a small distortion profile $\varepsilon_{\text{dist}}(\ell)$ on the operational horizon.
- **Semantic drift.** Realistic threat trajectories from $\mathcal{P}_S(L)$ induce curves $\gamma_M : [0,1] \to \mathcal{M}$ with geodesic span bounded by a semantic drift profile $\sigma_S^{\mathbf{M}}(L)$ as in Definition **??**. In practice, this drift profile is estimated from historical behavior and used to distinguish benign motion from novel tactics.

Internally, this manifold is treated as the system's "threat space": genomes, drift statistics, and risk policies all operate in $\mathcal{M}$.

9.5 Detector and error budget mapping

KRAKEN's detector stack—behavior genomes, risk-aware fusion, and deterministic replay—fits naturally into the sameness-detector abstraction of Definition **??**.

- **Statistic $\Sigma_{\mathbf{M}}$.** In the geometric view, the detection statistic $\Sigma_{\mathbf{M}} : \mathcal{M} \to \mathbb{R}$ is built from (i) distances to class-specific threat behavior genomes (shrinkage Mahalanobis templates in $\mathcal{M}$), (ii) drift statistics along recent trajectories, and (iii) simple auxiliary features (environmental covariates, sensor trust scores). A corresponding translator-side statistic $\Sigma_{\mathbf{T}} : \bigsqcup_i V_i \to \mathbb{R}$ is defined by composing modalities through G_S in the sense of Section **??**.
- **Decision rule A.** A CVaR-based fusion and thresholding rule A maps the statistic and context (mission phase, alert policies) to $\{0, 1, \text{abstain}\}$: trigger, suppress, or abstain.
- **Error decomposition.** Following Definition **??**, KRAKEN's end-to-end miss probability decomposes into:
  - *geometric/representation error* $E^{\mathbf{T}}_{\text{geom}}$, $E^{\mathbf{M}}_{\text{geom}}$ from translator drift and manifold distortion;
  - *linkage/observation error* E_{link} from noisy sensors and imperfect encoders;
  - calibration error ξ and abstention failure ζ from CVaR fusion and threshold policies;
  - independence slack δ_{indep} capturing dependence between failure modes.

  Internally, these terms are estimated by a combination of held-out evaluation, stress tests (sensor outages, adversarial conditions), and analytic bounds of the type derived in Section **??**.

A separate, application-focused document proves a KRAKEN-specific analogue of the compositional error-budget theorem: under Lipschitz and margin assumptions, stage-wise bounds on selection, translation, geometry, genomes, fusion, trust, and replay compose into an end-to-end miss probability bound $\varepsilon_{\text{total}}$ for a given mission profile.

9.6 How the general theorems apply

From the perspective of this paper, the most interesting feature of KRAKEN is that it uses *both* realizations at once:

- A translator-first view $\mathbf{T}_{\text{mar}}$ powers cross-modal prediction, redundancy, and graceful degradation when some sensors fail;
- A manifold-first view $\mathbf{M}_{\text{mar}}$ powers behavior genomes, drift detection, and zero-/few-shot generalization.

The functors F_S and G_S and the error-transfer theorems of Section **??** give a principled explanation of empirical observations made during KRAKEN's development:

- When the translator graph passes chartability audits (small $\varepsilon_{\text{trans}}$, δ_{asymm}), the induced manifold has bounded distortion $\varepsilon_{\text{dist}}$ and preserves threat semantics. This is precisely the regime where $F_S(\mathbf{T}_{\text{mar}})$ and $\mathbf{M}_{\text{mar}}$ become ε-equivalent.

- Conversely, when geometry audits flag increased distortion or eroded margins in $\mathcal{M}$ (e.g., after a distribution shift), projecting back through G_S identifies which translators, modalities, or threat classes are responsible. In categorical terms, this uses the stability of $\Delta_{\mathsf{SamTrans}}$ under G_S.
- Most importantly, the *detection* guarantees—bounds on miss probability, false positive rate, and abstention behavior—are invariant up to first-order slack under switching between translator- and manifold-centric implementations, as long as chartability and realizability assumptions hold. This matches the observed fact that translator-heavy and manifold-heavy configurations of the system achieve comparable detection performance when tuned to the same operational envelope.

9.7 Limitations and redactions

This case study is necessarily incomplete. It deliberately omits:

- detailed threat models, sensor geometries, and campaign specifics;
- numerical values of internal error budgets and performance metrics;
- implementation details of encoders, translators, genomes, and fusion policies.

Those belong in a separate, application-focused paper and, in some cases, cannot be made public.

Nonetheless, even in this abstracted form, the mapping between a real, deployed system and the semantic-sameness framework is clear:

1. a well-posed sameness structure S_{mar};
2. dual realizations $\mathbf{T}_{\text{mar}}$ and $\mathbf{M}_{\text{mar}}$ satisfying smooth chartability and geometric realizability;
3. an end-to-end detector whose error budget decomposes as in Definition **??**;
4. empirical audits that instantiate the chartability and distortion checks required by the theorems of Section **??**.

In this sense, KRAKEN (or any similarly structured system) serves as a proof of concept: the abstract, category-theoretic equivalence between translator-based and manifold-based views of semantic sameness is not only mathematically sound but also aligned with how one can design auditable, geometry-first detection systems in high-stakes settings.

References

[1] Davis, B. R. (2025). *Adaptive Embedding Format Switching and Semantic Translation in High-Volume Telemetry Systems*. Zenodo. `https://doi.org/10.5281/zenodo.17211403`

[2] Davis, B. R. (2025). *HERALD: High-resolution Early Recognition of Antigenic Landscape Divergence*. Zenodo. `https://doi.org/10.5281/zenodo.17640400`

[3] Davis, B. R. (2025). *The Davis Manifold: Geometry-First Detection with Compositional Error Budgets*. Zenodo. `https://doi.org/10.5281/zenodo.17642038`

[4] Davis, B. (2025). *PRISM v4.1: Risk-Aware Multi-Rail Payment Reconciliation with Bounded Semantic Translation and Spectral Consensus – Adapting NASA Safety-Critical Systems to Financial Infrastructure*. Zenodo. `https://doi.org/10.5281/zenodo.17489610`