Categorical Entropy Obstruction Theory IV

Finite Trace-Law Compression, Target Sufficiency, and Support-Relative Information Obstructions

Duy Hoang Nguyen
ORCID: 0009-0009-9228-7988

Abstract

This paper develops the finite trace-compression layer of Categorical Entropy Obstruction Theory. Given a finite algorithmic trace $S_{0:T}$, a deterministic compressed observation $O = c(S_{0:T})$, and a declared task target Y, it separates three questions that are often conflated: reconstruction of the trace, preservation of the target law, and preservation of task performance under a bounded loss.

The first contribution is a support-relative entropy obstruction theory for complete-trace, hidden-trace, bridge-relative, coordinatewise static, and visible-memory reconstruction. The second contribution is a task-relative sufficiency theory: deterministic target recovery is governed by $H(Y \mid O)$, stochastic target-law sufficiency by $I(Y; S_{0:T} \mid O)$, and bounded-loss degradation by Bayes-risk differences. The third contribution is a finite categorical packaging: deterministic observation refinement gives monotonicity as functoriality into the reverse ordered category $(\mathbb{R}_{\geq 0}, \geq)$, while coherent typed profiles give relabeling invariance and explain why scalar entropy shadows do not classify posterior trace profiles.

The theory is deliberately finite, law-level, and support-relative. It does not claim efficient decoder synthesis, generator-sensitive algorithmic complexity, a target-independent minimal quotient, or a universal categorical entropy functor. The main separation theorem shows that the observed target law $P_{O,Y}$, even up to law-isomorphism, is not a complete invariant for hidden-trace obstruction. Finite examples from dynamic programming and signature-dependent path systems illustrate the trace/task/memory distinctions.

Keywords. Categorical Entropy Obstruction Theory; finite trace-law compression; information-theoretic trace compression; finite stochastic systems; hidden trace reconstruction; deterministic target recovery; law-sufficiency; loss-relative sufficiency; conditional mutual information; visible-trace insufficiency; finite posterior profiles; support-relative obstruction.

Author identifier. Duy Hoang Nguyen, ORCID: 0009-0009-9228-7988.

Suggested citation. Duy Hoang Nguyen. *Categorical Entropy Obstruction Theory IV: Finite Trace-Law Compression, Target Sufficiency, and Support-Relative Information Obstructions.* 2026.

Metadata.

Title	Categorical Entropy Obstruction Theory IV: Finite Trace-Law Compression, Target Sufficiency, and Support-Relative Information Obstructions
Author	Duy Hoang Nguyen
ORCID	0009-0009-9228-7988
Year	2026
Series position	CEOT IV in the five-paper Categorical Entropy Obstruction Theory program
Immediate predecessor	CEOT III: Diagram-Collapse Obstructions for Finite Marked Sampleable Stochastic Presentations; DOI: 10.5281/zenodo.20518578
Primary mathematical role	Finite trace-law compression layer separating reconstruction obstruction, task-relative sufficiency, bounded-loss degradation, and finite categorical profile invariance

Predecessor DOI anchors. CEOT I: 10.5281/zenodo.20500996; CEOT II: 10.5281/zenodo.20506069; CEOT III: 10.5281/zenodo.20518578.

Entropy, divergence, and total-variation conventions. Throughout the paper, all logarithms are taken in a fixed base $b > 1$. Entropies and mutual informations are measured in the corresponding units. The base is immaterial for all zero, positivity, monotonicity, and non-completeness statements, since changing the base multiplies every entropy and mutual information by the same positive constant.

For a finite law P on a finite set $\mathcal{X}$, write

$$H_b(P) := -\sum_{x\in\mathcal{X}} P(x)\log_b P(x),$$

with the convention $0\log_b 0 = 0$. For finite laws P, Q on the same finite set, define

$$D_b(P\|Q) := \sum_{x:P(x)>0} P(x)\log_b \frac{P(x)}{Q(x)},$$

with $D_b(P\|Q) = +\infty$ if $P(x) > 0$ and $Q(x) = 0$ for some x. Mutual information and conditional mutual information are computed with the same base:

$$I_b(X;Y) := D_b(P_{X,Y}\|P_X P_Y),$$

$$I_b(X;Y \mid Z) := \mathbb{E}_Z D_b(P_{X,Y|Z}\|P_{X|Z}P_{Y|Z}).$$

When the base is fixed and no quantitative constant depends on it, we write H, I, and D for readability.

The total-variation distance is normalized as

$$\mathrm{TV}(P,Q) := \sup_A |P(A) - Q(A)| = \frac{1}{2}\sum_x |P(x) - Q(x)|.$$

With these conventions, Pinsker's inequality is

$$\mathrm{TV}(P,Q) \le \sqrt{\frac{\ln b}{2} D_b(P\|Q)}.$$

All conditional laws in these formulas are evaluated only on positive-probability conditioning fibers.

Obstruction notation. The standard CEOT IV obstruction symbols used below are:

Symbol	Formula	Meaning
$\mathrm{Ob}_{\mathrm{complete}}(O)$	$H(S_{0:T} \mid O)$	complete/full-trace reconstruction
$\mathrm{Ob}_{\mathrm{trace}}(O)$	$H(S_{\mathrm{int}} \mid O)$	internal-trace reconstruction
$\mathrm{Ob}_{\mathrm{bridge}}(O)$	$H(S_{\mathrm{int}} \mid O, S_0, S_T)$	bridge-relative reconstruction
$\mathrm{Ob}^{Y}_{\mathrm{det}}(O)$	$H(Y \mid O)$	deterministic-target recovery
$\mathrm{Ob}^{Y}_{\mathrm{law}}(O)$	$I(Y; S_{0:T} \mid O)$	stochastic target-law sufficiency
$\mathrm{Ob}^{Y}_{\ell}(O)$	$R_\ell(O) - R_\ell(S_{0:T})$	loss-relative decision obstruction
$\mathrm{Ob}_{\mathrm{mem}}(V)$	$H(M_{0:T} \mid V_{0:T})$	hidden-memory reconstruction
$\mathrm{Ob}_{\mathrm{mem}\to Y}$	$I(Y; M_{0:T} \mid V_{0:T})$	target-relevant hidden memory

Relation to previous papers. This paper continues CEOT I [10], which introduced the finite stochastic reconstruction-obstruction calculus, CEOT II [11], which studied scaling bridge obstructions and endpoint/interior reconstruction separation, and CEOT III [12], which replaced path-shaped systems by finite marked sampleable stochastic presentations and diagram-collapse posterior profiles. The present paper keeps the same finite, support-relative, profile-first discipline, but changes the object of study to compressed algorithmic traces and task-relative targets. Its obstruction values are law-level invariants of the induced trace-target law unless explicitly stated otherwise.

Remark 0.1 (Novelty boundary). The finite entropy identities used below are standard. The zero-reconstruction criterion $H_b(A \mid O) = 0$, the conditional-independence criterion $I_b(Y; S \mid O) = 0$, entropy monotonicity under refinement, Pinsker-type bounds, and finite Bayes-risk inequalities are used as analytic input, not claimed as new entropy identities. The CEOT IV contribution is the support-relative presentation layer built around these tools: it separates complete-trace, hidden-trace, bridge, deterministic-target, stochastic-target, loss-relative, and lifted-memory obstructions; organizes deterministic observations by coarser-to-finer refinement; factors scalar obstructions through coherent trace-target profiles; and proves finite non-completeness results showing that observed target behavior does not determine hidden trace obstruction. Thus CEOT IV is a finite obstruction taxonomy and profile-invariance framework for algorithmic trace compression, not a replacement for classical finite information theory.

Background lineage. The reconstruction criteria used here are finite information-theoretic criteria in the Shannon–Cover–Thomas tradition [1, 2, 3]. The target-relative and loss-relative viewpoints are closely related to classical sufficiency, statistical decision theory, and comparison of experiments [4, 5, 6, 7, 8, 14]. The lumpability comparison in Section 9.6 is related to classical finite Markov-chain aggregation [15, 9]. CEOT IV does not claim to replace these theories; it packages their finite entropy and sufficiency certificates around declared algorithmic trace presentations.

Contents

1 Introduction: From Trace Loss to Task-Relevant Loss

1.1 Motivation

Categorical Entropy Obstruction Theory begins with a finite reconstruction principle. If U is a finite hidden random variable and O is a finite observed random variable, then

$$H(U \mid O) = 0 \iff U \text{ is reconstructible from } O$$

up to null sets. As an isolated statement, this belongs to standard finite information theory [1, 2]. The purpose of CEOT is not to replace Shannon entropy, nor to introduce a new entropy identity.

Its purpose is to attach the pair (U, O) to an explicitly specified finite stochastic presentation, so that conditional entropy becomes a numerical certificate for a concrete reconstruction problem.

CEOT I developed this reconstruction viewpoint for finite stochastic morphisms, factorizations, trajectories, and bridges [10]. In the two-step factorization

$$X \overset{K}{\rightsquigarrow} Y \overset{L}{\rightsquigarrow} Z,$$

the hidden intermediate variable Y may fail to be reconstructible from the endpoints (X, Z) even when the endpoint composite is fixed. The corresponding obstruction is

$$H(Y \mid X, Z),$$

and its vanishing is exactly a support-relative reconstruction condition. For path-shaped trajectories, CEOT I similarly attaches bridge and hidden-path obstructions to intermediate states conditioned on endpoint data.

CEOT II moved from fixed finite trajectories to scaling families [11]. It showed that endpoint-level information and interior reconstruction can separate in asymptotic regimes. Endpoint uncertainty may become small under a chosen normalization while interior or bridge uncertainty remains macroscopically nonzero. Thus endpoint reconstruction, even if asymptotically strong, does not automatically certify reconstruction of hidden internal history.

CEOT III changed the shape of the finite object [12]. Instead of path-shaped bridges, it studied finite sampleable stochastic diagrams equipped with an observed marking. Observation becomes coordinate collapse: retained coordinates form the observed variable, while forgotten coordinates form the hidden completion variable. The basic object is no longer a single scalar entropy, but a posterior reconstruction profile

$$\mathsf{Post}_{U|O} = \left(P_O, \{P_{U|O=o}\}_{o \in \mathrm{supp}(O)}\right),$$

whose scalar entropy shadow is $H(U \mid O)$. CEOT III also established a finite, groupoid-level categorical invariance principle and a non-completeness phenomenon, using only the elementary relabeling-invariance layer of categorical language [13]: the same observed or endpoint behavior can coexist with different hidden posterior profiles and different hidden obstruction values.

The present paper develops the next layer. The object is now an algorithmic trace

$$S_{0:T} = (S_0, S_1, \ldots, S_T)$$

generated by a finite stochastic or randomized algorithmic presentation. Instead of asking only whether a hidden bridge, a hidden diagram completion, or an endpoint-conditioned interior state is reconstructible, we ask what information is lost when the trace is compressed to a finite observation

$$O = c(S_{0:T}).$$

The coordinatewise compressed trace

$$\bar{S}_{0:T} = (q_0(S_0), \ldots, q_T(S_T))$$

is the canonical roadmap case, but the formulation in terms of $O = c(S_{0:T})$ also allows global finite summaries of the entire computation.

The guiding question is therefore:

$$\boxed{\text{When does compression of an algorithmic trace destroy task-relevant information?}}$$

The answer is not determined by hidden trace loss alone. A compressed representation may fail to reconstruct the hidden computational history while still preserving the output, selected decision, accumulated cost, feasibility certificate, value, selected path, future outcome, or finite future-law parameter required by the task. This distinction is the central point of CEOT IV.

1.2 The main separation

Let $S_{0:T}$ be the complete trace, let $O = c(S_{0:T})$ be a deterministic finite compressed observation, and let Y be a task target. CEOT IV separates four different questions.

First, one can ask whether the entire trace is reconstructible from O. This is measured by the complete-trace obstruction

$$\mathrm{Ob}_{\mathrm{complete}}(O) := H(S_{0:T} \mid O).$$

Second, one can ask whether the hidden internal trace is reconstructible from O. This is measured by

$$\mathrm{Ob}_{\mathrm{trace}}(O) := H(S_{\mathrm{int}} \mid O).$$

Third, if the target is deterministic relative to positive trace support,

$$Y = f_+(S_{0:T}),$$

one can ask whether the target itself is recoverable from O. This is measured by

$$\mathrm{Ob}^{Y}_{\mathrm{det}}(O) := H(Y \mid O).$$

Fourth, if the target is generated by a stochastic kernel

$$L : \mathcal{S}_{0:T} \rightsquigarrow \mathcal{Y},$$

then Y may remain random even after the full trace is known. In that setting, the correct sufficiency obstruction is not usually $H(Y \mid O)$, but rather the conditional mutual information

$$\mathrm{Ob}^{Y}_{\mathrm{law}}(O) := I(Y; S_{0:T} \mid O).$$

This quantity vanishes exactly when O screens off Y from the full trace:

$$I(Y; S_{0:T} \mid O) = 0 \quad \Longleftrightarrow \quad Y \perp S_{0:T} \mid O.$$

Equivalently, on support,

$$P(Y \mid S_{0:T}, O) = P(Y \mid O).$$

Because $O = c(S_{0:T})$ is deterministic, the target kernel also satisfies

$$P(Y \mid S_{0:T}, O) = P(Y \mid S_{0:T})$$

on support. Thus stochastic sufficiency may be written in the operational form

$$P(Y \mid S_{0:T}) = P(Y \mid O)$$

on support.

The main separation is that hidden trace loss need not be task loss. One may have

$$H(S_{\mathrm{int}} \mid O) > 0$$

while still having

$$H(Y \mid O) = 0$$

for a deterministic target. One may likewise have

$$H(S_{\mathrm{int}} \mid O) > 0$$

while

$$I(Y; S_{0:T} \mid O) = 0$$

for a stochastic target. Thus CEOT IV rejects the over-strong criterion that an algorithmic compression is successful only when it reconstructs the entire hidden computational trace.

The hierarchy of trace reconstruction questions is also important. Conditioning on endpoints can only make hidden-trace reconstruction easier. Consequently the bridge, trace, and complete obstructions satisfy

$$\mathrm{Ob}_{\mathrm{bridge}}(O) \le \mathrm{Ob}_{\mathrm{trace}}(O) \le \mathrm{Ob}_{\mathrm{complete}}(O),$$

where

$$\mathrm{Ob}_{\mathrm{bridge}}(O) := H(S_{\mathrm{int}} \mid O, S_0, S_T).$$

This hierarchy imports the endpoint/interior discipline of CEOT I and CEOT II into algorithmic compression.

1.3 Why algorithmic compression needs target-relative obstruction

A compressed algorithmic record is rarely intended to reproduce every internal state of the computation. It may be intended to preserve only a final answer, an admissible decision, an optimality certificate, a feasibility witness, a value function, a selected path, or a predictive law. Therefore the relevant obstruction depends on the purpose assigned to the compressed observation.

For example, suppose the full trace has the form

$$S = (Y, Z), \qquad O = Y,$$

where Z is hidden randomness irrelevant to the target. Then

$$H(S \mid O) = H(Z \mid Y)$$

may be positive, while

$$H(Y \mid O) = 0.$$

The full trace is not reconstructible, but the deterministic target is perfectly preserved. This elementary model is not meant as a difficult information-theoretic example. Its purpose is structural: it shows that trace reconstruction and task preservation are different claims.

The stochastic version is similar. Suppose Y is generated by a target kernel $L : \mathcal{S}_{0:T} \rightsquigarrow \mathcal{Y}$. If O preserves the conditional law of Y carried by the full trace, then

$$I(Y; S_{0:T} \mid O) = 0,$$

even if hidden states remain unreconstructible. In this case O is target-sufficient relative to the stochastic criterion of this paper. Demanding $H(Y \mid O) = 0$ would be wrong, because Y may be intrinsically random even after $S_{0:T}$ is known.

This distinction is especially important for lifted-state algorithms. A full state may decompose as

$$S_t = (V_t, M_t),$$

where V_t is visible state and M_t is memory, signature, mode, label, or latent algorithmic state. Hidden memory may be unreconstructible:

$$H(M_{0:T} \mid V_{0:T}) > 0.$$

But this alone does not prove that memory matters for the task. The correct target-relevance obstruction is

$$I(Y; M_{0:T} \mid V_{0:T}).$$

Relative to the conditional-independence sufficiency criterion used in this paper,

$$V_{0:T} \text{ is target-sufficient for } Y \quad \iff \quad I(Y; M_{0:T} \mid V_{0:T}) = 0.$$

Thus a lifted state is not justified merely because it contains hidden information. It is justified, for the target Y, when its hidden component carries target-relevant information not already present in the visible trace.

1.4 Compressed observations and scope

The compressed observation in this paper is deterministic and finite:

$$O = c(S_{0:T}).$$

This choice is deliberate. It covers coordinatewise trace compression, global trace summaries, retained logs, deterministic sketches, endpoint-plus-certificate observations, and other finite summaries of a full algorithmic trace. It also keeps the reconstruction and monotonicity theory in the support-relative finite setting established by CEOT I–III.

The coordinatewise case remains the canonical roadmap model. In that case

$$O = \bar{S}_{0:T}, \qquad \bar{S}_t = q_t(S_t),$$

and one can define the local static obstruction

$$\mathrm{Ob}_{\mathrm{static}}(t) = H(S_t \mid \bar{S}_t).$$

If this local obstruction vanishes at every time, then the complete trace obstruction vanishes:

$$H(S_t \mid \bar{S}_t) = 0 \ \forall t \quad \implies \quad H(S_{0:T} \mid \bar{S}_{0:T}) = 0.$$

Indeed each S_t is then a function of $\bar{S}_t$ up to null sets, and the whole trace is a function of the compressed trace.

This paper does not treat randomized observation channels

$$C : S_{0:T} \rightsquigarrow O.$$

Such channels are natural in randomized sketching, noisy logging, and channel-valued compression, but they require a different refinement order, such as a Blackwell or garbling order [14]. Introducing that layer here would obscure the finite deterministic reconstruction calculus. Randomized observation channels are therefore left for a later channel-valued or measurable extension.

1.5 Posterior profiles rather than scalar-only obstruction

CEOT IV follows the profile-first discipline of CEOT III. For finite random variables A and O, the posterior reconstruction profile is

$$\mathsf{Post}_{A|O} = \Big(P_O, \{P_{A|O=o}\}_{o\in\mathrm{supp}(O)}\Big).$$

The scalar obstruction

$$H(A \mid O)$$

is a numerical shadow of this profile. It detects exact reconstructibility through its zero set, but it does not determine the entire posterior profile. Distinct posterior profiles may have the same scalar entropy, and the isomorphic observed target behavior may coexist with different hidden trace profiles.

For CEOT IV this matters in a particularly sharp way. There exist finite trace-compression presentations $\mathfrak{A}$ and $\mathfrak{A}'$ such that

$$P^{\mathfrak{A}}_{O,Y} \cong_{\text{law}} P^{\mathfrak{A}'}_{O',Y'},$$

but

$$\mathrm{Ob}^{\mathfrak{A}}_{\text{trace}}(O) \neq \mathrm{Ob}^{\mathfrak{A}'}_{\text{trace}}(O').$$

Thus observed target behavior, even up to support relabeling, does not determine hidden trace obstruction. This is the algorithmic analogue of the hidden-completion non-completeness phenomena isolated in CEOT III.

The finite categorical content of this paper has two explicit layers. First, for a fixed finite trace-target law, deterministic observations form a coarser-to-finer refinement category: a morphism $c \to c'$ means that c' refines c on positive trace support. Trace and target obstructions then define functors from this category to the ordered category $(\mathbb{R}_{\geq 0}, \geq)$, equivalently antitone numerical maps with respect to refinement. Second, finite trace-compression presentations form an isomorphism groupoid under law-preserving relabelings, and their scalar obstructions factor through a coherent finite trace-target profile groupoid:

$$\mathbf{AlgCompPres}^{\cong}_{\text{fin}} \longrightarrow \mathbf{TraceTargetProf}^{\cong}_{\text{fin,coh}} \longrightarrow \mathbb{R}^{\text{disc}}_{\geq 0}.$$

No universal categorical entropy functor is claimed. The categorical assertion is exactly the finite refinement/factorization structure stated above; it is not a higher-categorical classification theorem.

1.6 Contributions

The paper makes three main contributions.

Finite reconstruction obstructions. We define support-relative entropy obstructions for complete trace, hidden trace, bridge-relative trace, coordinatewise static state, and visible-memory reconstruction. The corresponding zero criteria are stated with explicit support-relative decoders, and strict non-converse examples separate bridge, hidden-trace, complete-trace, static, and memory reconstruction tasks.

Task-relative sufficiency. We separate deterministic target recovery, stochastic target-law sufficiency, fixed-loss sufficiency, universal bounded-loss sufficiency, and approximate obstruction interfaces. Deterministic targets are governed by $H(Y \mid O)$, stochastic law-sufficiency by $I(Y; S_{0:T} \mid O)$, and bounded-loss degradation by Bayes-risk differences. This separates trace reconstruction from target performance.

Finite categorical packaging. We construct the deterministic observation-refinement category and prove obstruction monotonicity as functoriality into the reverse ordered category $(\mathbb{R}_{\geq 0}, \geq)$. Separately, coherent typed profiles give relabeling invariance, factorization of the declared scalar obstruction family, and non-completeness of observed target laws for hidden-trace obstruction.

1.7 Claim discipline and non-goals

This paper is deliberately finite and support-relative. It does not claim that conditional entropy is a new invariant of probability theory. It uses conditional entropy and conditional mutual information as reconstruction and sufficiency certificates attached to specified finite presentations.

It also does not claim that an entropy-zero decoder is computationally efficient:

$$\text{entropy-zero decoder} \not\Rightarrow \text{efficient algorithm}.$$

The statements here are information-theoretic and support-relative. They certify existence of exact finite decoders or conditional-independence sufficiency, not polynomial-time recovery.

It does not claim that hidden trace loss implies task failure:

$$\text{hidden trace loss} \not\Rightarrow \text{target failure}.$$

Indeed, the separation between hidden trace loss and target preservation is one of the main points of the paper.

It does not attempt to replace Markov lumpability theory. Lumpability concerns preservation of visible transition laws. CEOT IV concerns reconstruction and target sufficiency relative to a specified hidden trace and task variable. These questions interact, but neither subsumes the other.

Finally, the paper does not claim a universal categorical entropy functor. Its categorical content is finite and explicit: obstruction maps are functorial into $(\mathbb{R}_{\geq 0}, \geq)$ on the coarser-to-finer deterministic observation-refinement category, and scalar obstruction values are invariant under law-preserving relabeling because they factor through coherent finite trace-target profiles.

1.8 Organization of the paper

Section 2 will define typed finite trace-compression presentations, deterministic compressed observations, coordinatewise compressed traces, deterministic targets, stochastic targets, and finite future-law parameters.

Section 3 will introduce posterior reconstruction profiles and explains why scalar entropy should be regarded as a shadow of a profile rather than as the entire invariant.

Section 4 will define complete-trace, hidden-trace, bridge, static, and memory obstructions, proves the support-relative zero criteria, and establishes the bridge/trace/complete hierarchy.

Section 5 will treat deterministic target recovery and proves that complete trace recovery implies deterministic target recovery, while the converse fails.

Section 6 will treat stochastic target sufficiency using conditional mutual information and explains why deterministic target recovery is a special case.

Section 7 will construct the coarser-to-finer deterministic observation-refinement category, prove functoriality into $(\mathbb{R}_{\geq 0}, \geq)$, and separate this category from finite coherent-profile groupoid invariance.

Section 8 will prove the visible-trace insufficiency theorem and distinguishes hidden memory from target-relevant memory.

Section 9 will give separation and non-completeness principles, including a trace-level non-completeness construction and lumpability witnesses.

Section 10 will present finite case studies from dynamic programming and signature-dependent path problems.

Section 11 will relate CEOT IV to CEOT I–III and explains how CEOT V should move to finite-factor and measurable reconstruction.

2 Finite Algorithmic Compression Presentations

This section fixes the finite objects used throughout the paper. The point is to separate three layers of data that are often conflated in informal discussions of algorithmic compression:

$$\text{the full trace law,} \qquad \text{the compressed observation,} \qquad \text{the task target.}$$

The trace law specifies what computation may occur. The observation specifies what part or summary of that computation is retained. The target specifies what the compression is meant to preserve. All later obstruction quantities are attached to this triple.

2.1 Finite and support-relative conventions

All state spaces in this paper are finite and nonempty. A finite stochastic kernel

$$K : X \rightsquigarrow Y$$

means a function assigning to each $x \in X$ a probability distribution $K(\cdot \mid x)$ on Y. We write

$$\operatorname{supp}(K(x, \cdot)) = \{y \in Y : K(y \mid x) > 0\}.$$

A deterministic map $f : X \to Y$ is regarded as the special kernel

$$K_f(y \mid x) = \mathbf{1}_{\{y=f(x)\}}.$$

Thus deterministic algorithms and randomized algorithms may be treated in one finite stochastic notation: deterministic transitions are Dirac kernels, while randomized transitions are non-degenerate kernels.

Spaces versus random variables. Calligraphic letters denote finite spaces, while capital Roman letters denote random variables taking values in those spaces. Thus

$$S_t : \Omega \to \mathcal{S}_t, \qquad O : \Omega \to \mathcal{O}, \qquad Y : \Omega \to \mathcal{Y}.$$

The full trace space is

$$\mathcal{S}_{0:T} := \mathcal{S}_0 \times \cdots \times \mathcal{S}_T,$$

and the random full trace is

$$S_{0:T} = (S_0, \ldots, S_T).$$

Maps such as compression maps, coordinate projections, relabelings, and target kernels are maps between finite spaces, not maps between random variables. When a support-relative object is used, the corresponding space is the positive-support set, for instance $\operatorname{supp}(O) \subseteq \mathcal{O}$.

All reconstruction statements are support-relative. If A and O are finite random variables, the assertion that A is reconstructible from O means that there exists a function

$$r : \operatorname{supp}(O) \to \operatorname{supp}(A)$$

such that

$$r(O) = A \quad \text{almost surely.}$$

No assertion is made about values of r outside $\operatorname{supp}(O)$. This convention is the same finite support discipline used in CEOT I–III: null fibers are irrelevant to reconstruction, while positive-support fibers control obstruction.

Convention 2.1 (Entropy base)**.** All logarithms are taken in a fixed base $b > 1$. Entropies and mutual informations are measured in the corresponding units. Changing b multiplies all entropy and mutual information values by the same positive constant; hence zero criteria, positivity criteria, monotonicity, and non-completeness statements are base-independent. In particular,

$$H(\text{Bernoulli}(1/2)) = \log_b 2,$$

which equals one bit when $b = 2$.

Remark 2.2 (Information-theoretic scope)*.* This paper studies information-theoretic compression of finite algorithmic traces. The term algorithmic refers to the origin and interpretation of the trace, not to a claim that an induced decoder is computationally efficient. Entropy-zero statements certify support-relative set-theoretic decoders; they do not certify polynomial-time algorithms, small circuits, low-memory procedures, or succinct descriptions of those decoders.

2.2 Finite algorithmic trace systems

A finite algorithmic trace system consists of a finite time horizon

$$T \geq 1,$$

finite nonempty state spaces

$$\mathcal{S}_0, \mathcal{S}_1, \ldots, \mathcal{S}_T,$$

an initial law

$$\mu_0 \in \mathrm{Prob}(\mathcal{S}_0),$$

and finite transition kernels

$$K_t : \mathcal{S}_t \rightsquigarrow \mathcal{S}_{t+1}, \qquad 0 \leq t \leq T-1.$$

These data induce random variables $S_t : \Omega \to \mathcal{S}_t$. For a concrete trace value

$$s_{0:T} = (s_0, s_1, \ldots, s_T) \in \mathcal{S}_{0:T},$$

the induced trace law is

$$P(S_{0:T} = s_{0:T}) = \mu_0(s_0) \prod_{t=0}^{T-1} K_t(s_{t+1} \mid s_t).$$

The positive trace support is therefore

$$\mathrm{supp}(S_{0:T}) = \left\{ s_{0:T} \in \mathcal{S}_{0:T} : \mu_0(s_0) > 0,\ K_t(s_{t+1} \mid s_t) > 0\ \forall t \right\}.$$

Remark 2.3 (Algorithmic interpretation)*.* The variable S_t is not required to be a physical state. It may encode a program counter, a graph vertex, a dynamic-programming table entry, a memory label, a randomized seed state, an accumulated signature, a certificate state, or any finite object needed to describe the algorithm at time t. The kernels K_t encode either algorithmic randomness or an externally randomized input process. If the algorithm is deterministic after the initial state is fixed, the K_t are Dirac kernels.

Remark 2.4 (Non-homogeneous and time-dependent systems)*.* The kernels K_t are allowed to depend on t. Thus the formalism includes non-homogeneous finite Markov processes, finite randomized algorithms with time-dependent update rules, and finite unfoldings of algorithms whose transition rule depends on the current stage. No stationarity assumption is used in CEOT IV.

2.3 Law-level versus generator-level information

A finite algorithmic trace system contains generator-level data

$$\mathfrak{G} = (\mathcal{S}_{0:T}, \mu_0, K_0, \dots, K_{T-1}).$$

After the deterministic observation map and target kernel are declared, these data induce a joint trace-target law

$$P_{S_{0:T},O,Y}.$$

The CEOT IV obstruction quantities studied in this paper are law-level quantities of this induced joint law. They measure support-relative reconstruction, stochastic target-law sufficiency, and loss-relative Bayes-risk degradation after the finite presentation has generated its trace-target law.

Theorem 2.5 (Law-level factorization of CEOT IV obstructions)**.** *Let two finite deterministic trace-compression presentations induce isomorphic joint laws*

$$P_{S_{0:T},O,Y} \cong_{\mathrm{law}} P_{S'_{0:T},O',Y'}$$

through bijections preserving the declared trace coordinates, compressed observations, target variables, and positive supports. Then the CEOT IV obstruction values transported by these bijections agree. In particular,

$$\mathrm{Ob}_{\mathrm{complete}}(O) = \mathrm{Ob}_{\mathrm{complete}}(O'), \qquad \mathrm{Ob}_{\mathrm{trace}}(O) = \mathrm{Ob}_{\mathrm{trace}}(O'), \qquad \mathrm{Ob}_{\mathrm{bridge}}(O) = \mathrm{Ob}_{\mathrm{bridge}}(O'),$$

$$\mathrm{Ob}^{Y}_{\mathrm{det}}(O) = \mathrm{Ob}^{Y'}_{\mathrm{det}}(O'), \qquad \mathrm{Ob}^{Y}_{\mathrm{law}}(O) = \mathrm{Ob}^{Y'}_{\mathrm{law}}(O').$$

If the loss functions are transported by the target-decision relabeling, then also

$$\mathrm{Ob}^{Y}_{\ell}(O) = \mathrm{Ob}^{Y'}_{\ell'}(O').$$

Proof. Each displayed obstruction is a finite conditional entropy, conditional mutual information, or Bayes-risk difference computed from the joint law of the involved variables. These quantities are invariant under support-preserving relabeling of the finite variables. The loss-relative statement follows because transporting the target and decision labels transports the admissible decision rules and their risks bijectively. □

Remark 2.6 (What CEOT IV does not distinguish)*.* The law-level factorization theorem deliberately forgets generator-level differences. Distinct transition kernels, distinct program structures, or distinct computational implementations can induce the same joint law $P_{S_{0:T},O,Y}$. CEOT IV then assigns the same obstruction values to them. Therefore CEOT IV should not be read as a theory of runtime, circuit size, memory complexity, succinct decoder synthesis, or computational complexity of constructing a decoder. It is a finite information-theoretic obstruction theory for declared trace laws and declared task targets.

2.4 Deterministic compressed observations

The compressed data retained from the trace is a finite deterministic observation

$$O = c(S_{0:T}),$$

where

$$c : \mathcal{S}_{0:T} \to \mathcal{O}$$

is a map into a finite observation space $\mathcal{O}$. The support of the observation is

$$\mathrm{supp}(O) = c(\mathrm{supp}(S_{0:T})).$$

The observation O may retain the whole trace, discard the whole trace, retain only endpoints, retain a certificate, retain a compressed log, or retain a global summary such as a cost, path label, or memory signature.

The deterministic observation assumption is part of the scope of this paper. Randomized observation channels

$$C : \mathcal{S}_{0:T} \rightsquigarrow \mathcal{O}$$

are natural for noisy logging, randomized sketches, and channel-valued compression, but they require a different refinement theory, for instance a Blackwell or garbling order [14]. CEOT IV keeps the deterministic finite setting in order to isolate exact support-relative reconstruction and task sufficiency. Randomized observation channels are reserved for later channel-valued or measurable extensions.

Definition 2.7 (Finite deterministic trace-compression presentation)**.** A finite deterministic trace-compression presentation is a tuple

$$\mathfrak{A} = (T, \{\mathcal{S}_t\}_{t=0}^{T}, \mu_0, \{K_t\}_{t=0}^{T-1}, \mathcal{O}, c, \mathcal{Y}, L),$$

where $(T, \{\mathcal{S}_t\}, \mu_0, \{K_t\})$ is a finite algorithmic trace system, $\mathcal{O}$ is a finite observation space, $c : \mathcal{S}_{0:T} \to \mathcal{O}$ is a deterministic compression map, $\mathcal{Y}$ is a finite target space, and

$$L : \mathcal{S}_{0:T} \rightsquigarrow \mathcal{Y}$$

is a finite target kernel. The random variables $S_0, \ldots, S_T, O, Y$ are generated by the induced joint law.

Definition 2.8 (Support-restricted target kernel)**.** Let

$$\mathcal{T}_+ := \mathrm{supp}(S_{0:T})$$

be the positive trace support induced by a finite deterministic trace-compression presentation. The support-restricted target support is

$$\mathcal{Y}_+ := \left\{ y \in \mathcal{Y} : \sum_{\tau \in \mathcal{T}_+} P(S_{0:T} = \tau) L(y \mid \tau) > 0 \right\}.$$

The support-restricted target kernel is

$$L_+ : \mathcal{T}_+ \rightsquigarrow \mathcal{Y}_+, \qquad L_+(y \mid \tau) := L(y \mid \tau)$$

for $\tau \in \mathcal{T}_+$ and $y \in \mathcal{Y}_+$. All CEOT IV target obstructions, target profiles, and profile isomorphisms depend on the ambient target kernel L only through L_+.

Convention 2.9 (Internal-trace convention)**.** For every horizon $T \geq 1$, define the internal trace space by the empty-product convention

$$\mathcal{S}_{\mathrm{int}} := \prod_{t=1}^{T-1} \mathcal{S}_t.$$

If $T = 1$, this product is the singleton set

$$\mathcal{S}_{\mathrm{int}} = \{*_{\mathrm{int}}\}.$$

The internal trace random variable is

$$S_{\mathrm{int}} = (S_1, \dots, S_{T-1}) \in \mathcal{S}_{\mathrm{int}},$$

with the convention that $S_{\mathrm{int}} = *_{\mathrm{int}}$ when $T = 1$. Thus every occurrence of S_{int} below is horizon-safe: no expression such as $S_{1:0}$ is used.

Lemma 2.10 (Finite witness realization lemma)**.** *Let W be any finite random variable with law P_W. Let*

$$c_W : \mathrm{supp}(W) \to \mathcal{O}$$

be a deterministic observation map, and let

$$L_W : \mathrm{supp}(W) \rightsquigarrow \mathcal{Y}$$

be a finite target kernel. Then the finite witness data (W, c_W, L_W) are realized by a finite CEOT IV trace-compression presentation.

More explicitly, take $T = 2$,

$$\mathcal{S}_0 = \{*_0\}, \qquad \mathcal{S}_1 = \mathrm{supp}(W), \qquad \mathcal{S}_2 = \{*_2\},$$

set

$$\mu_0(*_0) = 1,$$

define the transition kernels by

$$K_0(w \mid *_0) = P_W(w), \qquad K_1(*_2 \mid w) = 1,$$

and define

$$\widetilde{c}(*_0, w, *_2) = c_W(w), \qquad \widetilde{L}(y \mid *_0, w, *_2) = L_W(y \mid w).$$

Then the induced presentation satisfies

$$S_{\mathrm{int}} = S_1 \sim W, \qquad O = c_W(W), \qquad P(Y \mid S_{\mathrm{int}} = w) = L_W(\cdot \mid w).$$

Proof. The displayed data define a finite deterministic trace-compression presentation. The induced joint law is

$$P(*_0, w, *_2, o, y) = P_W(w)\mathbf{1}_{\{o=c_W(w)\}} L_W(y \mid w),$$

which is exactly the desired finite witness law written as a CEOT IV trace. Thus random-variable witnesses used later are shorthand for genuine finite trace kernels. □

The presentation induces a joint law on

$$\mathcal{S}_{0:T} \times \mathcal{O} \times \mathcal{Y}$$

by

$$P_{\mathfrak{A}}(s_{0:T}, o, y) = \mu_0(s_0) \prod_{t=0}^{T-1} K_t(s_{t+1} \mid s_t) \mathbf{1}_{\{o=c(s_{0:T})\}} L(y \mid s_{0:T}).$$

All conditional entropies and conditional mutual informations used later are computed with respect to this law.

Remark 2.11 (Why the target kernel is part of the presentation)*.* The same trace law and the same compression map can produce different obstruction statements depending on the assigned target. If the target is the full output, the relevant obstruction may differ from the obstruction for a selected decision, a feasibility certificate, a future-law parameter, or an optimal path. CEOT IV is therefore purpose-relative: the target kernel L is part of the mathematical object, not an afterthought.

2.5 Coordinatewise compressed traces

The canonical roadmap case is coordinatewise trace compression. For each time t, let

$$q_t : S_t \to \bar{S}_t$$

be a finite compression map. The compressed trace is

$$\bar{S}_{0:T} = (\bar{S}_0, \bar{S}_1, \ldots, \bar{S}_T) = (q_0(S_0), q_1(S_1), \ldots, q_T(S_T)).$$

This is recovered from the general observation notation by taking

$$O = \bar{S}_{0:T}$$

and

$$c(s_{0:T}) = (q_0(s_0), q_1(s_1), \ldots, q_T(s_T)).$$

The coordinatewise case is important because it supports a local static obstruction

$$\mathrm{Ob}_{\mathrm{static}}(t) = H(S_t \mid \bar{S}_t),$$

which measures whether the full state at time t can be reconstructed from the compressed state at that same time. The general observation $O = c(S_{0:T})$ is introduced not to replace the coordinatewise model, but to avoid excluding global finite summaries of the trace.

Example 2.12 (Endpoint observation)**.** If

$$O = (S_0, S_T),$$

then CEOT IV reduces part of its trace reconstruction problem to the endpoint/interior viewpoint of CEOT I–II. The hidden-trace obstruction is

$$H(S_{\mathrm{int}} \mid S_0, S_T),$$

which is the algorithmic analogue of a bridge obstruction.

Example 2.13 (Coordinatewise visible state)**.** Suppose

$$S_t = (V_t, M_t),$$

where V_t is visible state and M_t is hidden memory. Taking

$$q_t(V_t, M_t) = V_t$$

gives

$$O = V_{0:T}.$$

The memory reconstruction obstruction later studied in the lifted-state section is

$$H(M_{0:T} \mid V_{0:T}).$$

The target relevance of this memory is not measured by this entropy alone, but by

$$I(Y; M_{0:T} \mid V_{0:T}).$$

2.6 Deterministic targets

A target is deterministic if there exists a support-relative function

$$f_+ : \mathcal{T}_+ \to \mathcal{Y}$$

such that

$$Y = f_+(S_{0:T}).$$

Equivalently, the support-restricted target kernel is the Dirac kernel

$$L_+(y \mid s_{0:T}) = \mathbf{1}_{\{y=f_+(s_{0:T})\}}.$$

Typical deterministic targets include:

$$\text{final output}, \qquad \text{selected decision}, \qquad \text{selected path},$$

$$\text{accumulated cost}, \qquad \text{value certificate}, \qquad \text{feasibility certificate}.$$

For such targets, exact target recovery from O is measured by

$$H(Y \mid O).$$

The vanishing condition

$$H(Y \mid O) = 0$$

means that the target is a support-relative function of the observation, even if the full trace is not reconstructible.

Example 2.14 (Trace loss with deterministic target preservation)**.** Let $T = 1$, let $S_1 = (Y, Z)$, and let $O = Y$. If Z is positive-entropy randomness independent of Y, then

$$H(S_1 \mid O) = H(Y, Z \mid Y) = H(Z \mid Y) > 0,$$

but

$$H(Y \mid O) = 0.$$

Thus a deterministic target may be perfectly preserved while hidden trace information is lost.

2.7 Stochastic targets

A target is stochastic if it is generated by a non-Dirac kernel

$$L : \mathcal{S}_{0:T} \rightsquigarrow \mathcal{Y}.$$

In this case, even perfect knowledge of the full trace need not determine the realized target. Therefore exact target recovery in the sense of $H(Y \mid O) = 0$ is usually the wrong requirement.

The appropriate CEOT IV condition is target sufficiency: the observation O should preserve all information about Y carried by the full trace. This is measured by

$$I(Y; S_{0:T} \mid O).$$

The vanishing condition

$$I(Y; S_{0:T} \mid O) = 0$$

means

$$Y \perp S_{0:T} \mid O.$$

Thus O is target-sufficient precisely when, after O is known, the full trace carries no additional information about the stochastic target.

Lemma 2.15 (Target-kernel compatibility)**.** *Let $\mathfrak{A}$ be a finite deterministic trace-compression presentation, with $O = c(S_{0:T})$ and support-restricted target kernel L_+. Then, on positive trace support,*

$$P(Y \mid S_{0:T}, O) = P(Y \mid S_{0:T}) = L_+(\cdot \mid S_{0:T}).$$

Consequently,

$$I(Y; S_{0:T} \mid O) = 0$$

if and only if

$$P(Y \mid S_{0:T}) = P(Y \mid O)$$

on support.

Proof. Because $O = c(S_{0:T})$, conditioning on $(S_{0:T}, O)$ gives the same positive-support trace information as conditioning on $S_{0:T}$ alone. The target is generated from the full trace by the kernel L, but on positive trace support this kernel is exactly L_+. Therefore

$$P(Y \mid S_{0:T}, O) = L_+(\cdot \mid S_{0:T}) = P(Y \mid S_{0:T})$$

on support. The conditional mutual information $I(Y; S_{0:T} \mid O)$ vanishes exactly when Y and $S_{0:T}$ are conditionally independent given O, equivalently when

$$P(Y \mid S_{0:T}, O) = P(Y \mid O)$$

on support. Substituting the first identity gives the stated form. □

Remark 2.16 (Deterministic targets as a special case). If $Y = f_+(S_{0:T})$ is a support-relative deterministic target, then

$$I(Y; S_{0:T} \mid O) = H(Y \mid O),$$

because $H(Y \mid S_{0:T}, O) = 0$. Thus deterministic target recovery is the deterministic special case of stochastic target sufficiency.

2.8 Future outcomes and future-law parameters

In algorithmic applications, the phrase "preserve the future law" can mean two different things. The first is to preserve predictive sufficiency for a future random outcome. Let

$$Y_{\mathrm{fut}}$$

be a finite future outcome generated from the full trace. If the task is to preserve the predictive information about this random future outcome, the obstruction is

$$I(Y_{\mathrm{fut}}; S_{0:T} \mid O).$$

The zero condition says that O is sufficient for predicting Y_{fut} relative to the full trace.

The second meaning is to recover the future-law parameter itself. Define

$$\Theta = \mathcal{L}(Y_{\mathrm{fut}} \mid S_{0:T}).$$

Since $S_{0:T}$ is finite, Θ takes values in the finite realized set

$$\{\mathcal{L}(Y_{\mathrm{fut}} \mid S_{0:T} = s) : s \in \operatorname{supp}(S_{0:T})\}.$$

Thus Θ is a finite deterministic target. If the task is to recover the conditional future law itself, the correct obstruction is

$$H(\Theta \mid O).$$

This distinction prevents a common ambiguity: preserving prediction of a random outcome is a stochastic sufficiency problem, while recovering the law that governs the outcome is a deterministic target recovery problem.

2.9 Observation refinement

The later monotonicity results use a deterministic refinement order. For two deterministic observations

$$O = c(S_{0:T}), \qquad O' = c'(S_{0:T}),$$

we write

$$O \preceq O'$$

if

$$\sigma(O) \subseteq \sigma(O') \subseteq \sigma(S_{0:T}).$$

Equivalently, there exists a function

$$h : \mathrm{supp}(O') \to \mathrm{supp}(O)$$

such that

$$O = h(O') \quad \text{a.s.}$$

Thus O' is at least as informative about the trace as O. Under this order, conditional entropy obstructions decrease when the observation is refined.

Remark 2.17 (Refinement versus isomorphism). Observation refinement is not the same structure as presentation isomorphism. Refinement compares two observations by information content within a fixed trace law. Isomorphism, developed later, compares presentations up to law-preserving relabeling. CEOT IV keeps these two structures separate: refinement gives monotonicity; isomorphism gives invariance.

2.10 Section summary

A CEOT IV presentation consists of a finite trace law, a deterministic compressed observation, and a target kernel. The trace law determines $S_{0:T}$; the observation $O = c(S_{0:T})$ determines what data are retained; the target kernel $L : \mathcal{S}_{0:T} \rightsquigarrow \mathcal{Y}$ determines the task. The three fundamental questions prepared by this section are:

Can the trace be reconstructed from O?

Can the deterministic target be recovered from O?

Is O sufficient for the stochastic target law?

The next section turns these questions into posterior reconstruction profiles. Scalar entropy and conditional mutual information will then appear as shadows of those profiles rather than as unstructured standalone quantities.

3 Compressed Observations and Posterior Reconstruction Profiles

Section 2 defined the finite presentation

$$\mathfrak{A} = (T, \{\mathcal{S}_t\}_{t=0}^{T}, \mu_0, \{K_t\}_{t=0}^{T-1}, \mathcal{O}, c, \mathcal{Y}, L), \qquad O = c(S_{0:T}),$$

and the induced joint law

$$P_{\mathfrak{A}}(S_{0:T}, O, Y).$$

The present section introduces the profile layer attached to this law. This layer is essential for two reasons.

First, a conditional entropy value such as

$$H(S_{\mathrm{int}} \mid O)$$

is only a scalar summary of a family of posterior distributions. It detects exact reconstruction through its zero set, but it does not record how the hidden traces are distributed over observation fibers. Second, the isomorphic observed target behavior can coexist with different hidden posterior trace profiles. Therefore CEOT IV must keep the posterior profile as the primary object and treat scalar entropy as a numerical shadow of that profile.

This is the same methodological discipline used in CEOT III for finite marked stochastic diagrams. There, the observed marking induced a hidden-completion posterior profile. Here, the deterministic compressed observation $O = c(S_{0:T})$ induces complete-trace, hidden-trace, bridge, and target profiles.

3.1 General posterior reconstruction profile

Let A and O be finite random variables defined on the same probability space. In CEOT IV, A will usually be a trace random variable, a hidden internal trace, a bridge-internal trace, a memory variable, or a target variable. The observation variable is the compressed observation $O = c(S_{0:T})$, or an augmented observation such as (O, S_0, S_T).

Definition 3.1 (Posterior reconstruction profile)**.** The posterior reconstruction profile of A relative to O is

$$\mathsf{Post}_{A|O} := \left(P_O, \{P_{A|O=o}\}_{o\in\mathrm{supp}(O)}\right).$$

The associated scalar entropy obstruction is

$$\mathrm{Ob}(A \mid O) := H(A \mid O).$$

The profile contains more data than the scalar entropy. It records the observation law P_O and, for every positive-support observation fiber, the posterior distribution of the hidden or target variable on that fiber. The scalar entropy is recovered from the profile by

$$H(A \mid O) = \sum_{o\in\mathrm{supp}(O)} P_O(o)\, H(A \mid O = o).$$

Thus $H(A \mid O)$ is a functional of $\mathsf{Post}_{A|O}$, but the reverse implication is false: the scalar does not determine the profile.

Remark 3.2 (Support-relative convention)*.* The family $\{P_{A|O=o}\}$ is indexed only by $o \in \mathrm{supp}(O)$. Values outside the support of O have no reconstruction meaning. This is why all decoder and zero-obstruction statements in later sections are formulated on $\mathrm{supp}(O)$ rather than on an arbitrary ambient observation set.

Definition 3.3 (Profile support fibers)**.** For $o \in \mathrm{supp}(O)$, the positive posterior fiber of A over o is

$$\mathcal{F}_{A|O}(o) := \{a \in \mathrm{supp}(A) : P(A = a \mid O = o) > 0\}.$$

The profile is pointwise deterministic over o if $\mathcal{F}_{A|O}(o)$ is a singleton.

The zero-obstruction criteria proved in Section 4 are exactly the global form of pointwise determinism: $H(A \mid O) = 0$ if and only if every positive posterior fiber is a singleton. Section 3 records the profile object; Section 4 proves the corresponding reconstruction criteria.

3.2 Complete-trace profile

The complete-trace profile asks what entire algorithmic histories remain possible after the compressed observation is known. It is obtained from Definition 3.1 by setting

$$A = S_{0:T}.$$

Definition 3.4 (Complete-trace posterior profile)**.** The complete-trace posterior profile of a presentation $\mathfrak{A}$ relative to O is

$$\mathsf{Post}_{\mathrm{complete}}(\mathfrak{A}; O) := \mathsf{Post}_{S_{0:T}|O} = \Big(P_O, \{P_{S_{0:T}|O=o}\}_{o\in\mathrm{supp}(O)}\Big).$$

Its scalar complete-trace obstruction is

$$\mathrm{Ob}_{\mathrm{complete}}(O) := H(S_{0:T} \mid O).$$

The posterior distribution $P_{S_{0:T}|O=o}$ is supported on the trace fiber

$$c^{-1}(o) \cap \mathrm{supp}(S_{0:T}).$$

Thus the complete-trace profile records not only the size of an observation fiber, but also the probability weights inherited from the trace law.

Remark 3.5 (Complete reconstruction is usually too strong)*.* Complete-trace reconstruction is the strongest reconstruction demand considered in this paper. It asks for recovery of every state $S_0, \dots, S_T$. In algorithmic compression, this is often more than the task requires. Later sections therefore compare $\mathrm{Ob}_{\mathrm{complete}}(O)$ with hidden-trace, bridge, deterministic-target, and stochastic-target obstructions.

3.3 Hidden-trace profile

The hidden-trace profile is the main algorithmic trace profile of CEOT IV. It ignores the endpoints as objects to be reconstructed and focuses on the internal computational history

$$S_{\mathrm{int}}.$$

This is the algorithmic analogue of the interior or bridge-hidden variable appearing in CEOT I and CEOT II.

Definition 3.6 (Hidden-trace posterior profile)**.** The hidden-trace posterior profile of $\mathfrak{A}$ relative to O is

$$\mathsf{Post}_{\mathrm{trace}}(\mathfrak{A}; O) := \mathsf{Post}_{S_{\mathrm{int}}|O} = \Big(P_O, \{P_{S_{\mathrm{int}}|O=o}\}_{o\in\mathrm{supp}(O)}\Big).$$

Its scalar hidden-trace obstruction is

$$\mathrm{Ob}_{\mathrm{trace}}(O) := H(S_{\mathrm{int}} \mid O).$$

By Convention 2.9, this definition also covers the degenerate horizon $T = 1$. In that case S_{int} is the singleton variable $*_{\text{int}}$, so the hidden-trace obstruction is automatically zero. The genuinely internal reconstruction theory begins when $T \geq 2$.

The hidden-trace profile is not determined by the complete observation law P_O alone. It depends on the trace law, the transition kernels, and the compression map. Two presentations may induce the same law of O and even the same joint law of (O, Y) while inducing different hidden-trace profiles. This is the trace-level non-completeness principle proved in the later section on separation and non-completeness.

Remark 3.7 (Why hidden trace is distinguished from complete trace)*.* The complete trace includes endpoints S_0 and S_T. In many algorithmic situations the endpoints are already part of the observed or task-relevant data, while the internal computation is the uncertain object. This motivates separating

$$H(S_{\text{int}} \mid O)$$

from

$$H(S_{0:T} \mid O).$$

The latter dominates the former, but it may overstate the obstruction relevant to internal reconstruction.

3.4 Bridge profile

The bridge profile conditions not only on the compressed observation O, but also on the endpoints (S_0, S_T). It measures the remaining uncertainty in the internal trace after the compressed observation and endpoints are known.

Definition 3.8 (Bridge posterior profile)**.** The bridge posterior profile of $\mathfrak{A}$ relative to O is

$$\mathsf{Post}_{\text{bridge}}(\mathfrak{A}; O) := \mathsf{Post}_{S_{\text{int}} \mid O, S_0, S_T}$$

that is,

$$\mathsf{Post}_{\text{bridge}}(\mathfrak{A}; O) = \Big(P_{O,S_0,S_T}, \{P_{S_{\text{int}} \mid O=o, S_0=s_0, S_T=s_T}\}_{(o,s_0,s_T) \in \text{supp}(O,S_0,S_T)}\Big).$$

Its scalar bridge obstruction is

$$\text{Ob}_{\text{bridge}}(O) := H(S_{\text{int}} \mid O, S_0, S_T).$$

Proposition 3.9 (Degenerate bridge convention)**.** *If $T = 1$, then*

$$S_{\text{int}} = *_{\text{int}}$$

is deterministic. Consequently,

$$\text{Ob}_{\text{trace}}(O) = H(S_{\text{int}} \mid O) = 0$$

and

$$\text{Ob}_{\text{bridge}}(O) = H(S_{\text{int}} \mid O, S_0, S_T) = 0.$$

Moreover, both the hidden-trace profile and the bridge profile are singleton posterior profiles.

Proof. By Convention 2.9, $\mathcal{S}_{\text{int}} = \{*_{\text{int}}\}$ when $T = 1$. Therefore S_{int} is deterministic under every positive-probability conditioning event. Its conditional entropy is zero given O, and also zero given (O, S_0, S_T). The corresponding posterior distributions are Dirac masses at $*_{\text{int}}$. $\square$

The bridge profile is the closest CEOT IV analogue of the bridge reconstruction profile of CEOT I. It keeps the algorithmic compression variable O but also supplies endpoint data. Since conditioning can only reduce entropy, the bridge obstruction will satisfy

$$\mathrm{Ob}_{\mathrm{bridge}}(O) \leq \mathrm{Ob}_{\mathrm{trace}}(O).$$

The full hierarchy is proved in Section 4.

Example 3.10 (Endpoint-only observation)**.** If O is trivial and (S_0, S_T) are supplied separately, then

$$\mathrm{Ob}_{\mathrm{bridge}}(O) = H(S_{\mathrm{int}} \mid S_0, S_T),$$

which is the ordinary endpoint-conditioned interior obstruction. If instead $O = (S_0, S_T)$, then

$$\mathrm{Ob}_{\mathrm{trace}}(O) = H(S_{\mathrm{int}} \mid S_0, S_T),$$

so hidden-trace obstruction relative to endpoint observation coincides with the bridge obstruction.

3.5 Target profile

The target profile records what the compressed observation says about the task target. Its role depends on whether the target is deterministic or stochastic.

Definition 3.11 (Target posterior profile)**.** The target posterior profile of $\mathfrak{A}$ relative to O is

$$\mathsf{Post}_Y(\mathfrak{A}; O) := \mathsf{Post}_{Y|O} = \Big(P_O, \{P_{Y|O=o}\}_{o\in\mathrm{supp}(O)}\Big).$$

For a support-relative deterministic target $Y = f_+(S_{0:T})$, its scalar deterministic-target obstruction is

$$\mathrm{Ob}^{Y}_{\mathrm{det}}(O) := H(Y \mid O).$$

If Y is deterministic, then $\mathsf{Post}_Y(\mathfrak{A}; O)$ is the object whose pointwise degeneracy determines target recovery. In that case

$$H(Y \mid O) = 0$$

if and only if every positive observation fiber carries a single target value.

For stochastic targets, however, $\mathsf{Post}_{Y|O}$ alone does not express whether O has preserved all target-relevant information carried by the full trace. The relevant comparison is between

$$P_{Y|S_{0:T}} \quad \text{and} \quad P_{Y|O}.$$

The obstruction is

$$\mathrm{Ob}^{Y}_{\mathrm{law}}(O) := I(Y; S_{0:T} \mid O).$$

It vanishes exactly when

$$P(Y \mid S_{0:T}) = P(Y \mid O)$$

on support, using Lemma 2.15.

Definition 3.12 (Stochastic sufficiency profile)**.** For a stochastic target with support-restricted kernel L_+, the stochastic sufficiency profile relative to O is the collection

$$\mathsf{Suff}_Y(\mathfrak{A}; O) := \Big(P_O, \{P_{S_{0:T}|O=o}\}_{o\in\mathrm{supp}(O)}, L_+\Big).$$

Its scalar sufficiency obstruction is

$$I(Y; S_{0:T} \mid O).$$

This definition deliberately includes both the posterior trace profile and the support-restricted target kernel L_+. For stochastic sufficiency, knowing only the marginal profile $\mathsf{Post}_{Y|O}$ is not enough to determine whether the full trace carries additional information about Y after O is known. Ambient values of L outside positive trace support are not part of this profile.

Remark 3.13 (Deterministic targets are included). When $Y = f_+(S_{0:T})$ is support-relative deterministic, the stochastic sufficiency obstruction reduces to the deterministic target obstruction:

$$I(Y; S_{0:T} \mid O) = H(Y \mid O).$$

Thus the deterministic target profile is the degenerate-kernel case of the stochastic sufficiency framework.

3.6 Scalar entropy is only a shadow

The profile-first principle can now be stated precisely. The scalar entropy obstruction is obtained from the posterior profile by averaging the fiber entropies:

$$H(A \mid O) = \sum_{o \in \mathrm{supp}(O)} P_O(o)\, H(P_{A|O=o}).$$

Therefore the scalar is invariant under profile isomorphism, but it is not a complete invariant of the profile.

Proposition 3.14 (Scalar non-completeness)**.** *There exist finite pairs* (A, O) *and* (A', O') *such that*

$$H(A \mid O) = H(A' \mid O')$$

but

$$\mathsf{Post}_{A|O}$$

and

$$\mathsf{Post}_{A'|O'}$$

are not isomorphic as posterior reconstruction profiles.

Proof. Let O be uniformly distributed on two values and set $A = O$. Then every positive posterior fiber of A over O is a singleton, hence

$$H(A \mid O) = 0.$$

Let O' be uniformly distributed on three values and set $A' = O'$. Again every positive posterior fiber is a singleton, hence

$$H(A' \mid O') = 0.$$

The two profiles have the same scalar entropy, but they are not isomorphic because their observation supports have different cardinalities. Thus the scalar entropy value does not determine the posterior profile. □

Theorem 3.15 (Strong scalar non-completeness)**.** *There exist finite posterior profiles*

$$\mathsf{Post}_{A|O} \quad \textit{and} \quad \mathsf{Post}_{A'|O'}$$

such that:

(i) $\operatorname{supp}(O)$ *and* $\operatorname{supp}(O')$ *have the same cardinality.*

(ii) $\operatorname{supp}(A)$ *and* $\operatorname{supp}(A')$ *have the same cardinality.*

(iii)

$$H(A \mid O) = H(A' \mid O') > 0.$$

(iv) The two posterior profiles are not isomorphic.

Proof. Let O and O' both be constant observations, and let A and A' both take values in a three-point set. Then the posterior profile is just one probability distribution on three points.

Choose an interior probability vector

$$p = (p_1, p_2, p_3)$$

with $p_i > 0$, $p_1 + p_2 + p_3 = 1$, and p not uniform. The entropy function on the interior of the two-simplex is smooth and nonconstant. Away from the uniform point its gradient is nonzero in some tangent direction. Hence, by the implicit function theorem, the level set

$$\{q \in \Delta_3^\circ : H(q) = H(p)\}$$

contains infinitely many points near p. Only finitely many of these are permutations of p. Choose q on the same entropy level which is not a permutation of p.

Let A have law p, let A' have law q, and let O, O' be constant. Then

$$H(A \mid O) = H(p) = H(q) = H(A' \mid O') > 0.$$

The posterior profiles are not isomorphic because, with constant observations, profile isomorphism is exactly equality of the posterior distributions up to permutation. Since q is not a permutation of p, the profiles are not isomorphic. □

Corollary 3.16 (Hidden-trace scalar non-completeness)**.** *There exist finite deterministic trace-compression presentations* $\mathfrak{A}$ *and* $\mathfrak{A}'$ *with the same observation support size and the same hidden-trace support size such that*

$$\mathrm{Ob}^{\mathfrak{A}}_{\mathrm{trace}}(O) = \mathrm{Ob}^{\mathfrak{A}'}_{\mathrm{trace}}(O') > 0,$$

but

$$\mathsf{Post}^{\mathfrak{A}}_{S_{\mathrm{int}}|O} \ncong \mathsf{Post}^{\mathfrak{A}'}_{S'_{1:T-1}|O'}.$$

Proof. Realize the two posterior distributions p and q from Theorem 3.15 as hidden-trace laws with constant observation and trivial endpoints. For example, take $T = 2$,

$$S_0 = S_2 = *, \qquad S_1 \in \{1, 2, 3\},$$

with law p in the first presentation and law q in the second. Let O and O' be constant. Then the hidden-trace obstruction is exactly the entropy of the hidden state:

$$\mathrm{Ob}_{\mathrm{trace}}(O) = H(S_1).$$

The scalar values agree by construction, but the posterior profiles are not isomorphic. □

Remark 3.17 (Scalar equality is not structural equality). Even positive scalar obstruction values do not classify posterior uncertainty. Two compressions may lose the same numerical amount of information while losing it in different posterior shapes. CEOT IV therefore treats scalar entropy as a shadow of the posterior profile, not as the profile itself.

The preceding propositions are intentionally explicit. Their role is to prevent overreading scalar entropy. CEOT IV uses scalar entropy and conditional mutual information because their zero sets and monotonicity properties give useful obstruction criteria. It does not claim that scalar entropy classifies posterior reconstruction profiles.

A stronger non-completeness phenomenon appears later: even the observed target law $P_{O,Y}$, considered up to law isomorphism, does not determine the hidden-trace obstruction $H(S_{\mathrm{int}} \mid O)$. That result depends on the algorithmic presentation, not merely on abstract pairs of finite variables.

3.7 Profile isomorphism at the profile layer

The formal profile groupoid is developed in Section 7. At the elementary posterior-profile level, the relevant idea is support-relative relabeling. If $A : \Omega \to \mathcal{A}$ and $A' : \Omega' \to \mathcal{A}'$ are finite random variables and O, O' are observations, then posterior profiles

$$\mathsf{Post}_{A|O} \quad \text{and} \quad \mathsf{Post}_{A'|O'}$$

are isomorphic when there are bijections

$$\alpha : \mathrm{supp}(A) \to \mathrm{supp}(A'), \qquad \beta : \mathrm{supp}(O) \to \mathrm{supp}(O')$$

such that

$$P_{O'}(\beta(o)) = P_O(o)$$

and

$$P_{A'|O'=\beta(o)}(\alpha(a)) = P_{A|O=o}(a)$$

for all $o \in \mathrm{supp}(O)$ and all $a \in \mathrm{supp}(A)$.

Under posterior-profile isomorphism,

$$H(A \mid O) = H(A' \mid O').$$

The converse is false by Proposition 3.14. Thus the correct finite categorical direction is

$$\text{typed presentation} \longrightarrow \text{posterior profile object} \longrightarrow \text{scalar obstruction},$$

not the reverse.

3.8 Section summary

This section introduced the posterior profile layer for CEOT IV. The complete-trace profile records the posterior law of $S_{0:T}$ given O. The hidden-trace profile records the posterior law of S_{int} given O. The bridge profile records the posterior law of S_{int} given (O, S_0, S_T). The target profile records the posterior law of Y given O, while stochastic target sufficiency also requires comparison with the full trace through the kernel $L : \mathcal{S}_{0:T} \rightsquigarrow \mathcal{Y}$.

The scalar quantities

$$H(S_{0:T} \mid O), \quad H(S_{\mathrm{int}} \mid O), \quad H(S_{\mathrm{int}} \mid O, S_0, S_T), \quad H(Y \mid O), \quad I(Y; S_{0:T} \mid O)$$

are therefore shadows of profile-level objects. The next section proves the zero criteria and hierarchy for the trace reconstruction shadows.

4 Trace Reconstruction Obstructions

Section 3 introduced the posterior profile layer. We now extract the scalar reconstruction obstructions and prove their support-relative zero criteria. The point of this section is deliberately narrow: it concerns reconstruction of traces and trace fragments from a deterministic compressed observation. Target-relative recovery and stochastic sufficiency are treated later. This separation prevents the central error that CEOT IV is designed to avoid, namely the identification of hidden trace loss with algorithmic failure.

Throughout this section, let

$$\mathfrak{A} = (T, \{S_t\}_{t=0}^{T}, \mu_0, \{K_t\}_{t=0}^{T-1}, c, O, Y, L)$$

be a finite deterministic trace-compression presentation, and let

$$O = c(S_{0:T}).$$

All statements are support-relative: decoders are required only on the positive-probability support of the relevant observed variable.

4.1 The finite zero-entropy reconstruction lemma

We first record the elementary finite lemma used repeatedly below.

Lemma 4.1 (Finite zero-entropy reconstruction)**.** *Let A and B be finite random variables. Then*

$$H(A \mid B) = 0$$

if and only if there exists a map

$$r : \operatorname{supp}(B) \to \operatorname{supp}(A)$$

such that

$$A = r(B) \quad \textit{a.s.}$$

Equivalently, for every $b \in \operatorname{supp}(B)$, the posterior distribution $P_{A|B=b}$ is a Dirac mass.

Proof. Since A and B are finite,

$$H(A \mid B) = \sum_{b \in \operatorname{supp}(B)} P_B(b)\, H(A \mid B = b).$$

Each summand is nonnegative. Hence the conditional entropy vanishes if and only if

$$H(A \mid B = b) = 0$$

for every $b \in \operatorname{supp}(B)$. A finite distribution has entropy zero if and only if it is a Dirac mass. Thus for every $b \in \operatorname{supp}(B)$ there is a unique value $r(b)$ with

$$P(A = r(b) \mid B = b) = 1.$$

This defines the required support-relative decoder. The converse is immediate: if $A = r(B)$ almost surely, then every positive posterior fiber of A over B is a singleton, so $H(A \mid B) = 0$. $\square$

Remark 4.2 (Why support-relative decoders are sufficient). No value of a decoder outside $\operatorname{supp}(B)$ affects any obstruction in this paper. Thus CEOT IV uses support-relative reconstruction throughout, in the same finite null-set discipline used in CEOT I–III. This is not a measurability subtlety; it is simply the correct finite formulation.

Convention 4.3 (Support-relative decoder convention)**.** Every decoder in CEOT IV is support-relative unless explicitly stated otherwise. Thus a decoder from an observed variable B to a trace-derived or target variable A is a map

$$r : \operatorname{supp}(B) \to \operatorname{supp}(A)$$

satisfying

$$A = r(B) \quad \text{a.s.}$$

No value of r outside $\operatorname{supp}(B)$ is part of the obstruction data. If an ambient map

$$\widetilde{r} : \mathcal{B} \to \mathcal{A}$$

is desired, it may be chosen arbitrarily on $\mathcal{B} \setminus \operatorname{supp}(B)$, and all CEOT IV obstruction values remain unchanged.

Lemma 4.4 (Ambient extension is obstruction-invisible)**.** *Let A, B be finite random variables with supports contained in ambient finite sets $\mathcal{A}, \mathcal{B}$. If*

$$r : \operatorname{supp}(B) \to \operatorname{supp}(A)$$

is a support-relative decoder with

$$A = r(B) \quad \textit{a.s.},$$

then any extension

$$\widetilde{r} : \mathcal{B} \to \mathcal{A}, \qquad \widetilde{r}|_{\operatorname{supp}(B)} = r,$$

also satisfies

$$A = \widetilde{r}(B) \quad \textit{a.s.}$$

Conversely, any ambient decoder $\widetilde{r}$ that reconstructs A almost surely restricts to a support-relative decoder on $\operatorname{supp}(B)$*. Therefore CEOT IV zero-obstruction criteria depend only on support-relative decoders.*

Proof. Since $P(B \in \operatorname{supp}(B)) = 1$, values of $\widetilde{r}$ outside $\operatorname{supp}(B)$ are never evaluated with positive probability. Hence $\widetilde{r}(B) = r(B)$ almost surely. The converse is immediate by restriction to the positive support. □

Zero-theorem template. Every exact zero criterion below has the support-relative form

$$H(A \mid B) = 0 \quad \Longleftrightarrow \quad \exists r : \operatorname{supp}(B) \to \operatorname{supp}(A) \text{ with } A = r(B) \text{ a.s.}$$

When the text says that A is recovered, reconstructed, or determined by B, it always means this support-relative almost-sure statement. Null observation fibers impose no reconstruction constraints.

Remark 4.5 (Null fibers impose no constraints). If $b \notin \operatorname{supp}(B)$, then the ambient fiber over b may be nonempty, but it has probability zero under the declared law. CEOT IV exact reconstruction criteria impose no constraints on such null values. This is why the finite formulation is support-relative rather than ambient-universal.

4.2 Static obstruction

The general observation model $O = c(S_{0:T})$ allows a global finite summary of the trace. The roadmap case, however, is coordinatewise compression:

$$O = \bar{S}_{0:T} = (\bar{S}_0, \ldots, \bar{S}_T), \qquad \bar{S}_t = q_t(S_t).$$

In that case one can define a one-time obstruction.

Definition 4.6 (Static reconstruction obstruction)**.** In the coordinatewise compressed-trace case, the static obstruction at time t is

$$\mathrm{Ob}_{\mathrm{static}}(t) := H(S_t \mid \bar{S}_t).$$

The static obstruction asks whether the full algorithmic state at one time is recoverable from its one-time compressed state. It does not, by itself, measure joint trace reconstruction. The joint trace may contain temporal correlations, and the one-time posterior laws do not determine the full posterior law of $S_{0:T}$ given $\bar{S}_{0:T}$.

Proposition 4.7 (Static zero criterion)**.** *In the coordinatewise case, for a fixed time t,*

$$\mathrm{Ob}_{\mathrm{static}}(t) = 0$$

if and only if there exists a support-relative decoder

$$r_t : \mathrm{supp}(\bar{S}_t) \to \mathrm{supp}(S_t)$$

such that

$$S_t = r_t(\bar{S}_t) \quad \textit{a.s.}$$

Proof. Apply Lemma 4.1 with $A = S_t$ and $B = \bar{S}_t$. $\square$

Proposition 4.8 (Static zero implies complete coordinatewise zero)**.** *In the coordinatewise compression model*

$$O = \bar{S}_{0:T}, \qquad \bar{S}_t = q_t(S_t),$$

if

$$H(S_t \mid \bar{S}_t) = 0 \quad \textit{for every } t = 0, \ldots, T,$$

then

$$H(S_{0:T} \mid \bar{S}_{0:T}) = 0.$$

Proof. For each t, the support-relative zero criterion gives a decoder

$$r_t : \mathrm{supp}(\bar{S}_t) \to \mathrm{supp}(S_t)$$

with

$$S_t = r_t(\bar{S}_t) \quad \text{a.s.}$$

Hence

$$S_{0:T} = (r_0(\bar{S}_0), \ldots, r_T(\bar{S}_T)) \quad \text{a.s.}$$

so the complete coordinatewise trace is support-relatively decoded from $\bar{S}_{0:T}$. $\square$

Example 4.9 (Complete coordinatewise reconstruction without local static zero)**.** Let $T = 1$, let $B \sim \text{Bernoulli}(1/2)$, and set

$$S_0 = B, \qquad S_1 = B.$$

Define coordinatewise compressed states by

$$\bar{S}_0 = B, \qquad \bar{S}_1 = *.$$

Then $\bar{S}_{0:1} = (B, *)$ recovers the complete trace (B, B), hence

$$H(S_{0:1} \mid \bar{S}_{0:1}) = 0.$$

However,

$$H(S_1 \mid \bar{S}_1) = H(B \mid *) = \log_b 2 > 0.$$

Thus complete coordinatewise reconstruction does not imply local static zero at every time.

Remark 4.10 (Static obstruction is local and coordinatewise)*.* The static obstruction $H(S_t \mid \bar{S}_t)$ is a local coordinatewise quantity. It is not a synonym for hidden-trace obstruction $H(S_{\text{int}} \mid O)$ or complete-trace obstruction $H(S_{0:T} \mid O)$. Static zero at every time is a sufficient condition for complete coordinatewise reconstruction, but it is not necessary when information about S_t is carried by other compressed time coordinates.

4.3 Complete-trace obstruction

The strongest reconstruction question asks whether the whole trace is recoverable.

Definition 4.11 (Complete-trace obstruction)**.** The complete-trace obstruction associated to O is

$$\text{Ob}_{\text{complete}}(O) := H(S_{0:T} \mid O).$$

This obstruction is intentionally stronger than most algorithmic tasks require. It asks for complete recovery of every state in the computation, including endpoints and all internal states. Its vanishing gives a sufficient condition for any deterministic target recovery, but later sections show that it is far from necessary.

Theorem 4.12 (Complete-trace zero criterion)**.** *One has*

$$H(S_{0:T} \mid O) = 0$$

if and only if there exists a support-relative decoder

$$R_{\text{complete}} : \text{supp}(O) \to \text{supp}(S_{0:T})$$

such that

$$R_{\text{complete}}(O) = S_{0:T} \quad \textit{a.s.}$$

Proof. Apply Lemma 4.1 with $A = S_{0:T}$ and $B = O$. □

Corollary 4.13 (Complete reconstruction determines every trace function)**.** *If*

$$\text{Ob}_{\text{complete}}(O) = 0,$$

then every deterministic finite variable of the form

$$Z = h(S_{0:T})$$

is also reconstructible from O:

$$H(Z \mid O) = 0.$$

Proof. By Theorem 4.12, $S_{0:T} = R_{\mathrm{complete}}(O)$ almost surely. Hence

$$Z = h(S_{0:T}) = h(R_{\mathrm{complete}}(O))$$

almost surely, so Z is a function of O and $H(Z \mid O) = 0$. $\square$

4.4 Hidden-trace obstruction

The main roadmap obstruction of CEOT IV is not necessarily complete-trace reconstruction. It is hidden internal trace reconstruction.

Definition 4.14 (Hidden-trace obstruction)**.** The hidden-trace obstruction associated to O is

$$\mathrm{Ob}_{\mathrm{trace}}(O) := H(S_{\mathrm{int}} \mid O).$$

The hidden-trace obstruction asks whether the internal computational states are recoverable from the compressed observation. It is the algorithmic analogue of the bridge/interior obstructions from CEOT I–II and the hidden-completion obstruction from CEOT III. It is also the quantity that will later be separated from target recovery: internal history can be hidden while the target remains fully preserved.

Theorem 4.15 (Hidden-trace zero criterion)**.** *One has*

$$H(S_{\mathrm{int}} \mid O) = 0$$

if and only if there exists a support-relative decoder

$$R_{\mathrm{trace}} : \mathrm{supp}(O) \to \mathrm{supp}(S_{\mathrm{int}})$$

such that

$$R_{\mathrm{trace}}(O) = S_{\mathrm{int}} \quad a.s.$$

Proof. Apply Lemma 4.1 with $A = S_{\mathrm{int}}$ and $B = O$. $\square$

Remark 4.16 (Hidden trace versus complete trace)*.* The condition

$$H(S_{\mathrm{int}} \mid O) = 0$$

need not imply

$$H(S_{0:T} \mid O) = 0.$$

For example, if O records all internal states but forgets a random endpoint, then the hidden internal trace is reconstructible while the complete trace is not. Conversely, complete-trace reconstruction always implies hidden-trace reconstruction by deterministic projection.

4.5 Bridge obstruction

The bridge version conditions additionally on endpoints. It is included to preserve the link with CEOT I and CEOT II, where endpoint/interior separation is one of the central themes.

Definition 4.17 (Bridge obstruction)**.** The bridge obstruction associated to O is

$$\mathrm{Ob}_{\mathrm{bridge}}(O) := H(S_{\mathrm{int}} \mid O, S_0, S_T).$$

Thus $\mathrm{Ob}_{\mathrm{bridge}}(O)$ measures internal trace uncertainty after observing both the compressed observation and the endpoints. If endpoints are already functions of O, then this equals the hidden-trace obstruction. If endpoints contain additional information not present in O, the bridge obstruction can be strictly smaller.

Theorem 4.18 (Bridge zero criterion)**.** *One has*

$$H(S_{\mathrm{int}} \mid O, S_0, S_T) = 0$$

if and only if there exists a support-relative decoder

$$R_{\mathrm{bridge}} : \mathrm{supp}(O, S_0, S_T) \to \mathrm{supp}(S_{\mathrm{int}})$$

such that

$$R_{\mathrm{bridge}}(O, S_0, S_T) = S_{\mathrm{int}} \quad \textit{a.s.}$$

Proof. Apply Lemma 4.1 with $A = S_{\mathrm{int}}$ and $B = (O, S_0, S_T)$. □

Example 4.19 (Endpoint information can reduce internal uncertainty)**.** Let $T = 2$, let S_0, S_2 be independent fair bits, and let

$$S_1 = S_0 \oplus S_2,$$

where $\oplus$ denotes addition modulo two. Let O be the trivial observation. Then

$$H(S_1 \mid O) = H(S_1) = \log_b 2,$$

but

$$H(S_1 \mid O, S_0, S_2) = 0.$$

Thus bridge reconstruction can hold even when hidden-trace reconstruction from the compressed observation alone fails.

4.6 Memory obstruction

A common algorithmic reason to enlarge a state space is that a visible state does not contain enough information about a memory variable, signature, mode, or lifted label. CEOT IV records the reconstruction part of this phenomenon here and postpones target relevance to the later section on visible-trace insufficiency.

The formal definition of the memory reconstruction obstruction is deferred to Definition 8.3, after the lifted-state notation has been fixed. At this point the relevant quantity is the ordinary reconstruction entropy

$$H(M_{0:T} \mid V_{0:T}),$$

whenever a coordinate decomposition

$$S_{0:T} = (V_{0:T}, M_{0:T})$$

has been specified. Target relevance is a different question, measured later by

$$I(Y; M_{0:T} \mid V_{0:T}).$$

Proposition 4.20 (Memory zero criterion preview)**.** *One has*

$$H(M_{0:T} \mid V_{0:T}) = 0$$

if and only if there exists a support-relative decoder

$$R_{\mathrm{mem}} : \mathrm{supp}(V_{0:T}) \to \mathrm{supp}(M_{0:T})$$

such that

$$R_{\mathrm{mem}}(V_{0:T}) = M_{0:T} \quad a.s.$$

Proof. Apply Lemma 4.1 with $A = M_{0:T}$ and $B = V_{0:T}$. □

Remark 4.21 (Memory hiddenness is not memory relevance)*.* The obstruction $H(M_{0:T} \mid V_{0:T})$ only says whether the memory trace is reconstructible from the visible trace. It does not say whether that memory is relevant to a chosen target. Target relevance is measured later by

$$I(Y; M_{0:T} \mid V_{0:T}).$$

This distinction is essential: hidden memory may be irrelevant to the task, or it may carry target-relevant information that the visible trace loses.

4.7 Local zero implies complete-trace zero

The coordinatewise case has one useful implication: if every one-time compression is lossless, then the whole coordinatewise compressed trace is lossless.

Theorem 4.22 (Local zero implies complete-trace zero)**.** *Assume the coordinatewise compressed-trace model*

$$O = \bar{S}_{0:T} = (\bar{S}_0, \dots, \bar{S}_T), \qquad \bar{S}_t = q_t(S_t).$$

If

$$H(S_t \mid \bar{S}_t) = 0 \quad \textit{for every } t = 0, \dots, T,$$

then

$$H(S_{0:T} \mid \bar{S}_{0:T}) = 0.$$

Proof. By Proposition 4.7, for each t there exists a support-relative decoder r_t such that

$$S_t = r_t(\bar{S}_t) \quad \text{a.s.}$$

Define

$$R(\bar{s}_0, \dots, \bar{s}_T) := (r_0(\bar{s}_0), \dots, r_T(\bar{s}_T))$$

on the positive support of $\bar{S}_{0:T}$. Then

$$R(\bar{S}_{0:T}) = S_{0:T} \quad \text{a.s.}$$

Theorem 4.12 gives

$$H(S_{0:T} \mid \bar{S}_{0:T}) = 0.$$

□

Remark 4.23 (No converse from marginal profiles). The theorem says that pointwise lossless coordinate compression implies lossless trace compression. It does not say that the scalar values

$$H(S_t \mid \bar{S}_t), \quad t = 0, \dots, T,$$

determine

$$H(S_{0:T} \mid \bar{S}_{0:T}).$$

The latter depends on the joint posterior law of the whole trace, not merely on one-time marginal posterior laws.

4.8 The obstruction hierarchy

The bridge, hidden-trace, and complete-trace obstructions form a basic hierarchy.

Theorem 4.24 (Bridge/trace/complete hierarchy)**.** *For every finite deterministic compressed observation $O = c(S_{0:T})$,*

$$\mathrm{Ob}_{\mathrm{bridge}}(O) \leq \mathrm{Ob}_{\mathrm{trace}}(O) \leq \mathrm{Ob}_{\mathrm{complete}}(O).$$

Equivalently,

$$H(S_{\mathrm{int}} \mid O, S_0, S_T) \leq H(S_{\mathrm{int}} \mid O) \leq H(S_{0:T} \mid O).$$

Proof. If $T = 1$, then $\mathrm{Ob}_{\mathrm{trace}}(O) = \mathrm{Ob}_{\mathrm{bridge}}(O) = 0$ by Proposition 3.9, and the hierarchy is immediate. Assume $T \geq 2$ for the non-degenerate case.

The first inequality is monotonicity of conditional entropy under additional conditioning:

$$H(S_{\mathrm{int}} \mid O, S_0, S_T) \leq H(S_{\mathrm{int}} \mid O).$$

For the second inequality, use the chain rule:

$$H(S_{0:T} \mid O) = H(S_{\mathrm{int}}, S_0, S_T \mid O).$$

Expanding the right hand side gives

$$H(S_{\mathrm{int}} \mid O) + H(S_0, S_T \mid O, S_{\mathrm{int}}).$$

The second term is nonnegative, hence

$$H(S_{\mathrm{int}} \mid O) \leq H(S_{0:T} \mid O).$$

□

Corollary 4.25 (Zero implications)**.** *The following implications hold:*

$$\mathrm{Ob}_{\mathrm{complete}}(O) = 0 \quad \Longrightarrow \quad \mathrm{Ob}_{\mathrm{trace}}(O) = 0 \quad \Longrightarrow \quad \mathrm{Ob}_{\mathrm{bridge}}(O) = 0.$$

The converses can fail.

Proof. The implications follow from Theorem 4.24. The precise non-converse witnesses are recorded in Theorem 4.26 below. □

Theorem 4.26 (Strictness and non-converses in the bridge/trace/complete hierarchy)**.** *The implications*

$$\mathrm{Ob}_{\mathrm{complete}}(O) = 0 \Longrightarrow \mathrm{Ob}_{\mathrm{trace}}(O) = 0 \Longrightarrow \mathrm{Ob}_{\mathrm{bridge}}(O) = 0$$

cannot be reversed in general. More sharply, each inequality in

$$\mathrm{Ob}_{\mathrm{bridge}}(O) \le \mathrm{Ob}_{\mathrm{trace}}(O) \le \mathrm{Ob}_{\mathrm{complete}}(O)$$

can be strict in finite deterministic trace-compression presentations.

Proof. All examples use $T = 2$ and a constant observation $O = *$.

First, let $X \sim \mathrm{Bernoulli}(1/2)$ and set

$$S_0 = X, \qquad S_1 = X, \qquad S_2 = X.$$

Then

$$\mathrm{Ob}_{\mathrm{bridge}}(O) = H(S_1 \mid O, S_0, S_2) = H(X \mid X, X) = 0,$$

but

$$\mathrm{Ob}_{\mathrm{trace}}(O) = H(S_1 \mid O) = H(X) = \log_b 2.$$

Thus bridge-zero does not imply hidden-trace-zero.

Second, let $X \sim \mathrm{Bernoulli}(1/2)$ and set

$$S_0 = X, \qquad S_1 = *, \qquad S_2 = *.$$

Then

$$\mathrm{Ob}_{\mathrm{trace}}(O) = H(S_1 \mid O) = 0,$$

but

$$\mathrm{Ob}_{\mathrm{complete}}(O) = H(S_0, S_1, S_2 \mid O) = H(X) = \log_b 2.$$

Thus hidden-trace-zero does not imply complete-trace-zero.

Finally, to witness simultaneous strictness, let X, Z be independent fair bits and set

$$S_0 = X, \qquad S_1 = X, \qquad S_2 = Z, \qquad O = *.$$

Then

$$\mathrm{Ob}_{\mathrm{bridge}}(O) = H(S_1 \mid O, S_0, S_2) = 0,$$

while

$$\mathrm{Ob}_{\mathrm{trace}}(O) = H(S_1 \mid O) = \log_b 2,$$

and

$$\mathrm{Ob}_{\mathrm{complete}}(O) = H(S_0, S_1, S_2 \mid O) = H(X, Z) = 2 \log_b 2.$$

Therefore

$$0 = \mathrm{Ob}_{\mathrm{bridge}}(O) < \mathrm{Ob}_{\mathrm{trace}}(O) < \mathrm{Ob}_{\mathrm{complete}}(O).$$

$\square$

Remark 4.27 (The hierarchy is not a ladder of equivalent tasks)*.* The bridge/trace/complete hierarchy is an inequality hierarchy, not an equivalence hierarchy. A zero bridge obstruction says only that the internal trace is determined after the compressed observation and the endpoints are known. It does not say that the internal trace is determined from the compressed observation alone. Likewise, a zero hidden-trace obstruction says only that the internal trace is recovered; it does not say that endpoints or all complete-trace coordinates are recovered. The three obstructions therefore certify three different reconstruction tasks.

4.9 Marginal static data do not determine trace obstruction

The distinction between static and trace obstruction is not merely terminological. Static one-time posterior data do not determine joint trace posterior data.

Proposition 4.28 (Same static scalar data, different trace obstruction)**.** *There exist two coordinatewise compressed trace systems with the same one-time scalar obstructions*

$$H(S_t \mid \bar{S}_t)$$

for every t, but with different complete-trace obstructions

$$H(S_{0:T} \mid \bar{S}_{0:T}).$$

Proof. It is enough to take $T = 1$. Let both systems have trivial compressed variables $\bar{S}_0 = \bar{S}_1 = *$, and let both S_0 and S_1 be fair bits marginally. Then in both systems

$$H(S_0 \mid \bar{S}_0) = \log_b 2, \qquad H(S_1 \mid \bar{S}_1) = \log_b 2.$$

In the first system, set $S_1 = S_0$. Then

$$H(S_{0:1} \mid \bar{S}_{0:1}) = H(S_0, S_1) = \log_b 2.$$

In the second system, let S_0 and S_1 be independent fair bits. Then

$$H(S_{0:1} \mid \bar{S}_{0:1}) = H(S_0, S_1) = 2 \log_b 2.$$

Thus the one-time scalar obstruction data are identical, but the joint trace obstruction differs. □

Remark 4.29 (Why CEOT IV uses profiles)*.* Proposition 4.28 is a trace-level version of the profile-first discipline. One-time scalar shadows do not determine joint reconstruction profiles. Consequently CEOT IV treats the posterior profile

$$\mathsf{Post}_{S_{0:T}|O}$$

and its hidden-trace analogue as primary, and scalar entropy obstructions as computable shadows with useful zero criteria.

4.10 What this section does not prove

The reconstruction criteria above do not decide whether a compressed algorithm is good for a particular task. They decide whether specific trace variables are reconstructible from the compressed observation. A positive hidden-trace obstruction

$$H(S_{\text{int}} \mid O) > 0$$

certifies hidden internal trace loss, not target failure. A zero complete-trace obstruction is sufficient for every deterministic target of the trace, but it is usually stronger than needed.

The next two sections therefore turn from trace reconstruction to task-relative obstruction. The next section treats deterministic target recovery. The following section treats stochastic target sufficiency.

4.11 Section summary

This section proved the support-relative zero criteria for complete-trace, hidden-trace, bridge, static, and memory reconstruction. It established the hierarchy

$$\mathrm{Ob}_{\mathrm{bridge}}(O) \le \mathrm{Ob}_{\mathrm{trace}}(O) \le \mathrm{Ob}_{\mathrm{complete}}(O),$$

proved that local coordinatewise zero obstruction implies complete-trace zero obstruction in the coordinatewise model, and showed that static one-time scalar data do not determine joint trace obstruction. These are reconstruction results only. The paper now turns to the target-relative layer, where the decisive question is not whether the hidden trace is recoverable, but whether the compressed observation preserves the information needed for the task.

5 Deterministic Target Recovery

5.1 From reconstructing traces to recovering targets

The previous section studied reconstruction of trace variables. It asked whether the complete trace

$$S_{0:T},$$

the hidden internal trace

$$S_{\mathrm{int}},$$

or a bridge-conditioned internal trace can be recovered from a compressed observation O. These are trace reconstruction questions. They are deliberately stronger than many algorithmic tasks require.

In an algorithm, the object of interest is often not the entire internal history. It may be a final output, an optimal action, a selected path, a cost, a certificate, or a finite value label. Such a quantity is represented in this section by a deterministic target

$$Y = f_+(S_{0:T}).$$

The deterministic target may depend on the entire trace, but it may also ignore most of the trace. The central question is therefore not

$$\text{Can one reconstruct } S_{0:T} \text{ from } O?$$

but rather

$$\text{Can one reconstruct } Y \text{ from } O?$$

This distinction is the first task-relative layer of CEOT IV.

The main lesson is simple but structurally important:

$$\boxed{\text{complete trace recovery implies deterministic target recovery, but not conversely.}}$$

Thus positive hidden trace obstruction is not by itself an algorithmic failure certificate. It is only a certificate that some hidden trace information is lost. To determine whether the compression fails a deterministic task, one must measure the target-specific obstruction

$$H(Y \mid O),$$

not only a trace obstruction such as $H(S_{\mathrm{int}} \mid O)$ or $H(S_{0:T} \mid O)$.

5.2 Deterministic target obstruction

Definition 5.1 (Support-relative deterministic target)**.** Let $\mathfrak{A}$ be a finite deterministic trace-compression presentation with full trace $S_{0:T}$, positive trace support

$$\mathcal{T}_+ = \operatorname{supp}(S_{0:T}),$$

and compressed observation $O = c(S_{0:T})$. A target Y is called deterministic relative to the presentation if there exists a support-relative map

$$f_+ : \mathcal{T}_+ \to \mathcal{Y}$$

such that

$$Y = f_+(S_{0:T}) \quad \text{a.s.}$$

Equivalently, the support-restricted target kernel L_+ is Dirac on positive trace support:

$$L_+(y \mid \tau) = \mathbf{1}_{\{y=f_+(\tau)\}}$$

for all $\tau \in \mathcal{T}_+$. Any values of an ambient extension

$$f : \mathcal{S}_{0:T} \to \mathcal{Y}$$

outside $\mathcal{T}_+$ are obstruction-invisible.

Convention 5.2 (Notation for deterministic targets)**.** Throughout CEOT IV, the symbol

$$f_+ : \mathcal{T}_+ \to \mathcal{Y}_+$$

is reserved for a support-relative deterministic target of the full trace:

$$Y = f_+(S_{0:T}) \quad \text{a.s.}$$

Ambient extensions

$$f : \mathcal{S}_{0:T} \to \mathcal{Y}$$

may exist, but their values outside $\mathcal{T}_+ = \operatorname{supp}(S_{0:T})$ are obstruction-invisible and are not part of the support-relative target data.

Auxiliary deterministic maps that are not the full-trace target map are denoted by g, h, r, φ, ψ, or by decorated support-relative symbols such as φ_+ and ψ_+. In particular, for a lifted trace decomposition

$$S_{0:T} = (V_{0:T}, M_{0:T}),$$

a deterministic lifted target should be written as

$$Y = \varphi_+(V_{0:T}, M_{0:T})$$

on positive lifted-trace support, equivalently as

$$Y = f_+(S_{0:T})$$

with

$$f_+(\tau) = \varphi_+(\rho_V(\tau), \rho_M(\tau)).$$

If the target depends only on the visible trace, we write

$$Y = \psi_+(V_{0:T}),$$

or equivalently $f_+(\tau) = \psi_+(\rho_V(\tau))$ on positive support.

Lemma 5.3 (Deterministic target as a Dirac target kernel)**.** *A target Y is support-relative deterministic if and only if its support-restricted target kernel L_+ is Dirac on positive trace support. Equivalently, there exists a unique map*

$$f_+ : \mathcal{T}_+ \to \mathcal{Y}_+$$

up to target-support relabeling such that

$$L_+(y \mid \tau) = \mathbf{1}_{\{y=f_+(\tau)\}}$$

for all $\tau \in \mathcal{T}_+$ and all $y \in \mathcal{Y}_+$.

Proof. If $Y = f_+(S_{0:T})$ almost surely on positive trace support, then conditional on $S_{0:T} = \tau$ the target law is the point mass at $f_+(\tau)$, so the displayed Dirac formula holds. Conversely, if L_+ is Dirac for every positive-support trace τ, define $f_+(\tau)$ to be the unique support target where the point mass is located. Then $Y = f_+(S_{0:T})$ almost surely. Uniqueness is only support-relative: values outside $\mathcal{T}_+$ and target labels outside $\mathcal{Y}_+$ are not part of the obstruction data. □

Definition 5.4 (Deterministic target obstruction)**.** For a support-relative deterministic target $Y = f_+(S_{0:T})$, define the deterministic target obstruction of the compressed observation O by

$$\mathrm{Ob}^{Y}_{\mathrm{det}}(O) := H(Y \mid O).$$

The corresponding posterior profile is

$$\mathsf{Post}_{Y|O} = \left(P_O, \{P_{Y|O=o}\}_{o\in\mathrm{supp}(O)}\right).$$

The scalar obstruction $\mathrm{Ob}^{Y}_{\mathrm{det}}(O)$ is the entropy shadow of the target posterior profile. Its vanishing means that the target is determined by the compressed observation on every positive-probability observation fiber. Its positivity means that at least one positive observation fiber contains two target values with positive conditional probability.

5.3 Fiber criterion for deterministic recovery

The deterministic target obstruction can be read directly from the fibers of the observation map.

Theorem 5.5 (Target zero criterion)**.** *Let $Y = f_+(S_{0:T})$ be a support-relative deterministic target and let $O = c(S_{0:T})$ be a finite deterministic compressed observation. Then the following are equivalent:*

(i) $H(Y \mid O) = 0$.

(ii) *There exists a support-relative decoder*

$$d : \mathrm{supp}(O) \to \mathrm{supp}(Y)$$

such that

$$d(O) = Y \quad a.s.$$

(iii) *For every $o \in \mathrm{supp}(O)$, the target $f_+(s_{0:T})$ is constant over the support fiber*

$$\{s_{0:T} \in \mathrm{supp}(S_{0:T}) : c(s_{0:T}) = o\}.$$

Proof. The equivalence between (i) and (ii) is the finite support-relative zero entropy criterion applied to the pair (Y, O). Indeed, $H(Y \mid O) = 0$ if and only if for every $o \in \operatorname{supp}(O)$ the conditional law $P_{Y|O=o}$ is a point mass. Defining $d(o)$ to be that unique point gives $d(O) = Y$ almost surely. Conversely, if such a decoder exists, then $P_{Y|O=o}$ is a point mass for every positive observation value, hence $H(Y \mid O) = 0$.

For support-relative deterministic $Y = f_+(S_{0:T})$, the conditional law $P_{Y|O=o}$ is the pushforward of $P_{S_{0:T}|O=o}$ under f_+. It is a point mass exactly when f_+ is constant on the positive support fiber over o. This gives the equivalence with (iii). □

Remark 5.6 (Support-relative nature)*.* The fiber condition is only imposed on positive-probability fibers. No value of the target outside $\operatorname{supp}(S_{0:T})$ affects the obstruction. This matches the support-relative convention used throughout CEOT I–III and in the trace reconstruction criteria of Section 4.

5.4 Trace recovery implies target recovery

Complete trace recovery is enough to recover every deterministic target of the trace.

Theorem 5.7 (Complete trace recovery implies deterministic target recovery)**.** *If $Y = f_+(S_{0:T})$ is a support-relative deterministic target, then*

$$H(Y \mid O) \leq H(S_{0:T} \mid O).$$

In particular,

$$\mathrm{Ob}_{\mathrm{complete}}(O) = 0 \quad \Longrightarrow \quad \mathrm{Ob}^{Y}_{\mathrm{det}}(O) = 0.$$

Proof. Since $Y = f_+(S_{0:T})$, conditioning on O gives the entropy monotonicity inequality for a function of a random variable:

$$H(Y \mid O) \leq H(S_{0:T} \mid O).$$

Equivalently, if $H(S_{0:T} \mid O) = 0$, then by Theorem 4.12 there exists R with $S_{0:T} = R(O)$ a.s. Hence

$$Y = f_+(S_{0:T}) = f_+(R(O)) \quad \text{a.s.},$$

so Y is a function of O and $H(Y \mid O) = 0$. □

Theorem 5.7 is intentionally one-sided. It says that complete trace recovery is sufficient for deterministic target recovery. It does not say that complete trace recovery is necessary.

Corollary 5.8 (Universal deterministic target preservation)**.** *The condition*

$$H(S_{0:T} \mid O) = 0$$

is equivalent to deterministic recovery of every finite support-relative target of the form $Y = f_+(S_{0:T})$.

Proof. If $S_{0:T}$ is recoverable from O, then every deterministic target of $S_{0:T}$ is recoverable by composition. Conversely, take the deterministic target to be $Y = S_{0:T}$. □

Thus complete trace recovery is a universal deterministic-target criterion. CEOT IV is interested in the weaker, purpose-relative question for a specified target Y.

5.5 Hidden-trace recovery may be sufficient

Some deterministic targets need only the internal hidden trace together with the compressed observation. In that case hidden-trace recovery is already sufficient.

Proposition 5.9 (Hidden-trace recovery implies target recovery under trace dependence)**.** *Assume that*

$$Y = g(S_{\mathrm{int}}, O) \quad \text{a.s.}$$

for some finite map g. If

$$H(S_{\mathrm{int}} \mid O) = 0,$$

then

$$H(Y \mid O) = 0.$$

Proof. By the hidden-trace zero criterion, there exists a support-relative decoder R_{trace} such that

$$S_{\mathrm{int}} = R_{\mathrm{trace}}(O) \quad \text{a.s.}$$

Therefore

$$Y = g(S_{\mathrm{int}}, O) = g(R_{\mathrm{trace}}(O), O) \quad \text{a.s.}$$

Hence Y is a function of O, and so $H(Y \mid O) = 0$. □

This proposition explains why CEOT IV keeps both complete-trace and hidden-trace obstructions. If a target depends on endpoints, hidden internal states, and compressed observable data in different ways, then the correct sufficient reconstruction condition depends on the target's support-relative dependence structure.

Corollary 5.10 (Observed targets have zero obstruction)**.** *If $Y = h(O)$ a.s. for some finite map h, then*

$$\mathrm{Ob}^{Y}_{\mathrm{det}}(O) = 0.$$

Proof. Take $d = h$ in Theorem 5.5. □

5.6 Non-converse: target recovery without trace recovery

The converse of Theorem 5.7 fails. A compressed observation may preserve the deterministic target while losing hidden trace information.

Example 5.11 (Irrelevant hidden randomness)**.** Let $T = 2$ and let endpoints be trivial:

$$S_0 = *, \qquad S_2 = *.$$

Let

$$Y \sim \mathrm{Bernoulli}(1/2), \qquad Z \sim \mathrm{Bernoulli}(1/2), \qquad Z \perp Y,$$

and set

$$S_1 = (Y, Z), \qquad O = Y.$$

Take the deterministic target to be Y. Then

$$H(Y \mid O) = 0,$$

but

$$H(S_1 \mid O) = H(Y, Z \mid Y) = H(Z \mid Y) = \log_b 2.$$

Thus the target is perfectly recovered from the compressed observation even though the hidden internal trace is not reconstructed.

Proposition 5.12 (Deterministic target recovery does not imply trace recovery). *There exist finite trace-compression presentations with*

$$\mathrm{Ob}^{Y}_{\mathrm{det}}(O) = 0$$

but

$$\mathrm{Ob}_{\mathrm{trace}}(O) > 0 \quad \textit{and} \quad \mathrm{Ob}_{\mathrm{complete}}(O) > 0.$$

Proof. Example 5.11 gives $H(Y \mid O) = 0$ and $H(S_1 \mid O) = \log_b 2 > 0$. Since $S_{0:T}$ contains S_1, also $H(S_{0:T} \mid O) > 0$. □

This is the basic deterministic separation principle:

$$\boxed{\text{deterministic target recovery} \not\Longrightarrow \text{hidden trace recovery.}}$$

It is one of the main reasons CEOT IV treats task-relative obstruction separately from trace reconstruction obstruction.

5.7 Changing the target changes the obstruction

A compressed observation may fail for one deterministic target while succeeding for a coarser one. This is not a defect of the theory; it is the purpose-relative character of the theory.

Proposition 5.13 (Target coarsening monotonicity). *Let $Y = f_+(S_{0:T})$ be a support-relative deterministic target and let*

$$Y' = r(Y)$$

be a coarsening of the target. Then

$$H(Y' \mid O) \leq H(Y \mid O).$$

In particular,

$$H(Y \mid O) = 0 \quad \Longrightarrow \quad H(Y' \mid O) = 0.$$

Proof. Since $Y' = r(Y)$, Y' is a function of Y. Conditional entropy cannot increase under applying a deterministic function to the variable being reconstructed, hence

$$H(Y' \mid O) \leq H(Y \mid O).$$

□

Example 5.14 (Correct decision without full value recovery). Suppose a trace determines a finite value label V_{val} and a decision

$$Y' = r(V_{\mathrm{val}})$$

obtained by thresholding or choosing an argmin class. A compressed observation may determine the decision Y' while failing to determine the full value label V_{val}. In CEOT terms,

$$H(Y' \mid O) = 0$$

may hold while

$$H(V_{\mathrm{val}} \mid O) > 0.$$

Thus the obstruction is attached to the declared target, not to an implicit stronger quantity.

5.8 Target fibers and quotient interpretation

For deterministic targets, the support-relative map $f_+ : \mathcal{T}_+ \to \mathcal{Y}$ induces a quotient of the support of the full trace into target fibers. The observation $O = c(S_{0:T})$ induces another quotient into observation fibers. The zero criterion says that observation fibers refine target fibers on support.

Proposition 5.15 (Observation fibers refine target fibers)**.** *Let $Y = f_+(S_{0:T})$ and $O = c(S_{0:T})$. Then*

$$H(Y \mid O) = 0$$

if and only if for all $s, s' \in \operatorname{supp}(S_{0:T})$,

$$c(s) = c(s') \quad \Longrightarrow \quad f(s) = f(s').$$

Equivalently, there exists a support-relative map d such that

$$f = d \circ c$$

on $\operatorname{supp}(S_{0:T})$.

Proof. This is a restatement of the fiber criterion in Theorem 5.5. If $H(Y \mid O) = 0$, define $d(o)$ to be the unique target value attained on the positive support fiber $c^{-1}(o) \cap \operatorname{supp}(S_{0:T})$. Then $f = d \circ c$ on support. Conversely, if such a d exists, then $Y = d(O)$ a.s., so $H(Y \mid O) = 0$. □

This quotient formulation is useful when the compression is deterministic. It says that deterministic target recovery is exactly compatibility of two finite partitions on the actually realized trace support: the observation partition must be at least as fine as the target partition.

5.9 Relation to posterior profiles

Although $H(Y \mid O)$ is the scalar obstruction used for deterministic target recovery, the profile-level object remains

$$\mathsf{Post}_{Y|O} = \left(P_O, \{P_{Y|O=o}\}_{o \in \operatorname{supp}(O)}\right).$$

The scalar vanishes exactly when each conditional law in this profile is a point mass. However, scalar equality alone does not determine the profile. Two targets may have the same value of $H(Y \mid O)$ but different conditional distributions over observation fibers. Consequently, as in CEOT III and in Section 3, CEOT IV treats the posterior profile as the invariant and the scalar entropy as a computable shadow.

For deterministic targets, the profile also records where the compression fails. If $H(Y \mid O) > 0$, then some positive observation value o has a non-degenerate conditional target law $P_{Y|O=o}$. Such an o is a concrete obstruction fiber: within that compressed observation class, the target has not been determined.

5.10 What deterministic target recovery does not cover

The theory in this section applies only to support-relative deterministic targets $Y = f_+(S_{0:T})$. It is not the correct criterion for targets that remain random even after the full trace is known. If the target is generated by a non-degenerate kernel

$$L : \mathcal{S}_{0:T} \rightsquigarrow \mathcal{Y},$$

then the condition $H(Y \mid O) = 0$ is usually too strong and often conceptually wrong. In that case the right question is not whether the random target value is determined by O, but whether O preserves all the trace information relevant to the target law. The next section treats that question using conditional mutual information:

$$I(Y; S_{0:T} \mid O).$$

5.11 Section summary

This section defined the deterministic target obstruction

$$\mathrm{Ob}^Y_{\mathrm{det}}(O) = H(Y \mid O),$$

proved the support-relative target zero criterion, and showed that complete trace recovery implies deterministic target recovery. It also showed that the converse fails: a compressed observation can recover a target while losing hidden trace information. Finally, it made explicit the purpose-relative nature of CEOT IV: changing the target changes the obstruction. The next section moves from deterministic recovery to stochastic target sufficiency.

6 Stochastic Law-Sufficiency and Loss-Relative Sufficiency

6.1 Why deterministic recovery is not the right criterion

Section 5 treated the case in which the target is a support-relative function of the full trace,

$$Y = f_+(S_{0:T}).$$

In that setting the correct target obstruction is the conditional entropy

$$\mathrm{Ob}^Y_{\mathrm{det}}(O) = H(Y \mid O),$$

and vanishing means that the compressed observation determines the target value.

Many algorithmic targets are not deterministic in this sense. A randomized algorithm may output a randomized decision after a trace has been generated. A stochastic dynamic program may assign to a trace a probability law for a future state rather than a single future state. A simulation or sampling procedure may retain a trace-dependent distribution from which a later quantity is drawn. In such cases the target is generated by a kernel

$$L : \mathcal{S}_{0:T} \rightsquigarrow \mathcal{Y},$$

and even knowing the full trace may not determine the realized target value. Thus the condition

$$H(Y \mid O) = 0$$

is usually too strong. It asks the compressed observation to determine the random draw itself, not merely to preserve the trace-dependent information governing the target law.

The correct question is therefore not:

$$\text{is } Y \text{ determined by } O?$$

It is:

$$\boxed{\text{does } O \text{ contain all information in } S_{0:T} \text{ relevant to the law of } Y?}$$

This is a sufficiency question. The compressed observation is sufficient for the stochastic target precisely when, after conditioning on O, the full trace carries no additional information about Y.

6.2 Target-kernel compatibility

Throughout this paper the compressed observation is deterministic:

$$O = c(S_{0:T}).$$

The target kernel is a finite stochastic map

$$L : \mathcal{S}_{0:T} \rightsquigarrow \mathcal{Y}.$$

Equivalently, for each positive trace value $\tau \in \mathcal{T}_+ = \mathrm{supp}(S_{0:T})$ one has a probability distribution

$$L_+(\cdot \mid \tau) \in \mathrm{Prob}(\mathcal{Y}_+).$$

The ambient values of L outside $\mathcal{T}_+$ are irrelevant to CEOT IV sufficiency obstructions. The joint law is

$$P(S_{0:T} = s, O = o, Y = y) = P(S_{0:T} = s)\mathbf{1}_{\{o=c(s)\}}L_+(y \mid s).$$

Lemma 6.1 (Target-kernel compatibility)**.** *On the support of the induced law,*

$$P(Y \mid S_{0:T}, O) = P(Y \mid S_{0:T}).$$

Equivalently,

$$P(Y = y \mid S_{0:T} = s, O = c(s)) = L_+(y \mid s)$$

for every $s \in \mathrm{supp}(S_{0:T})$ *and* $y \in \mathcal{Y}_+$*.*

Proof. The observation O is a deterministic function of $S_{0:T}$. Hence conditioning on both $(S_{0:T}, O)$ adds no information beyond conditioning on $S_{0:T}$ alone. Substituting the joint law above gives the displayed identity on every positive-support event. □

This lemma is small but important. It ensures that the conditional-independence criterion below can be stated either as

$$P(Y \mid S_{0:T}, O) = P(Y \mid O)$$

or, equivalently, as

$$P(Y \mid S_{0:T}) = P(Y \mid O)$$

on support.

6.3 The coarsest law-sufficient quotient

The stochastic target obstruction used below is a full-law sufficiency obstruction. Before reducing it to a scalar conditional mutual information, there is a canonical quotient that captures exactly the trace information relevant to the declared target law. The word “minimal” is used only relative to the fixed positive trace support, the fixed support-restricted target kernel, and the support-relative quotient-refinement order defined next. It is not a target-independent quotient of the algorithmic presentation.

Let

$$\mathcal{T}_+ := \mathrm{supp}(S_{0:T}).$$

Definition 6.2 (Support-relative quotient-refinement order)**.** A support-relative deterministic quotient of the trace is a surjective finite map

$$q : \mathcal{T}_+ \to \mathcal{Q}.$$

For two such quotients

$$q : \mathcal{T}_+ \to \mathcal{Q}, \qquad q' : \mathcal{T}_+ \to \mathcal{Q}',$$

write

$$q \preceq q'$$

if q' refines q, equivalently if there exists a finite map

$$r : q'(\mathcal{T}_+) \to q(\mathcal{T}_+)$$

such that

$$q = r \circ q'.$$

Thus $q \preceq q'$ means that q is no more informative than q'.

Definition 6.3 (Law-sufficient quotient)**.** A support-relative quotient

$$q : \mathcal{T}_+ \to \mathcal{Q}$$

is law-sufficient for the support-restricted target kernel L_+ if the induced observation

$$Q = q(S_{0:T})$$

satisfies

$$I(Y; S_{0:T} \mid Q) = 0.$$

Equivalently, $L_+(\cdot \mid \tau)$ is constant on every positive fiber of q.

For the support-restricted target kernel L_+, define an equivalence relation on the positive-support trace set by

$$\tau \sim_{L_+} \tau' \quad \Longleftrightarrow \quad L_+(\cdot \mid \tau) = L_+(\cdot \mid \tau').$$

For brevity, write q_L for q_{L_+}, since the ambient kernel only matters through its support restriction. Let

$$q_L : \mathcal{T}_+ \to \mathcal{T}_+/{\sim_{L_+}}$$

be the quotient map and define the finite law-quotient variable

$$Q_L := q_L(S_{0:T}).$$

Equivalently, Q_L records the trace-conditioned target law

$$\mathcal{L}(Y \mid S_{0:T}).$$

Definition 6.4 (Relative canonical law quotient)**.** Fix the positive trace support $\mathcal{T}_+ = \operatorname{supp}(S_{0:T})$ and the support-restricted target kernel

$$L_+ : \mathcal{T}_+ \rightsquigarrow \mathcal{Y}_+.$$

The relative canonical law quotient is the quotient

$$q_L : \mathcal{T}_+ \to \mathcal{T}_+/{\sim_{L_+}}$$

where

$$\tau \sim_{L_+} \tau' \quad \Longleftrightarrow \quad L_+(\cdot \mid \tau) = L_+(\cdot \mid \tau').$$

The word canonical always means canonical relative to the pair $(\mathcal{T}_+, L_+)$.

Theorem 6.5 (Support- and target-relative coarsest law-sufficient quotient)**.** *The following statement is relative to the fixed positive trace support $\mathcal{T}_+$ and the fixed support-restricted target kernel L_+. It does not define a target-independent quotient of the algorithmic presentation.*

The quotient observation

$$Q_L = q_L(S_{0:T})$$

is law-sufficient for Y:

$$I(Y; S_{0:T} \mid Q_L) = 0.$$

Moreover, q_L is the coarsest law-sufficient quotient in the support-relative quotient-refinement order. Namely, if

$$q : \mathcal{T}_+ \to \mathcal{Q}$$

is any law-sufficient quotient, then

$$q_L \preceq q,$$

equivalently there exists a finite map

$$r : q(\mathcal{T}_+) \to q_L(\mathcal{T}_+)$$

such that

$$q_L = r \circ q.$$

For any deterministic observation

$$O = c(S_{0:T}),$$

the following are equivalent:

(i)

$$I(Y; S_{0:T} \mid O) = 0.$$

(ii) The target kernel is constant on every positive observation fiber:

$$c(\tau) = c(\tau') \quad \Longrightarrow \quad L_+(\cdot \mid \tau) = L_+(\cdot \mid \tau')$$

for all $\tau, \tau' \in \mathcal{T}_+$.

(iii) The law quotient factors through O: there exists a map

$$r : \mathrm{supp}(O) \to \mathcal{T}_+/{\sim_{L_+}}$$

such that

$$q_L(\tau) = r(c(\tau))$$

for every $\tau \in \mathcal{T}_+$.

Proof. By construction, Q_L records exactly the value of the conditional law $L(\cdot \mid S_{0:T})$. Hence

$$P(Y \mid S_{0:T}) = P(Y \mid Q_L)$$

on support. Lemma 6.1 then gives

$$I(Y; S_{0:T} \mid Q_L) = 0.$$

Now let $q : \mathcal{T}_+ \to \mathcal{Q}$ be any law-sufficient quotient. By Definition 6.3, or equivalently by the stochastic sufficiency criterion below, $L_+(\cdot \mid \tau)$ is constant on every positive fiber of q. Therefore

$$q(\tau) = q(\tau') \quad \Longrightarrow \quad L_+(\cdot \mid \tau) = L_+(\cdot \mid \tau'),$$

and hence

$$q_L(\tau) = q_L(\tau').$$

Thus q_L is determined by q, so there exists a unique finite map

$$r : q(\mathcal{T}_+) \to q_L(\mathcal{T}_+)$$

with $q_L = r \circ q$. This is exactly $q_L \preceq q$.

For a deterministic observation $O = c(S_{0:T})$, Theorem 6.10 says that $I(Y; S_{0:T} \mid O) = 0$ is equivalent to

$$L_+(\cdot \mid \tau) = P(Y \mid O = c(\tau))$$

for every positive-support trace τ. This holds exactly when $L_+(\cdot \mid \tau)$ is constant on each positive observation fiber, which is equivalent to saying that q_L factors through O. □

Remark 6.6 (The law quotient is not task-independent). The quotient q_L is not an invariant of the trace law alone. It is an invariant of the declared support-restricted target kernel on the positive trace support. Two different targets L_1, L_2 on the same trace support may induce different coarsest law-sufficient quotients,

$$q_{L_1} \neq q_{L_2}.$$

Likewise, changing the positive trace support can change the quotient, even if an ambient target kernel is left unchanged outside support. Thus the coarsest law-sufficient quotient is target-relative and support-relative, not an absolute compression of the algorithm.

Example 6.7 (Different targets give different coarsest quotients)**.** Let $\mathcal{T}_+ = \{a, b, c, d\}$. Let Y_1 be deterministic with

$$f_1(a) = f_1(b) = 0, \qquad f_1(c) = f_1(d) = 1.$$

Then the relative canonical law quotient for Y_1 has fibers $\{a, b\}$ and $\{c, d\}$. Let Y_2 be deterministic with

$$f_2(a) = f_2(c) = 0, \qquad f_2(b) = f_2(d) = 1.$$

Then the relative canonical law quotient for Y_2 has fibers $\{a, c\}$ and $\{b, d\}$. The positive trace support is the same, but the coarsest law-sufficient quotient changes with the target. Therefore the quotient is not an intrinsic algorithmic quotient; it is a target-relative quotient.

Corollary 6.8 (Coarsest deterministic-target quotient)**.** *If*

$$Y = f_+(S_{0:T})$$

is a support-relative deterministic target, then the coarsest law-sufficient quotient is the deterministic target quotient

$$\tau \sim_{f_+} \tau' \iff f_+(\tau) = f_+(\tau').$$

Consequently, a deterministic observation $O = c(S_{0:T})$ recovers the deterministic target if and only if this target quotient factors through O, equivalently

$$q_f \preceq c|_{\mathcal{T}_+}.$$

Proof. For a support-relative deterministic target, the kernel $L_+(\cdot \mid \tau)$ is a Dirac mass at $f_+(\tau)$. Two such Dirac masses are equal exactly when the corresponding target values are equal. The factorization statement is Theorem 6.5 specialized to deterministic kernels, together with Theorem 5.5. □

6.4 The stochastic law-sufficiency obstruction

Definition 6.9 (Stochastic law-sufficiency obstruction)**.** Let $\mathfrak{A}$ be a finite deterministic trace-compression presentation with support-restricted target kernel L_+. The *stochastic law-sufficiency obstruction* of O for Y is

$$\mathrm{Ob}^{Y}_{\mathrm{law}}(O) := I(Y; S_{0:T} \mid O).$$

Expanded in entropic form,

$$I(Y; S_{0:T} \mid O) = H(Y \mid O) - H(Y \mid S_{0:T}, O).$$

By Lemma 6.1, this becomes

$$I(Y; S_{0:T} \mid O) = H(Y \mid O) - H(Y \mid S_{0:T}).$$

Thus the obstruction measures exactly how much additional uncertainty about Y is removed by the full trace after the compressed observation is already known.

The quantity is always nonnegative:

$$\mathrm{Ob}^{Y}_{\mathrm{law}}(O) \geq 0.$$

It vanishes precisely when the full trace gives no additional information about the target beyond the compressed observation.

6.5 Sufficiency zero criterion

Theorem 6.10 (Stochastic law-sufficiency criterion)**.** *For a finite deterministic trace-compression presentation,*

$$I(Y; S_{0:T} \mid O) = 0$$

if and only if

$$Y \perp S_{0:T} \mid O.$$

Equivalently, for every $o \in \mathrm{supp}(O)$ and every $s \in \mathrm{supp}(S_{0:T})$ with $c(s) = o$ and $P(S_{0:T} = s) > 0$,

$$P(Y \mid S_{0:T} = s) = P(Y \mid O = o).$$

Proof. For finite random variables, conditional mutual information vanishes if and only if the corresponding conditional independence relation holds. Hence

$$I(Y; S_{0:T} \mid O) = 0 \quad \iff \quad Y \perp S_{0:T} \mid O.$$

Conditional independence is equivalent, on positive support, to

$$P(Y \mid S_{0:T}, O) = P(Y \mid O).$$

Using Lemma 6.1, this is equivalent to

$$P(Y \mid S_{0:T}) = P(Y \mid O)$$

on every positive-support fiber of O. □

The theorem says that O is law-sufficient for the stochastic target exactly when every trace inside the same observation fiber induces the same target law after averaging over that fiber. More explicitly, if $o \in \mathrm{supp}(O)$, then

$$P(Y \mid O = o) = \sum_{s \in c^{-1}(o)} P(Y \mid S_{0:T} = s) P(S_{0:T} = s \mid O = o).$$

The obstruction vanishes exactly when every positive-support trace s in the fiber $c^{-1}(o)$ already has this same law:

$$P(Y \mid S_{0:T} = s) = P(Y \mid O = o).$$

Thus stochastic law-sufficiency is a fiberwise equality of target laws, not a pointwise recovery of target samples.

Corollary 6.11 (Complete trace recovery implies stochastic law-sufficiency). *If*

$$H(S_{0:T} \mid O) = 0,$$

then

$$I(Y; S_{0:T} \mid O) = 0$$

for every target kernel $L : \mathcal{S}_{0:T} \rightsquigarrow \mathcal{Y}$.

Proof. If $H(S_{0:T} \mid O) = 0$, then $S_{0:T} = R(O)$ a.s. for some support-relative decoder R by Theorem 4.12. Therefore conditioning on O already determines $S_{0:T}$, so Y and $S_{0:T}$ are conditionally independent given O. □

The converse fails. A compressed observation can be sufficient for a particular stochastic target while losing large amounts of trace information irrelevant to that target. This is the stochastic analogue of the deterministic non-converse in Section 5.

6.6 Law-sufficiency versus loss-relative sufficiency

The obstruction

$$I(Y; S_{0:T} \mid O)$$

is a full-law sufficiency obstruction. It asks whether the compressed observation preserves the entire conditional law of Y carried by the full trace. Many decision tasks require less. They may only require the same optimal action, the same expected cost, the same feasibility decision, or the same certificate validity.

Let $\mathcal{D}$ be a finite decision set and let

$$\ell : \mathcal{Y} \times \mathcal{D} \to \mathbb{R}$$

be a finite loss function. Define the full-trace Bayes risk

$$R_\ell(S_{0:T}) := \inf_{\delta : \mathcal{S}_{0:T} \to \mathcal{D}} \mathbb{E}\big[\ell(Y, \delta(S_{0:T}))\big]$$

and the observation-level Bayes risk

$$R_\ell(O) := \inf_{\delta : \mathcal{O} \to \mathcal{D}} \mathbb{E}\big[\ell(Y, \delta(O))\big].$$

The loss-relative CEOT IV obstruction is

$$\mathrm{Ob}_\ell^Y(O) := R_\ell(O) - R_\ell(S_{0:T}).$$

Proposition 6.12 (Law-sufficiency implies loss-relative sufficiency)**.** *For every finite loss function ℓ,*

$$\mathrm{Ob}_\ell^Y(O) \geq 0.$$

Moreover, if

$$I(Y; S_{0:T} \mid O) = 0,$$

then

$$\mathrm{Ob}_\ell^Y(O) = 0$$

for every finite loss function ℓ.

Proof. Since every decision rule based on O is also a decision rule based on $S_{0:T}$, because $O = c(S_{0:T})$, the full trace cannot have larger Bayes risk:

$$R_\ell(S_{0:T}) \leq R_\ell(O).$$

This proves nonnegativity.

If $I(Y; S_{0:T} \mid O) = 0$, then

$$P(Y \mid S_{0:T}, O) = P(Y \mid O)$$

on support. Therefore the posterior risk of any action after observing the full trace depends only on O. A Bayes rule based on $S_{0:T}$ can be chosen to be a function of O. Hence the two Bayes risks coincide. □

Remark 6.13 (The converse depends on the loss)*.* For a fixed loss function, $\mathrm{Ob}_\ell^Y(O) = 0$ need not imply

$$I(Y; S_{0:T} \mid O) = 0.$$

A compression may preserve the optimal decision for a particular loss while failing to preserve the full conditional law of Y. CEOT IV therefore separates full-law sufficiency from task-specific decision sufficiency.

Theorem 6.14 (Hierarchy of law, universal-loss, and fixed-loss sufficiency)**.** *For a finite deterministic trace-compression presentation, consider the following three statements.*

(i) Full-law sufficiency:

$$I(Y; S_{0:T} \mid O) = 0.$$

(ii) Universal finite-loss sufficiency:

$$\mathrm{Ob}_\ell^Y(O) = 0$$

for every finite decision space and every finite real-valued loss function ℓ.

(iii) Fixed-loss sufficiency:

$$\mathrm{Ob}_{\ell_0}^Y(O) = 0$$

for one declared finite decision space and one declared finite loss function ℓ_0.

Then

$$\textit{full-law sufficiency} \Longrightarrow \textit{universal finite-loss sufficiency} \Longrightarrow \textit{fixed-loss sufficiency.}$$

The final implication cannot be reversed in general. Moreover, universal finite-loss sufficiency is equivalent to full-law sufficiency whenever the admitted finite decision problems contain a separating class of losses for probability laws on $\mathcal{Y}$.

Proof. The first implication is Proposition 6.12. The second implication is immediate by specializing the universal statement to the declared loss ℓ_0. The non-reversal is shown in Example 6.15. For the final claim, if two finite conditional target laws differ, a separating finite loss can be chosen whose Bayes risk distinguishes them. Hence universal zero loss-relative obstruction forces the conditional target law seen from the full trace to agree with the conditional target law seen from O, which is exactly $I(Y; S_{0:T} \mid O) = 0$. □

Example 6.15 (Fixed-loss sufficiency without law-sufficiency)**.** Let $S \in \{s_1, s_2\}$ with both states having positive probability, and let O be constant. Let $Y \in \{0, 1\}$ satisfy

$$P(Y = 1 \mid S = s_1) = 0.6, \qquad P(Y = 1 \mid S = s_2) = 0.9.$$

Then $P(Y \mid S)$ is not constant on the unique observation fiber, so

$$I(Y; S \mid O) > 0.$$

Thus O is not law-sufficient. However, for binary zero-one loss, the Bayes action is 1 after observing either $S = s_1$ or $S = s_2$, and it is also 1 after observing the constant O. Therefore access to S gives no Bayes-risk improvement for this fixed loss:

$$\mathrm{Ob}^{Y}_{\ell_{0\text{-}1}}(O) = 0.$$

This shows that fixed-loss sufficiency is strictly weaker than law-sufficiency.

Remark 6.16 (Bounded-loss assumption in approximate guarantees). The exact implication from law-sufficiency to zero loss-relative obstruction is finite and does not require an additional boundedness hypothesis beyond finite real-valued loss values. Quantitative approximate control is different. The total-variation bound on Bayes-risk degradation requires a uniform bound $0 \le \ell \le L_{\max}$. Without such a bound, small conditional mutual information or small conditional total variation does not by itself give a uniform finite bound on loss degradation. Thus the approximate loss statements in CEOT IV are bounded-loss certificates, not universal operational guarantees for arbitrary unbounded losses.

6.7 Relation to statistical sufficiency and Blackwell comparison

The stochastic target-law criterion

$$I(Y; S_{0:T} \mid O) = 0$$

is a conditional-independence sufficiency statement for the declared target Y. It says that the full trace carries no additional information about Y once the compressed observation O is known. It is not, by itself, a target-independent Blackwell comparison of experiments.

Proposition 6.17 (Target-law sufficiency is target-relative)**.** *Let S be a finite trace variable and let $O = c(S)$ be a finite observation. If*

$$I(Y; S \mid O) = 0,$$

then O is law-sufficient for the declared target Y:

$$P_{Y|S,O} = P_{Y|O} \quad a.s.$$

However, this does not imply that O is sufficient for every other target generated from the same trace. In general, there can exist a target Z such that

$$I(Z; S \mid O) > 0.$$

Proof. The first statement is the finite equivalence between zero conditional mutual information and conditional independence. For the non-universal part, let $S = (A, B)$, where A and B are independent fair bits, and let $O = A$. Let $Y = A$. Then

$$I(Y; S \mid O) = I(A; (A, B) \mid A) = 0.$$

But for $Z = B$,

$$I(Z; S \mid O) = I(B; (A, B) \mid A) = H(B) = \log_b 2 > 0.$$

Thus sufficiency for Y is not target-independent sufficiency. □

Remark 6.18 (Not a universal Blackwell dominance statement)*.* The criterion

$$I(Y; S_{0:T} \mid O) = 0$$

is a target-relative law-sufficiency criterion. It should not be read as a universal Blackwell dominance statement over all possible targets, experiments, and losses. Blackwell comparison becomes relevant only after one fixes the experiment, the decision class, and the family of loss functions under comparison. CEOT IV uses bounded-loss Bayes-risk obstruction as a finite operational shadow, not as a replacement for the full Blackwell order.

6.8 Deterministic recovery as a special case

The stochastic criterion subsumes deterministic target recovery.

Theorem 6.19 (Deterministic target as a special case)**.** *If*

$$Y = f_+(S_{0:T})$$

is a support-relative deterministic target, then

$$I(Y; S_{0:T} \mid O) = H(Y \mid O).$$

Consequently,

$$\mathrm{Ob}^Y_{\mathrm{law}}(O) = \mathrm{Ob}^Y_{\mathrm{det}}(O).$$

Proof. Since $Y = f_+(S_{0:T})$, the target is determined by the full trace on positive trace support. Hence

$$H(Y \mid S_{0:T}, O) = 0.$$

Therefore

$$I(Y; S_{0:T} \mid O) = H(Y \mid O) - H(Y \mid S_{0:T}, O) = H(Y \mid O).$$

□

This theorem explains why Section 5 is not separate from the stochastic theory but rather its deterministic boundary case. In the deterministic case, sufficiency becomes recoverability. In the non-degenerate stochastic case, sufficiency means preservation of the conditional target law.

6.9 Two elementary examples

Example 6.20 (Random target independent of the trace)**.** Let $Y \sim \text{Bernoulli}(1/2)$ be independent of the full trace $S_{0:T}$, and let O be any deterministic observation of the trace. Then

$$H(Y \mid O) = \log_b 2,$$

so deterministic recovery fails unless the base and law are degenerate. However,

$$I(Y; S_{0:T} \mid O) = 0,$$

because Y is independent of the trace even after conditioning on any trace-determined observation. Thus requiring $H(Y \mid O) = 0$ would falsely declare failure, while the stochastic sufficiency criterion correctly declares that no target-relevant trace information has been lost.

Example 6.21 (Trace-dependent noisy target)**.** Let S be a single binary trace variable with

$$P(S = 0) = P(S = 1) = \frac{1}{2},$$

and let the observation be trivial:

$$O = *.$$

Let Y be generated by the kernel

$$P(Y = 1 \mid S = 1) = p, \qquad P(Y = 1 \mid S = 0) = q,$$

with $p \neq q$ and $0 < p, q < 1$. Then Y is not determined by S, so deterministic recovery is not the relevant notion. Nevertheless the trace changes the target law. Since the two conditional target laws differ on the same observation fiber, Theorem 6.10 gives

$$I(Y; S \mid O) > 0.$$

Thus the trivial observation is not sufficient for the stochastic target.

Together these examples isolate the point of CEOT IV. Stochastic target sufficiency is neither the same as target determinism nor the same as trace reconstruction. It asks exactly whether the compressed observation preserves the target law induced by the trace.

6.10 Future outcomes and finite future-law parameters

Section 2 distinguished a future random outcome from the finite parameter representing its trace-dependent law. Suppose a future random variable is generated by

$$Y_{\text{fut}} \sim L_{\text{fut}}(\cdot \mid S_{0:T}).$$

If the task is to preserve prediction of the future outcome itself, the sufficiency condition is

$$I(Y_{\text{fut}}; S_{0:T} \mid O) = 0.$$

This says that, after observing O, the full trace does not further improve the predictive law of the future outcome.

If instead the task is to recover the law assigned to the future outcome, define the finite law parameter

$$\Theta = \mathcal{L}(Y_{\mathrm{fut}} \mid S_{0:T}).$$

Since the trace is finite, Θ takes values in the finite realized set

$$\{\mathcal{L}(Y_{\mathrm{fut}} \mid S_{0:T} = s) : s \in \operatorname{supp}(S_{0:T})\}.$$

Then law recovery is deterministic target recovery:

$$H(\Theta \mid O) = 0.$$

This distinction prevents a common category error. The future sample Y_{fut} may remain random even when the full trace is known, while the future-law parameter Θ may be a deterministic function of the full trace.

6.11 Relation to posterior profiles

For stochastic targets, the posterior target profile

$$\mathsf{Post}_{Y|O} = \Big(P_O, \{P_{Y|O=o}\}_{o \in \operatorname{supp}(O)}\Big)$$

records the target law seen after compression. However, sufficiency also compares this compressed target law with the trace-conditioned kernel

$$P(Y \mid S_{0:T} = s) = L(\cdot \mid s).$$

Thus the stochastic sufficiency profile is not merely $\mathsf{Post}_{Y|O}$; it includes the comparison between target laws inside each observation fiber.

One may regard the relevant finite object as the fiberwise family

$$\Big(P_O, \{P_{S_{0:T}|O=o}\}_o, \{L(\cdot \mid s)\}_{s \in \operatorname{supp}(S_{0:T})}\Big).$$

The scalar shadow of this comparison is

$$I(Y; S_{0:T} \mid O).$$

As before, CEOT IV treats the profile-level data as primary and the scalar obstruction as the computable certificate extracted from it.

6.12 Normalized obstruction indices

Raw CEOT IV obstructions are measured in entropy or mutual-information units. They are the primary quantities for exact zero criteria, positivity criteria, and monotonicity. Normalized indices are useful for comparing different finite systems, but they are not canonical until a denominator convention has been declared.

Definition 6.22 (Normalized obstruction index with declared denominator)**.** Let A be a finite trace-derived or target-derived random variable and let O be a finite observation. A normalized reconstruction obstruction is a pair

$$\widehat{\mathrm{Ob}}_{A,\mathsf{D}}(O) := \frac{H(A \mid O)}{\mathsf{D}(A)},$$

where $\mathsf{D}(A) > 0$ is a declared normalization denominator. Common choices are

$$\mathsf{D}_{\mathrm{ent}}(A) = H(A), \qquad \mathsf{D}_{\mathrm{supp}}(A) = \log_b |\operatorname{supp}(A)|,$$

and, when an ambient finite alphabet $\mathcal{A}$ is part of the typed profile,

$$\mathsf{D}_{\mathrm{amb}}(A) = \log_b |\mathcal{A}|.$$

If $\mathsf{D}(A) = 0$, the normalized index is left undefined unless an explicit zero-denominator convention is declared.

For the main finite comparisons below, CEOT IV uses entropy-normalization. Thus, when no denominator subscript is displayed in this subsection, the convention is

$$\widehat{\mathrm{Ob}}_{A,\mathrm{ent}}(O) := \begin{cases} \dfrac{H(A \mid O)}{H(A)}, & H(A) > 0, \\ 0, & H(A) = 0. \end{cases}$$

The zero-denominator value is a local convention for this subsection, not a canonical normalization rule.

For bridge-relative comparisons, entropy-normalization relative to a baseline conditioning variable is

$$\widehat{\mathrm{Ob}}_{A,\mathrm{ent}}(O; C) := \begin{cases} \dfrac{H(A \mid O, C)}{H(A \mid C)}, & H(A \mid C) > 0, \\ 0, & H(A \mid C) = 0. \end{cases}$$

This measures the fraction of the C-relative uncertainty of A left unresolved after adding O.

The entropy-normalized complete-trace, hidden-trace, bridge, and memory reconstruction indices are

$$\begin{aligned} \widehat{\mathrm{Ob}}_{\mathrm{complete,ent}}(O) &:= \widehat{\mathrm{Ob}}_{S_{0:T},\mathrm{ent}}(O), \\ \widehat{\mathrm{Ob}}_{\mathrm{trace,ent}}(O) &:= \widehat{\mathrm{Ob}}_{S_{\mathrm{int}},\mathrm{ent}}(O), \\ \widehat{\mathrm{Ob}}_{\mathrm{bridge,ent}}(O) &:= \widehat{\mathrm{Ob}}_{S_{\mathrm{int}},\mathrm{ent}}(O; S_0, S_T), \end{aligned}$$

and

$$\widehat{\mathrm{Ob}}_{\mathrm{mem,ent}}(V) := \widehat{\mathrm{Ob}}_{M_{0:T},\mathrm{ent}}(V_{0:T}).$$

For readability, the shorter symbols $\widehat{\mathrm{Ob}}_{\mathrm{complete}}$, $\widehat{\mathrm{Ob}}_{\mathrm{trace}}$, $\widehat{\mathrm{Ob}}_{\mathrm{bridge}}$, and $\widehat{\mathrm{Ob}}_{\mathrm{mem}}$ may be used only after this entropy-normalization convention has been invoked.

For stochastic target relevance, CEOT IV uses two explicitly named normalizations. The target-normalized index is

$$\widehat{\mathrm{Ob}}^{Y,\mathrm{tar}}_{\mathrm{law}}(O) := \begin{cases} \dfrac{I(Y; S_{0:T} \mid O)}{H(Y \mid O)}, & H(Y \mid O) > 0, \\ 0, & H(Y \mid O) = 0, \end{cases}$$

and the trace-normalized index is

$$\widehat{\mathrm{Ob}}^{Y,\mathrm{tr}}_{\mathrm{law}}(O) := \begin{cases} \dfrac{I(Y; S_{0:T} \mid O)}{H(S_{0:T} \mid O)}, & H(S_{0:T} \mid O) > 0, \\ 0, & H(S_{0:T} \mid O) = 0. \end{cases}$$

The target-normalized version asks what fraction of the residual target uncertainty after O is still explained by the full trace. The trace-normalized version asks what fraction of the residual trace uncertainty is target-relevant.

Proposition 6.23 (Normalized indices are convention-dependent)**.** *The unnormalized obstruction $H(A \mid O)$ is intrinsic to the joint law of (A, O). A normalized index*

$$\widehat{\mathrm{Ob}}_{A,\mathsf{D}}(O)$$

is not determined by the joint law alone unless the denominator convention D is declared as part of the typed profile. Different valid denominators can give different numerical normalized indices for the same pair (A, O).

Proof. The numerator $H(A \mid O)$ is determined by $P_{A,O}$. The denominator may be $H(A)$, $\log_b |\operatorname{supp}(A)|$, or $\log_b |\mathcal{A}|$, and these quantities need not agree. Therefore the ratio is not a canonical scalar unless the denominator convention is specified. □

Proposition 6.24 (Bounds for normalized indices)**.** *All entropy-normalized reconstruction indices above satisfy*

$$0 \le \widehat{\mathrm{Ob}} \le 1.$$

Moreover,

$$0 \le \widehat{\mathrm{Ob}}_{\mathrm{law}}^{Y,\mathrm{tar}}(O) \le 1,$$

and

$$0 \le \widehat{\mathrm{Ob}}_{\mathrm{law}}^{Y,\mathrm{tr}}(O) \le 1.$$

Proof. For reconstruction indices,

$$0 \le H(A \mid O) \le H(A).$$

For the conditional version,

$$0 \le H(A \mid O, C) \le H(A \mid C).$$

For the target-normalized stochastic index,

$$I(Y; S_{0:T} \mid O) \le H(Y \mid O).$$

For the trace-normalized stochastic index,

$$I(Y; S_{0:T} \mid O) \le H(S_{0:T} \mid O).$$

The stated bounds follow from the definitions and the local zero-denominator convention. □

Remark 6.25 (Normalization data belongs to the typed profile)*.* Normalized obstruction indices factor through a typed profile only after the normalization rule N is included. The unnormalized entropy obstruction is law-level. The normalized index is law-level plus convention-level.

Remark 6.26 (Raw and normalized obstructions have different roles)*.* Raw obstructions are primary because exact reconstruction and exact sufficiency are zero statements:

$$H(A \mid O) = 0, \qquad I(Y; S_{0:T} \mid O) = 0.$$

Normalized indices are secondary comparison tools. They should not replace raw obstructions in the main theorems, because normalization can obscure absolute entropy scale and because the denominator convention is part of the typed profile data.

6.13 Approximate CEOT IV obstructions

The exact CEOT IV criteria are zero criteria:

$$H(A \mid O) = 0, \qquad I(Y; S_{0:T} \mid O) = 0.$$

For approximate compression analysis, the same quantities can be read with an explicit tolerance, but entropy tolerance and operational decoder error should be kept distinct. Let A be a finite trace-derived variable. The observation O is called entropy-ε reconstructive for A if

$$H_b(A \mid O) \leq \varepsilon.$$

It is called ε-law-sufficient for the stochastic target Y if

$$I_b(Y; S_{0:T} \mid O) \leq \varepsilon.$$

For a finite loss $\ell : \mathcal{Y} \times \mathcal{D} \to \mathbb{R}$, it is called ε-loss-sufficient if

$$\mathrm{Ob}_\ell^Y(O) \leq \varepsilon.$$

Thus the exact CEOT IV theory is recovered by setting $\varepsilon = 0$.

Definition 6.27 (Operational reconstruction error)**.** Let A and O be finite random variables. The optimal support-relative decoder error of A from O is

$$\delta^*(A \mid O) := \inf_{d:\mathrm{supp}(O)\to\mathrm{supp}(A)} \mathbb{P}[d(O) \neq A].$$

The infimum is attained by any MAP decoder

$$d_{\mathrm{MAP}}(o) \in \arg\max_{a\in\mathrm{supp}(A)} P(A = a \mid O = o),$$

and

$$\delta^*(A \mid O) = \mathbb{E}_O \left[1 - \max_{a\in\mathrm{supp}(A)} P(A = a \mid O)\right].$$

The observation O is operationally δ-reconstructive for A if

$$\delta^*(A \mid O) \leq \delta.$$

Theorem 6.28 (Entropy tolerance gives a MAP-decoder certificate)**.** *Assume entropy is computed in logarithm base $b > 1$. If*

$$H_b(A \mid O) \leq \varepsilon,$$

then the MAP decoder satisfies

$$\mathbb{P}[d_{\mathrm{MAP}}(O) \neq A] = \delta^*(A \mid O) \leq \frac{\varepsilon}{\log_b 2}.$$

Thus small conditional entropy gives an explicit finite decoder with small average error.

Proof. For each positive observation value o, set

$$p_o^* := \max_a P(A = a \mid O = o), \qquad e_o := 1 - p_o^*.$$

The MAP decoder has conditional error e_o on the fiber $O = o$. For every finite probability vector p,

$$H_b(p) \geq (1 - \max_i p_i) \log_b 2.$$

Applying this inequality to $P_{A|O=o}$ gives

$$e_o \leq \frac{H_b(A \mid O = o)}{\log_b 2}.$$

Averaging over $o \in \operatorname{supp}(O)$ yields

$$\delta^*(A \mid O) = \mathbb{E}[e_O] \leq \frac{H_b(A \mid O)}{\log_b 2} \leq \frac{\varepsilon}{\log_b 2}.$$

□

Proposition 6.29 (Fano-type upper bound from decoder error)**.** *Let $m = |\operatorname{supp}(A)|$. If $m \geq 2$ and*

$$\delta^*(A \mid O) \leq \delta,$$

then

$$H_b(A \mid O) \leq h_b(\delta) + \delta \log_b(m-1),$$

where

$$h_b(\delta) := -\delta \log_b \delta - (1-\delta) \log_b(1-\delta)$$

is the binary entropy in base b, with the usual continuous endpoint convention.

Proof. This is the finite Fano inequality applied conditionally and averaged over O, using an optimal decoder attaining $\delta^*(A \mid O)$. It records the converse operational scale: small decoder error forces small conditional entropy up to the usual finite-alphabet Fano term. □

Proposition 6.30 (Small law obstruction implies approximate conditional independence)**.** *Assume entropy and mutual information are computed in logarithm base $b > 1$. If*

$$I_b(Y; S_{0:T} \mid O) \leq \varepsilon,$$

then

$$\mathbb{E}_O \ \mathrm{TV}\Big(P_{Y,S_{0:T}|O}, P_{Y|O} P_{S_{0:T}|O}\Big) \leq \sqrt{\frac{\ln b}{2}\, \varepsilon}.$$

In particular, small stochastic law obstruction means that, averaged over observation fibers, the target and the full trace are close to conditionally independent.

Proof. For each positive-support value o, Pinsker's inequality gives

$$\mathrm{TV}\Big(P_{Y,S_{0:T}|O=o}, P_{Y|O=o} P_{S_{0:T}|O=o}\Big) \leq \sqrt{\frac{\ln b}{2} D_b\Big(P_{Y,S_{0:T}|O=o} \Big\| P_{Y|O=o} P_{S_{0:T}|O=o}\Big)}.$$

Taking expectation over O and using Jensen's inequality for the concave square-root function gives

$$\mathbb{E}_O \ \mathrm{TV}(\cdots) \leq \sqrt{\frac{\ln b}{2}\, \mathbb{E}_O D_b(\cdots)}.$$

The expected conditional divergence is exactly $I_b(Y; S_{0:T} \mid O)$. □

Theorem 6.31 (Conditional total variation controls bounded-loss degradation)**.** *Assume*

$$0 \le \ell(y,d) \le L_{\max}$$

for all y, d. Let

$$\mathrm{Ob}_\ell^Y(O) := R_\ell(O) - R_\ell(S_{0:T}),$$

where $R_\ell(Z)$ denotes the Bayes risk when decisions may depend on Z. Then

$$0 \le \mathrm{Ob}_\ell^Y(O) \le 2L_{\max}\,\mathbb{E}_O\,\mathrm{TV}\Big(P_{Y,S_{0:T}|O}, P_{Y|O}P_{S_{0:T}|O}\Big)\,.$$

Here $\mathrm{TV}(P,Q) = \frac{1}{2}\|P-Q\|_1$ is the normalized total-variation distance.

Proof. The lower bound follows because a decision rule based on the full trace contains every decision rule based only on O as a special case.

Fix a positive observation value o and set

$$P_o := P_{Y,S_{0:T}|O=o}, \qquad Q_o := P_{Y|O=o}P_{S_{0:T}|O=o}.$$

Under Q_o, the target is conditionally independent of the full trace once $O = o$ is fixed, so access to $S_{0:T}$ gives no Bayes-risk advantage over access to O alone. For any decision rule depending on $S_{0:T}$, the expectation of a loss bounded by $L_{\max}$ changes by at most

$$2L_{\max}\,\mathrm{TV}(P_o, Q_o)$$

when the underlying law is changed from Q_o to P_o. Taking the infimum over full-trace decision rules and averaging over positive observation fibers gives the displayed upper bound. ☐

Corollary 6.32 (Approximate law-sufficiency controls bounded-loss degradation)**.** *Assume entropy and mutual information are computed in base $b > 1$, and assume*

$$0 \le \ell(y,d) \le L_{\max}$$

for all y, d. If

$$I_b(Y; S_{0:T} \mid O) \le \varepsilon,$$

then

$$\mathrm{Ob}_\ell^Y(O) \le 2L_{\max}\sqrt{\frac{\ln b}{2}\,\varepsilon}.$$

Proof. Apply Proposition 6.30 and then Theorem 6.31. ☐

Theorem 6.33 (Approximate obstruction interface theorem)**.** *Let $A, O, S_{0:T}, Y$ be finite random variables in a CEOT IV presentation, and compute entropy and mutual information in base $b > 1$. The following are the approximate implication interfaces used in this paper.*

(i) Reconstruction scale. If

$$H_b(A \mid O) \le \varepsilon,$$

then the MAP decoder satisfies

$$\delta^*(A \mid O) \le \frac{\varepsilon}{\log_b 2}.$$

(ii) Converse reconstruction scale. If $m = |\operatorname{supp}(A)| \geq 2$ and

$$\delta^*(A \mid O) \leq \delta,$$

then

$$H_b(A \mid O) \leq h_b(\delta) + \delta \log_b(m-1).$$

(iii) Law-sufficiency scale. If

$$I_b(Y; S_{0:T} \mid O) \leq \varepsilon,$$

then

$$\mathbb{E}_O \operatorname{TV}\Big(P_{Y,S_{0:T}|O}, P_{Y|O}P_{S_{0:T}|O}\Big) \leq \sqrt{\frac{\ln b}{2}\varepsilon}.$$

(iv) Bounded-loss scale. If additionally

$$0 \leq \ell(y, d) \leq L_{\max},$$

then

$$\mathrm{Ob}_\ell^Y(O) \leq 2L_{\max}\sqrt{\frac{\ln b}{2}\varepsilon}.$$

No other implication between approximate reconstruction, approximate law-sufficiency, and approximate fixed-loss sufficiency is asserted without additional assumptions.

Proof. The four items are respectively Theorem 6.28, Proposition 6.29, Proposition 6.30, and Corollary 6.32. The final sentence is a scope statement: the examples and remarks below show that the three approximate scales measure different properties. □

Remark 6.34 (No single approximate obstruction scale). There is no single canonical ε-scale that simultaneously controls trace reconstruction, stochastic law-sufficiency, and loss-relative performance. The same numerical value ε has different meanings in entropy units, total-variation units, and loss units. Approximate statements in CEOT IV must therefore specify which scale is being used and which interface theorem converts one scale into another.

Example 6.35 (Law-sufficiency without trace reconstruction)**.** Let Y be independent of $S_{0:T}$, and let O be constant. Then

$$I(Y; S_{0:T} \mid O) = 0,$$

so O is exactly law-sufficient for Y. But if $S_{0:T}$ is non-degenerate, then

$$H(S_{0:T} \mid O) = H(S_{0:T}) > 0.$$

Thus target law-sufficiency gives no reconstruction guarantee for the full trace.

Example 6.36 (Fixed-loss sufficiency without approximate law-sufficiency)**.** Example 6.15 has

$$\mathrm{Ob}_{\ell_{0\text{-}1}}^Y(O) = 0 \qquad \text{but} \qquad I(Y; S \mid O) > 0.$$

Hence zero fixed-loss obstruction for a declared loss does not imply exact or approximate law-sufficiency.

Remark 6.37 (Average error, not pointwise reconstruction). The MAP bound derived from $H_b(A \mid O) \le \varepsilon$ is an average decoder-error guarantee:

$$\mathbb{P}[d_{\mathrm{MAP}}(O) \neq A] \le \frac{\varepsilon}{\log_b 2}.$$

It is not a uniform fiberwise guarantee unless additional lower bounds on observation-fiber probabilities, or direct pointwise entropy bounds on every positive fiber, are assumed. A rare observation fiber may carry large posterior uncertainty while the average conditional entropy remains small.

Remark 6.38 (Role of approximate obstructions). Positive obstruction means exact failure. Small positive obstruction has a quantitative interpretation only after choosing the relevant scale: entropy units for reconstruction, conditional total variation for law-sufficiency, and loss scale for decision tasks. CEOT IV therefore keeps exact zero criteria as the structural core and uses approximate indices as operational refinements. In particular, approximate reconstruction, approximate law-sufficiency, and approximate loss-sufficiency are distinct notions connected only by the explicitly stated interface bounds.

6.14 Section summary

This section replaced deterministic target recovery by stochastic law-sufficiency. For a target kernel

$$L : \mathcal{S}_{0:T} \rightsquigarrow \mathcal{Y},$$

the full-law obstruction is

$$\mathrm{Ob}^{Y}_{\mathrm{law}}(O) = I(Y; S_{0:T} \mid O).$$

It vanishes exactly when

$$Y \perp S_{0:T} \mid O,$$

or equivalently, since $O = c(S_{0:T})$,

$$P(Y \mid S_{0:T}) = P(Y \mid O)$$

on support. The coarsest law-sufficient object that must be preserved is the canonical law quotient

$$Q_L = \mathcal{L}(Y \mid S_{0:T}),$$

and O is law-sufficient exactly when this quotient factors through O. Loss-relative sufficiency is weaker: full-law sufficiency implies zero loss-relative obstruction for every finite loss, while the converse may fail for a fixed loss. The section also introduced normalized obstruction indices for cross-system comparison while keeping raw entropy and mutual information as the primary exact-certification quantities. Deterministic recovery is the degenerate-kernel case in which

$$I(Y; S_{0:T} \mid O) = H(Y \mid O).$$

The next section studies how complete-trace, hidden-trace, deterministic-target, stochastic law-sufficiency, and loss-relative obstructions behave under deterministic refinement of observations and under finite profile isomorphism.

7 Observation Refinement, Monotonicity, and Profile Invariance

7.1 Purpose of the section

Sections 4–6 defined the main CEOT IV obstructions for a fixed deterministic compressed observation

$$O = c(S_{0:T}).$$

This section studies what happens when the observation is refined. If a second deterministic observation

$$O' = c'(S_{0:T})$$

contains at least as much information about the trace as O, then every reconstruction or sufficiency obstruction should weakly decrease. This is the monotonicity side of CEOT IV.

The section also separates this order-theoretic statement from the finite categorical invariance statement. Refinement is a poset relation between observations of a fixed trace system. Presentation isomorphism is a groupoid-level relabeling relation between finite trace-compression presentations. These two structures serve different purposes and should not be conflated.

The section has four goals:

(i) define the deterministic observation-refinement poset;

(ii) prove monotonicity of complete-trace, hidden-trace, bridge, deterministic-target, and stochastic-target obstructions;

(iii) define presentation and profile isomorphisms;

(iv) prove that the scalar obstructions factor through finite posterior profile isomorphism classes.

7.2 Deterministic observation-refinement order

Fix a finite trace law on $S_{0:T}$. Let

$$O = c(S_{0:T}), \qquad O' = c'(S_{0:T})$$

be deterministic finite observations of the same trace.

Definition 7.1 (Deterministic observation refinement)**.** We say that O' *refines* O, and write

$$O \preceq O',$$

if

$$\sigma(O) \subseteq \sigma(O') \subseteq \sigma(S_{0:T}).$$

Equivalently, O' contains at least as much finite trace information as O.

In the finite setting this condition has a concrete fiberwise form.

Lemma 7.2 (Finite factor criterion for refinement)**.** *For deterministic finite observations $O = c(S_{0:T})$ and $O' = c'(S_{0:T})$, the following are equivalent:*

(a) $O \preceq O'$;

(b) there exists a map

$$r : \mathrm{supp}(O') \to \mathrm{supp}(O)$$

such that

$$O = r(O') \quad a.s.;$$

(c) every positive-support fiber of O' is contained in a positive-support fiber of O.

Proof. For finite random variables, inclusion of generated σ-algebras is equivalent to one variable being a measurable function of the other. Since the spaces are finite, measurability is the same as being a set-theoretic function on support. Thus $\sigma(O) \subseteq \sigma(O')$ if and only if $O = r(O')$ a.s. for some finite map r. The fiber statement is the same condition written in terms of the partitions induced by O and O'. $\square$

Lemma 7.3 (Finite conditional entropy monotonicity under refinement)**.** *Let A, O, O' be finite random variables. If*

$$O \preceq O',$$

meaning that $O = r(O')$ almost surely for some support-relative finite map r, then

$$H(A \mid O') \leq H(A \mid O).$$

More generally, if B is any finite random variable jointly distributed with (A, O, O'), then

$$H(A \mid O', B) \leq H(A \mid O, B).$$

Proof. By Lemma 7.2, O is a finite function of O' on support. Hence the conditioning partition generated by O' refines the conditioning partition generated by O. Finite conditional entropy is the average entropy of the conditional laws over these fibers. Refining a conditioning partition cannot increase that average. The same argument applies after adjoining the common finite variable B to both conditioning variables. $\square$

Definition 7.4 (Deterministic observation-refinement category)**.** Fix a finite trace law on $S_{0:T}$. The deterministic observation-refinement category

$$\mathsf{Obs}_{\mathrm{ref}}(S_{0:T})$$

has as objects deterministic maps

$$c : \mathcal{S}_{0:T} \to \mathcal{O}_c$$

viewed through their induced observations $O_c = c(S_{0:T})$. A morphism

$$c \longrightarrow c'$$

is a support-relative finite factor map

$$r_{c,c'} : \mathrm{supp}(O_{c'}) \to \mathrm{supp}(O_c)$$

such that

$$O_c = r_{c,c'}(O_{c'}) \quad \text{a.s.}$$

Equivalently,

$$O_c \preceq O_{c'},$$

so the target observation c' is at least as informative as the source observation c. If

$$c \longrightarrow c' \quad \text{and} \quad c' \longrightarrow c'',$$

with factor maps $r_{c,c'}$ and $r_{c',c''}$, their composite is represented by

$$r_{c,c'} \circ r_{c',c''} : \mathrm{supp}(O_{c''}) \to \mathrm{supp}(O_c).$$

Identity morphisms are represented by identity maps on positive observation supports. Thus $\mathsf{Obs}_{\mathrm{ref}}(S_{0:T})$ is the category associated with the finite observation-refinement preorder, with arrows directed from coarser observations to finer observations.

Convention 7.5 (Ordered target category for obstruction values)**.** The symbol

$$(\mathbb{R}_{\geq 0}, \geq)$$

denotes the thin category whose objects are nonnegative real numbers and whose unique morphism

$$a \longrightarrow b$$

exists exactly when

$$a \geq b.$$

Thus a numerical obstruction map that decreases under refinement becomes a covariant functor into $(\mathbb{R}_{\geq 0}, \geq)$. Equivalently, the same map is antitone as an ordinary map into the usual ordered set $(\mathbb{R}_{\geq 0}, \leq)$.

Proposition 7.6 (Obstructions as covariant functors into the reverse order)**.** *Fix a finite trace-target law and let*

$$\mathsf{Obs}_{\mathrm{ref}}(S_{0:T})$$

be the coarser-to-finer deterministic observation-refinement category. For any trace-derived or target-derived finite variable A, the map

$$c \longmapsto H(A \mid O_c)$$

defines a covariant functor

$$\mathrm{Ob}_A : \mathsf{Obs}_{\mathrm{ref}}(S_{0:T}) \longrightarrow (\mathbb{R}_{\geq 0}, \geq).$$

Equivalently, if $c \to c'$ means that c' refines c, then

$$H(A \mid O_c) \geq H(A \mid O_{c'}).$$

The same reverse-order convention applies to the complete-trace, hidden-trace, bridge, deterministic-target, and stochastic-target obstruction maps whenever the corresponding obstruction variable or target kernel is fixed.

Proof. A morphism $c \to c'$ means that O_c is a deterministic function of $O_{c'}$ on positive support, equivalently

$$\sigma(O_c) \subseteq \sigma(O_{c'}).$$

Conditional entropy decreases under refinement:

$$H(A \mid O_{c'}) \leq H(A \mid O_c).$$

Hence in the reverse ordered category $(\mathbb{R}_{\geq 0}, \geq)$ there is a morphism

$$H(A \mid O_c) \longrightarrow H(A \mid O_{c'}).$$

Identity and composition follow from the thin-category structure of the preorder.

$$\begin{array}{ccc} c & \longrightarrow & c' \\ \downarrow \mathrm{Ob}_A & & \downarrow \mathrm{Ob}_A \\ H(A \mid O_c) & \longrightarrow & H(A \mid O_{c'}) \end{array}$$ where the bottom arrow exists because $H(A \mid O_c) \geq H(A \mid O_{c'})$.

□

Remark 7.7 (Covariant only after reversing the numerical order). The obstruction map is not covariant into the usual ordered category $(\mathbb{R}_{\geq 0}, \leq)$. It is antitone with respect to the usual numerical order and covariant only after the codomain order is reversed. This is the reason the codomain is written as $(\mathbb{R}_{\geq 0}, \geq)$ rather than $(\mathbb{R}_{\geq 0}, \leq)$.

The coordinatewise roadmap case is recovered as follows. Suppose

$$O = \bar{S}_{0:T} = (q_0(S_0), \ldots, q_T(S_T))$$

and

$$O' = \bar{S}'_{0:T} = (q'_0(S_0), \ldots, q'_T(S_T)).$$

If every q_t factors through q'_t on support, then

$$\bar{S}_{0:T} \preceq \bar{S}'_{0:T}.$$

The general definition allows also global summaries $O' = c'(S_{0:T})$ that are not coordinatewise.

7.3 Trace-obstruction monotonicity

The basic monotonicity principle is that conditioning on a finer observation cannot increase conditional entropy.

Theorem 7.8 (Trace-obstruction monotonicity). *Let O and O' be deterministic finite observations of $S_{0:T}$ with*

$$O \preceq O'.$$

Then:

$$H(S_{0:T} \mid O') \leq H(S_{0:T} \mid O),$$

$$H(S_{\text{int}} \mid O') \leq H(S_{\text{int}} \mid O),$$

and

$$H(S_{\text{int}} \mid O', S_0, S_T) \leq H(S_{\text{int}} \mid O, S_0, S_T).$$

Equivalently,

$$\mathrm{Ob}_{\text{complete}}(O') \leq \mathrm{Ob}_{\text{complete}}(O),$$

$$\mathrm{Ob}_{\text{trace}}(O') \leq \mathrm{Ob}_{\text{trace}}(O),$$

and

$$\mathrm{Ob}_{\text{bridge}}(O') \leq \mathrm{Ob}_{\text{bridge}}(O).$$

Proof. Apply Lemma 7.3 with $A = S_{0:T}$ and then with $A = S_{\text{int}}$. This gives

$$H(S_{0:T} \mid O') \le H(S_{0:T} \mid O)$$

and

$$H(S_{\text{int}} \mid O') \le H(S_{\text{int}} \mid O).$$

For the bridge term, apply the same lemma with $A = S_{\text{int}}$ and with the common additional conditioning variable $B = (S_0, S_T)$:

$$H(S_{\text{int}} \mid O', S_0, S_T) \le H(S_{\text{int}} \mid O, S_0, S_T).$$

□

Corollary 7.9 (Order-reversing trace obstruction map)**.** *For a fixed finite trace law, the maps*

$$O \mapsto \mathrm{Ob}_{\text{complete}}(O), \qquad O \mapsto \mathrm{Ob}_{\text{trace}}(O), \qquad O \mapsto \mathrm{Ob}_{\text{bridge}}(O)$$

are antitone on the deterministic observation-refinement poset.

This is the formal version of the intuitive statement that a finer log, summary, or compressed trace cannot make reconstruction harder. It may fail to make reconstruction complete, but it cannot increase the finite entropy obstruction.

7.4 Deterministic target monotonicity

Now suppose the target is support-relative deterministic:

$$Y = f_+(S_{0:T}).$$

The deterministic target obstruction is

$$\mathrm{Ob}^{Y}_{\text{det}}(O) = H(Y \mid O).$$

It is monotone under observation refinement.

Theorem 7.10 (Deterministic target monotonicity)**.** *Let $Y = f_+(S_{0:T})$ be a support-relative deterministic target. If*

$$O \preceq O',$$

then

$$H(Y \mid O') \le H(Y \mid O).$$

Equivalently,

$$\mathrm{Ob}^{Y}_{\text{det}}(O') \le \mathrm{Ob}^{Y}_{\text{det}}(O).$$

Proof. Apply Lemma 7.3 with $A = Y$. The deterministic assumption is not needed for this entropy inequality itself; it is needed only for interpreting $H(Y \mid O)$ as a deterministic target-recovery obstruction. □

Remark 7.11 (Zero recovery is preserved by refinement)*.* If

$$H(Y \mid O) = 0$$

and $O \preceq O'$, then also

$$H(Y \mid O') = 0.$$

Thus once a deterministic target is recoverable from a compressed observation, it remains recoverable from every deterministic refinement of that observation.

7.5 Stochastic target monotonicity

For stochastic targets one must state the hypotheses explicitly. Let

$$Y \sim L(\cdot \mid S_{0:T})$$

be generated by a finite target kernel, and let O, O' be deterministic functions of $S_{0:T}$ with

$$O \preceq O'.$$

Then the Markov structure is

$$O, O' \leftarrow S_{0:T} \rightarrow Y.$$

Equivalently,

$$P(Y \mid S_{0:T}, O, O') = P(Y \mid S_{0:T})$$

on support.

Theorem 7.12 (Stochastic target monotonicity)**.** *Assume $Y \sim L(\cdot \mid S_{0:T})$, and let O, O' be deterministic observations with*

$$O \preceq O'.$$

Then

$$I(Y; S_{0:T} \mid O') \leq I(Y; S_{0:T} \mid O).$$

Equivalently,

$$\mathrm{Ob}^{Y}_{\mathrm{law}}(O') \leq \mathrm{Ob}^{Y}_{\mathrm{law}}(O).$$

Proof. Write $S = S_{0:T}$. Since O and O' are deterministic functions of S, Lemma 6.1 gives

$$H(Y \mid S, O) = H(Y \mid S) \quad \text{and} \quad H(Y \mid S, O') = H(Y \mid S).$$

Therefore

$$I(Y; S \mid O) = H(Y \mid O) - H(Y \mid S)$$

and

$$I(Y; S \mid O') = H(Y \mid O') - H(Y \mid S).$$

By Lemma 7.3 applied to $A = Y$,

$$H(Y \mid O') \leq H(Y \mid O),$$

because $O \preceq O'$. Subtracting the common term $H(Y \mid S)$ gives the result. □

Corollary 7.13 (Zero stochastic sufficiency is preserved by refinement)**.** *If*

$$I(Y; S_{0:T} \mid O) = 0$$

and $O \preceq O'$, then

$$I(Y; S_{0:T} \mid O') = 0.$$

Proof. Conditional mutual information is nonnegative. By Theorem 7.12,

$$0 \leq I(Y; S_{0:T} \mid O') \leq I(Y; S_{0:T} \mid O) = 0.$$

□

Thus deterministic refinement cannot destroy stochastic target sufficiency. It may remove additional irrelevant hidden trace uncertainty, but it cannot reintroduce target-relevant information loss once the target law has already been screened off.

7.6 Profile monotonicity versus profile non-completeness

Monotonicity concerns changing the observation within a fixed trace law. It says that if O' refines O, then the scalar obstructions decrease. This is different from profile completeness. The scalar value

$$H(A \mid O)$$

does not determine the posterior profile

$$\mathsf{Post}_{A|O} = \Big(P_O, \{P_{A|O=o}\}_{o \in \mathrm{supp}(O)}\Big).$$

Nor does the observed target law $P_{O,Y}$, even up to law isomorphism, determine a hidden-trace profile. Section 4 already isolated the scalar obstructions, and Section 3 emphasized that each scalar is only a shadow of a posterior profile.

The point of the present section is narrower: once a particular presentation and a particular target are fixed, deterministic refinement cannot increase the obstruction. The later non-completeness section will show that two different presentations can share the isomorphic compressed target behavior while having different hidden trace obstructions. These are compatible facts:

$$\text{refinement gives monotonicity inside one presentation,}$$

whereas

$$\text{non-completeness compares different presentations or profiles.}$$

7.7 Support-relative presentation isomorphism

We next isolate finite relabeling invariance. Since every CEOT IV obstruction is computed from positive trace support, observation support, bridge support, and the target kernel restricted to positive trace support, the relevant notion of presentation isomorphism is support-relative rather than ambient. Null states, null transitions, null observations, and target values outside positive support do not affect the obstruction theory.

Let

$$\mathfrak{A} = (T, \{\mathcal{S}_t\}, \mu_0, \{K_t\}, \mathcal{O}, c, \mathcal{Y}, L)$$

and

$$\mathfrak{A}' = (T, \{\mathcal{S}'_t\}, \mu'_0, \{K'_t\}, \mathcal{O}', c', \mathcal{Y}', L')$$

be finite deterministic trace-compression presentations with the same time horizon T. Write

$$\mathcal{T}_+ = \mathrm{supp}_{\mathfrak{A}}(S_{0:T}), \qquad \mathcal{T}'_+ = \mathrm{supp}_{\mathfrak{A}'}(S'_{0:T}),$$

and let

$$\rho_I, \rho_O, \rho_B$$

denote the internal-trace, observation, and bridge maps induced by $\mathfrak{A}$ on $\mathcal{T}_+$, with primed notation for $\mathfrak{A}'$.

Definition 7.14 (Support-relative presentation isomorphism)**.** A *support-relative presentation isomorphism*

$$\mathfrak{A} \cong_{\mathrm{sr}} \mathfrak{A}'$$

consists of bijections

$$\Phi : \mathcal{T}_+ \to \mathcal{T}'_+, \qquad \alpha_I : \mathrm{supp}_{\mathfrak{A}}(S_{\mathrm{int}}) \to \mathrm{supp}_{\mathfrak{A}'}(S'_{\mathrm{int}}), \qquad \psi : \mathrm{supp}_{\mathfrak{A}}(O) \to \mathrm{supp}_{\mathfrak{A}'}(O'),$$

$$\gamma : \mathrm{supp}_{\mathfrak{A}}(B) \to \mathrm{supp}_{\mathfrak{A}'}(B'), \qquad \eta : \mathrm{supp}_{\mathfrak{A}}(Y) \to \mathrm{supp}_{\mathfrak{A}'}(Y'),$$

such that, for every $\tau \in \mathcal{T}_+$ and every $y \in \mathrm{supp}_{\mathfrak{A}}(Y)$,

$$P_{\mathfrak{A}'}(S'_{0:T} = \Phi(\tau)) = P_{\mathfrak{A}}(S_{0:T} = \tau),$$

$$\alpha_I \circ \rho_I = \rho'_I \circ \Phi, \qquad \psi \circ \rho_O = \rho'_O \circ \Phi, \qquad \gamma \circ \rho_B = \rho'_B \circ \Phi,$$

and

$$L'_+(\eta(y) \mid \Phi(\tau)) = L_+(y \mid \tau).$$

No condition is imposed on null trace values, null observation values, null transitions, or ambient target values outside positive support.

Proposition 7.15 (Support-relative joint-law invariance)**.** *If*

$$\mathfrak{A} \cong_{\mathrm{sr}} \mathfrak{A}',$$

then the positive-support joint laws of

$$(S_{0:T}, O, Y) \qquad \textit{and} \qquad (S'_{0:T}, O', Y')$$

are carried to one another by

$$(\tau, o, y) \mapsto (\Phi(\tau), \psi(o), \eta(y)).$$

Consequently, all conditional entropy and conditional mutual information obstructions defined in Sections 4–6 are preserved.

Proof. The first condition transports the positive trace law. The commuting equation with ρ_O transports the deterministic observation. The target-kernel equation transports the conditional law of the target on positive trace support. Hence the induced positive-support joint law of $(S_{0:T}, O, Y)$ is transported to that of $(S'_{0:T}, O', Y')$ by the displayed bijection.

Conditional entropy and conditional mutual information are finite joint-law invariants under bijective relabeling. Hence the scalar CEOT IV obstructions are preserved. □

Remark 7.16 (Ambient versus support-relative isomorphism)*.* A stronger ambient presentation isomorphism, requiring equality of initial laws, transition kernels, compression maps, and target kernels on the full ambient spaces, is sufficient for CEOT IV invariance. It is not necessary. The obstruction theory is support-relative: it depends only on the positive trace law, the positive observation and bridge fibers, and the target kernel restricted to positive trace support.

7.8 Trace-target profile objects

Support-relative presentation isomorphism is still stronger than equality of obstruction profiles. Two different presentations may induce the same posterior reconstruction data even when they are not support-relatively isomorphic. The correct profile-level object is therefore defined directly from the induced joint law of $(S_{0:T}, O, Y)$.

Let

$$B := (O, S_0, S_T)$$

be the bridge-conditioning variable and write

$$\begin{aligned} &\mathcal{T}_+ := \mathrm{supp}(S_{0:T}), \quad \mathcal{I}_+ := \mathrm{supp}(S_{\mathrm{int}}), \quad \mathcal{O}_+ := \mathrm{supp}(O), \\ &\mathcal{B}_+ := \mathrm{supp}(B), \qquad \mathcal{Y}_+ := \mathrm{supp}(Y). \end{aligned}$$

Definition 7.17 (Coherent finite CEOT IV trace-target profile)**.** The coherent finite CEOT IV trace-target profile induced by a presentation is the tuple

$$\Pi_{\mathrm{coh}}(\mathfrak{A}) = (T, \mathcal{T}_+, \mathcal{I}_+, \mathcal{O}_+, \mathcal{B}_+, \mathcal{Y}_+, P_{\mathcal{T}}, \rho_I, \rho_O, \rho_B, L_+),$$

where $P_{\mathcal{T}}$ is the positive-support trace law on $\mathcal{T}_+ = \mathrm{supp}(S_{0:T})$,

$$\rho_I : \mathcal{T}_+ \to \mathcal{I}_+, \qquad \rho_O : \mathcal{T}_+ \to \mathcal{O}_+, \qquad \rho_B : \mathcal{T}_+ \to \mathcal{B}_+$$

are the support-relative internal-trace, observation, and bridge projection maps, and

$$L_+(\cdot \mid \tau) = P_{Y|S_{0:T}=\tau} \qquad (\tau \in \mathcal{T}_+)$$

is the target kernel restricted to positive trace support.

The posterior laws used by the scalar obstructions are derived data. Namely,

$$P_O = (\rho_O)_* P_{\mathcal{T}}, \qquad P_B = (\rho_B)_* P_{\mathcal{T}},$$

$$\pi_o^{\mathrm{complete}}(\tau) = P_{\mathcal{T}}(\tau \mid \rho_O(\tau) = o),$$
$$\pi_o^{\mathrm{trace}} = (\rho_I)_* \pi_o^{\mathrm{complete}},$$
$$\pi_b^{\mathrm{bridge}} = (\rho_I)_* P_{\mathcal{T}}(\,\cdot \mid \rho_B(\tau) = b),$$

and

$$\pi_o^Y(y) = \sum_{\tau:\rho_O(\tau)=o} \pi_o^{\mathrm{complete}}(\tau) L_+(y \mid \tau).$$

All displayed conditional laws are evaluated only on positive-support fibers.

The inclusion of L_+ is essential. The marginal posterior target profile $\{P_{Y|O=o}\}$ alone determines $H(Y \mid O)$, but it does not determine the stochastic sufficiency obstruction

$$I(Y; S_{0:T} \mid O),$$

which compares the fiberwise full-trace target law $P_{Y|S_{0:T}}$ with the observation-level target law $P_{Y|O}$. The structural maps ρ_I, ρ_O, ρ_B are also essential: they enforce that complete-trace, internal-trace, bridge, and target posterior profiles are coherent shadows of one trace law, rather than independent posterior families.

Remark 7.18 (Elementary posterior profiles versus coherent trace profiles)*.* The elementary posterior profile of a pair (A, O) only has the relabeling structure of the two supports $\mathrm{supp}(A)$ and $\mathrm{supp}(O)$. This is sufficient for scalar examples involving one pair of random variables. By contrast, the CEOT IV trace-target profile is typed and coherent: complete trace, internal trace, observation, bridge, and target kernel all live over the same positive trace support. Therefore its isomorphisms must preserve the projection maps ρ_I, ρ_O, ρ_B and the restricted kernel L_+.

7.9 Profile isomorphism

Definition 7.19 (Coherent trace-target profile isomorphism)**.** Let Π_{coh} and Π'_{coh} be coherent finite CEOT IV trace-target profiles with the same time horizon. A coherent trace-target profile isomorphism

$$\Pi_{\mathrm{coh}} \cong \Pi'_{\mathrm{coh}}$$

consists of bijections

$$\alpha_{\mathcal{T}} : \mathcal{T}_+ \to \mathcal{T}'_+, \qquad \alpha_{\mathcal{I}} : \mathcal{I}_+ \to \mathcal{I}'_+, \qquad \beta : \mathcal{O}_+ \to \mathcal{O}'_+,$$

$$\gamma : \mathcal{B}_+ \to \mathcal{B}'_+, \qquad \eta : \mathcal{Y}_+ \to \mathcal{Y}'_+,$$

such that the trace law is transported,

$$P'_{\mathcal{T}}(\alpha_{\mathcal{T}}(\tau)) = P_{\mathcal{T}}(\tau),$$

the structural maps commute,

$$\alpha_{\mathcal{I}} \circ \rho_I = \rho'_I \circ \alpha_{\mathcal{T}},$$

$$\beta \circ \rho_O = \rho'_O \circ \alpha_{\mathcal{T}},$$

$$\gamma \circ \rho_B = \rho'_B \circ \alpha_{\mathcal{T}},$$

and the target kernels are transported,

$$L'_+(\eta(y) \mid \alpha_{\mathcal{T}}(\tau)) = L_+(y \mid \tau)$$

for all $\tau \in \mathcal{T}_+$ and $y \in \mathcal{Y}_+$.

Lemma 7.20 (Projection compatibility is mandatory)**.** *Let Π_{coh} and Π'_{coh} be coherent trace-target profiles. A relabeling of the individual positive-support sets preserves the CEOT IV trace-target profile only if it commutes with the projection maps:*

$$\alpha_{\mathcal{I}} \circ \rho_I = \rho'_I \circ \alpha_{\mathcal{T}},$$

$$\beta \circ \rho_O = \rho'_O \circ \alpha_{\mathcal{T}},$$

$$\gamma \circ \rho_B = \rho'_B \circ \alpha_{\mathcal{T}}.$$

Without these equations, the relabeled data need not represent the same internal-trace, observation, and bridge structure over a common complete-trace support.

Proof. The maps ρ_I, ρ_O, ρ_B are part of the coherent profile data. They specify how internal traces, observations, and bridges are obtained from the same positive complete-trace support. If a relabeling fails to commute with one of these maps, then a complete trace τ and its relabeled image $\alpha_{\mathcal{T}}(\tau)$ are assigned incompatible internal, observation, or bridge labels. Hence the relabeling does not preserve the coherent profile object. □

Proposition 7.21 (The coherent trace-target profiles form a groupoid)**.** *Coherent finite CEOT IV trace-target profiles and coherent trace-target profile isomorphisms form a groupoid.*

Proof. The identity morphism is given by identity bijections on all positive-support sets. The composite of two profile isomorphisms is obtained by composing the corresponding bijections. Preservation of $P_{\mathcal{T}}$, commutativity with ρ_I, ρ_O, ρ_B, and transport of L_+ are all stable under composition. Since all morphisms are bijective, every morphism has an inverse satisfying the same equations. □

Proposition 7.22 (Presentation isomorphism implies coherent profile isomorphism)**.** *Every support-relative presentation isomorphism induces a coherent trace-target profile isomorphism.*

Proof. By Proposition 7.15, a support-relative presentation isomorphism transports the positive-support trace law and the full joint law of $(S_{0:T}, O, Y)$ to the corresponding relabeled laws. The relabeling of the complete trace commutes with the internal-trace, observation, and bridge projection maps, and the target-kernel preservation condition transports L_+. Hence the induced relabelings satisfy every equation in Definition 7.19. $\square$

The converse is generally false. A coherent trace-target profile may forget how the trace law was generated by a particular chain of transition kernels. This is intentional. CEOT IV measures reconstruction and sufficiency after the finite presentation has induced the relevant trace-target law; it does not assert that the coherent profile remembers the entire algorithmic generator.

7.10 Groupoid-level factorization

Let

$$\mathbf{AlgCompPres}^{\cong}_{\mathrm{fin}}$$

denote the groupoid of finite deterministic trace-compression presentations and support-relative presentation isomorphisms. Let

$$\mathbf{TraceTargetProf}^{\cong}_{\mathrm{fin,coh}}$$

denote the groupoid of coherent finite CEOT IV trace-target profiles and coherent profile isomorphisms.

The construction sending a presentation to its coherent trace-target profile defines a map of groupoids:

$$\Pi_{\mathrm{coh}} : \mathbf{AlgCompPres}^{\cong}_{\mathrm{fin}} \longrightarrow \mathbf{TraceTargetProf}^{\cong}_{\mathrm{fin,coh}}.$$

Theorem 7.23 (Coherent profile generation of the CEOT IV obstruction family)**.** *For a coherent CEOT IV trace-target profile*

$$\Pi_{\mathrm{coh}} = (T, \mathcal{T}_+, \mathcal{I}_+, \mathcal{O}_+, \mathcal{B}_+, \mathcal{Y}_+, P_{\mathcal{T}}, \rho_I, \rho_O, \rho_B, L_+),$$

the complete-trace, hidden-trace, bridge, deterministic-target, and stochastic-target obstructions are determined by the finite functionals

$$\mathrm{Ob}_{\mathrm{complete}} = H(S_{0:T} \mid O), \qquad O = \rho_O(S_{0:T}),$$

$$\mathrm{Ob}_{\mathrm{trace}} = H(\rho_I(S_{0:T}) \mid \rho_O(S_{0:T})),$$

$$\mathrm{Ob}_{\mathrm{bridge}} = H(\rho_I(S_{0:T}) \mid \rho_B(S_{0:T})),$$

for support-relative deterministic targets $Y = f_+(S_{0:T})$*,*

$$\mathrm{Ob}^{Y}_{\mathrm{det}} = H(f_+(S_{0:T}) \mid \rho_O(S_{0:T})),$$

and for stochastic targets with kernel L_+*,*

$$\mathrm{Ob}^{Y}_{\mathrm{law}} = \sum_{\tau \in \mathcal{T}_+} P_{\mathcal{T}}(\tau) D_{\mathrm{KL}} \left(L_+(\cdot \mid \tau) \middle\| \sum_{\tau' : \rho_O(\tau') = \rho_O(\tau)} P_{\mathcal{T}}(\tau' \mid \rho_O(\tau') = \rho_O(\tau)) L_+(\cdot \mid \tau') \right).$$

Therefore the scalar CEOT IV obstruction family factors through coherent trace-target profile isomorphism.

Proof. All displayed quantities are computed from the finite law $P_{\mathcal{T}}$, the structural maps ρ_O, ρ_I, ρ_B, and the target kernel L_+. The posterior complete-trace, internal-trace, bridge, and target laws are finite conditional laws derived from these primitive coherent data. Hence each obstruction is a functional of the coherent profile.

If two coherent profiles are isomorphic, the isomorphism preserves $P_{\mathcal{T}}$, commutes with ρ_O, ρ_I, ρ_B, and transports L_+. Entropy, conditional entropy, conditional mutual information, and finite relative entropy are invariant under bijective relabeling of finite supports. Thus all scalar obstruction values are preserved. □

7.11 Typed profile extensions for loss and memory obstructions

The coherent trace-target profile

$$\Pi_{\mathrm{coh}}$$

is sufficient for the reconstruction and law-sufficiency obstruction family. Loss-relative, memory-relative, and normalized obstruction values require additional typed data. CEOT IV therefore uses the following profile ladder.

Definition 7.24 (Typed CEOT IV profile ladder)**.** Let

$$\Pi_{\mathrm{rec/law}} := (T, \mathcal{T}_+, \mathcal{I}_+, \mathcal{O}_+, \mathcal{B}_+, \mathcal{Y}_+, P_{\mathcal{T}}, \rho_I, \rho_O, \rho_B, L_+)$$

be the reconstruction/law-sufficiency coherent profile.

(i) A memory profile is

$$\Pi_{\mathrm{mem}} := (\Pi_{\mathrm{rec/law}}, \mathcal{V}_+, \mathcal{M}_+, \rho_V, \rho_M),$$

where

$$\rho_V : \mathcal{T}_+ \to \mathcal{V}_+, \qquad \rho_M : \mathcal{T}_+ \to \mathcal{M}_+$$

are the support-relative visible-trace and memory-trace maps.

(ii) A loss profile is

$$\Pi_\ell := (\Pi_{\mathrm{rec/law}}, \mathcal{D}, \ell),$$

where $\mathcal{D}$ is a finite decision space and

$$\ell : \mathcal{Y}_+ \times \mathcal{D} \to [0, L_{\max}]$$

is a bounded finite loss.

(iii) A normalized-obstruction profile is

$$\Pi_{\mathrm{norm}} := (\Pi_{\mathrm{rec/law}}, \mathsf{N}),$$

where N specifies the denominator used for each normalized index, for example $H(A)$, $\log_b |\operatorname{supp}(A)|$, or another declared finite normalization.

(iv) A full CEOT IV profile is

$$\Pi_{\mathrm{full}} := (\Pi_{\mathrm{rec/law}}, \mathcal{V}_+, \mathcal{M}_+, \rho_V, \rho_M, \mathcal{D}, \ell, \mathsf{N}).$$

Theorem 7.25 (Typed profile factorization)**.** *The CEOT IV scalar obstruction families factor through the corresponding typed profiles as follows.*

(i) The reconstruction and law-sufficiency family

$$\mathrm{Ob}_{\mathrm{complete}}, \qquad \mathrm{Ob}_{\mathrm{trace}}, \qquad \mathrm{Ob}_{\mathrm{bridge}}, \qquad \mathrm{Ob}^{Y}_{\mathrm{det}}, \qquad \mathrm{Ob}^{Y}_{\mathrm{law}}$$

factors through $\Pi_{\mathrm{rec/law}}$.

(ii) The memory reconstruction and memory-relevance family

$$H(M_{0:T} \mid V_{0:T}), \qquad I(Y; M_{0:T} \mid V_{0:T})$$

factors through Π_{mem}.

(iii) The bounded loss-relative obstruction

$$\mathrm{Ob}^{Y}_{\ell}(O) = R_{\ell}(O) - R_{\ell}(S_{0:T})$$

factors through Π_{ℓ}.

(iv) Normalized obstruction indices factor through Π_{norm}*, not through* $\Pi_{\mathrm{rec/law}}$ *alone unless the normalization convention is fixed externally.*

(v) The full scalar CEOT IV obstruction family factors through Π_{full}.

Proof. Part (i) is Theorem 7.23.

For (ii), the maps ρ_V, ρ_M, together with $P_{\mathcal{T}}$, induce the joint law of $(V_{0:T}, M_{0:T})$. Together with L_+, they induce the joint law of $(Y, V_{0:T}, M_{0:T})$. Hence both $H(M_{0:T} \mid V_{0:T})$ and $I(Y; M_{0:T} \mid V_{0:T})$ are finite functionals of Π_{mem}.

For (iii), $P_{\mathcal{T}}$, ρ_O, and L_+ determine $P_{Y|O}$ and $P_{Y|S_{0:T}}$. The additional decision space $\mathcal{D}$ and loss ℓ determine the Bayes risks

$$R_{\ell}(O) = \sum_{o} P_O(o) \min_{d \in \mathcal{D}} \sum_{y} P_{Y|O=o}(y) \ell(y, d),$$

and

$$R_{\ell}(S_{0:T}) = \sum_{\tau \in \mathcal{T}_+} P_{\mathcal{T}}(\tau) \min_{d \in \mathcal{D}} \sum_{y} L_+(y \mid \tau) \ell(y, d).$$

Therefore $\mathrm{Ob}^{Y}_{\ell}(O)$ is a functional of Π_{ℓ}.

Part (iv) is immediate because a normalized index is not determined until its denominator convention is declared. Part (v) collects the preceding items. □

Remark 7.26 (No profile should be used beyond its typed data). A profile factorization statement is only as strong as the typed data included in the profile. The reconstruction/law profile $\Pi_{\mathrm{rec/law}}$ determines reconstruction entropy and stochastic target-law sufficiency, but it does not by itself determine loss-relative obstruction unless $(\mathcal{D}, \ell)$ is supplied. It also does not determine visible-memory relevance unless the maps ρ_V, ρ_M are supplied. Thus every CEOT IV factorization theorem must name the profile through which it factors.

7.12 Finite categorical content of CEOT IV

The categorical content of CEOT IV is finite and explicit. It consists of two separate structures: a refinement category controlling monotonicity and a coherent typed-profile groupoid controlling relabeling invariance.

Theorem 7.27 (Finite categorical content of CEOT IV)**.** *For a fixed finite CEOT IV trace-target presentation, the categorical content of the theory consists of the following finite structures.*

(i) The deterministic observation-refinement category

$$\mathsf{Obs}_{\mathrm{ref}}(S_{0:T}),$$

whose morphisms are support-relative refinements $c \to c'$ from coarser observations to finer observations.

(ii) For each trace-derived or target-derived obstruction variable A, a covariant obstruction functor

$$\mathrm{Ob}_A : \mathsf{Obs}_{\mathrm{ref}}(S_{0:T}) \longrightarrow (\mathbb{R}_{\geq 0}, \geq), \qquad c \longmapsto H(A \mid O_c).$$

Conditional mutual-information obstruction maps obey the same reverse-order functoriality whenever the target kernel and the relevant trace law are fixed.

(iii) The coherent typed-profile groupoid, whose morphisms are support-preserving relabelings of the finite typed profile data. At the reconstruction/law level this is represented by

$$\mathbf{TraceTargetProf}^{\cong}_{\mathrm{fin,coh}},$$

and the typed extensions record the additional memory maps, losses, and normalization conventions needed for the larger CEOT IV obstruction family.

(iv) A profile-factorization map assigning to each typed coherent profile its declared scalar obstruction family:

$$\Pi_{\mathrm{full}} \longmapsto \mathrm{Ob}_{\mathcal{I}}(\Pi_{\mathrm{full}}) \in \mathbb{R}^{\mathcal{I}}_{\geq 0},$$

where $\mathcal{I}$ indexes the declared obstruction family. This map is invariant under typed profile isomorphism.

(v) The refinement category and the profile groupoid are logically separate. Refinement gives monotonicity; profile isomorphism gives relabeling invariance.

Proof. Part (i) is the deterministic support-relative observation-refinement construction. Part (ii) follows from Proposition 7.6: conditional entropy decreases under refinement, and the codomain order is reversed. Part (iii) is the coherent typed-profile construction. Part (iv) is Theorem 7.25: the relevant finite law, support maps, target kernel, memory maps, decision space, loss, and normalization convention determine the declared scalar obstruction family. Part (v) follows from the non-implication results separating refinement from isomorphism. $\square$

$$\begin{array}{llll} \mathsf{Obs}_{\mathrm{ref}}(S_{0:T}) & \xrightarrow{\mathrm{Ob}_A} & (\mathbb{R}_{\geq 0}, \geq) & \text{monotonicity under refinement} \\ \mathbf{TraceTargetProf}^{\cong}_{\mathrm{fin,coh}} & \xrightarrow{\mathrm{Ob}_{\mathcal{I}}} & \mathbb{R}^{\mathcal{I}}_{\geq 0} & \text{invariance under relabeling.} \end{array}$$

The first row is order-theoretic. The second row is groupoid-invariant. They are not the same structure.

Remark 7.28 (What the word categorical does not mean here). The word categorical in CEOT IV refers to the finite refinement and profile-factorization structures above. It does not assert a universal entropy functor on arbitrary probability spaces, a classification of all finite stochastic processes, a higher-categorical reconstruction theorem, or a generator-sensitive algorithmic complexity invariant. All categorical claims in this paper are support-relative, finite, and explicitly typed.

7.13 Separation between invariance and refinement

The two structures in this section should be kept separate.

First, presentation/profile isomorphism is a relabeling-invariance relation. It asks whether two finite descriptions represent the same reconstruction or sufficiency profile up to bijection of supports. This is groupoid-level structure.

Second, deterministic observation refinement is an order relation inside a fixed trace law. It asks whether one observation contains at least as much information as another:

$$O \preceq O' \quad \iff \quad \sigma(O) \subseteq \sigma(O') \subseteq \sigma(S_{0:T}).$$

This is poset-level structure.

The CEOT IV obstruction maps behave differently with respect to these structures:

under isomorphism, they are invariant,

while

under refinement, they are antitone numerical maps.

This separation prevents two common mistakes. One mistake is to treat refinement as merely a categorical relabeling. It is not; it changes the amount of observed information. The other mistake is to treat isomorphism invariance as a monotonicity principle. It is not; isomorphic presentations have exactly the same obstruction values.

7.14 Non-implications between refinement and isomorphism

The refinement category and the profile groupoid serve different purposes. Refinement is an order relation inside a fixed trace law. Profile isomorphism is a relabeling equivalence between coherent finite profiles. The following table records the forbidden inferences.

False inference	**Reason**	**Correct principle**
$O \preceq O' \Rightarrow O \cong O'$	Refinement may strictly add information.	$\mathrm{Ob}(O) \geq \mathrm{Ob}(O')$.
$O \cong O' \Rightarrow O \preceq O'$	Isomorphism is relabeling, not information inclusion.	$\mathrm{Ob}(O) = \mathrm{Ob}(O')$.
$\Pi_{\mathrm{coh}} \cong \Pi'_{\mathrm{coh}} \Rightarrow$ same generator	The profile forgets transition-kernel generation.	Same obstruction family only.
$P_{O,Y} \cong P_{O',Y'} \Rightarrow \mathrm{Ob}_{\mathrm{trace}}(O) = \mathrm{Ob}_{\mathrm{trace}}(O')$	Observed target law forgets hidden posterior fibers.	Requires coherent profile isomorphism.

Proposition 7.29 (Refinement and isomorphism are independent structures)**.** *Observation refinement and support-relative profile isomorphism do not imply one another.*

(i) There exist observations $O \preceq O'$ such that O and O' are not isomorphic as observation profiles.

(ii) There exist observations with isomorphic marginal laws and equal scalar obstruction values that are not comparable by refinement inside the same trace law.

Proof. For (i), let $S \sim \text{Bernoulli}(1/2)$, let $O = *$, and let $O' = S$. Then $O \preceq O'$, since O is a function of O'. However, $\text{supp}(O)$ has one point while $\text{supp}(O')$ has two points, so the observations cannot be isomorphic by a support bijection. The refinement strictly decreases complete-trace obstruction:

$$H(S \mid O) = \log_b 2, \qquad H(S \mid O') = 0.$$

For (ii), let $S = (X, Z)$, where X, Z are independent fair bits, and let

$$O = X, \qquad O' = Z.$$

Then O and O' have isomorphic marginal laws and

$$H(S \mid O) = H(Z \mid X) = \log_b 2, \qquad H(S \mid O') = H(X \mid Z) = \log_b 2.$$

However, neither O is a function of O' nor O' is a function of O almost surely. Thus they are not comparable in the refinement order, even though their scalar obstruction values agree. □

Remark 7.30 (Monotonicity versus invariance). Refinement is the structure behind monotonicity:

$$O \preceq O' \quad \Longrightarrow \quad \text{Ob}(O) \geq \text{Ob}(O').$$

Profile isomorphism is the structure behind invariance:

$$\Pi_{\text{coh}} \cong \Pi'_{\text{coh}} \quad \Longrightarrow \quad \text{Ob}(\Pi_{\text{coh}}) = \text{Ob}(\Pi'_{\text{coh}}).$$

These statements should never be interchanged. Refinement may strictly change obstruction values. Isomorphism preserves them by relabeling.

7.15 Section summary

This section proved that deterministic refinement of compressed observations weakly decreases every CEOT IV obstruction developed so far. If

$$O \preceq O',$$

then

$$\text{Ob}_{\text{complete}}(O') \leq \text{Ob}_{\text{complete}}(O), \qquad \text{Ob}_{\text{trace}}(O') \leq \text{Ob}_{\text{trace}}(O), \qquad \text{Ob}_{\text{bridge}}(O') \leq \text{Ob}_{\text{bridge}}(O),$$

$$\text{Ob}^Y_{\text{det}}(O') \leq \text{Ob}^Y_{\text{det}}(O), \qquad \text{Ob}^Y_{\text{law}}(O') \leq \text{Ob}^Y_{\text{law}}(O).$$

It also defined finite presentation and profile isomorphisms and proved the groupoid-level factorization

$$\mathbf{AlgCompPres}^{\cong}_{\text{fin}} \longrightarrow \mathbf{TraceTargetProf}^{\cong}_{\text{fin,coh}} \longrightarrow \mathbb{R}^{\text{disc}}_{\geq 0}.$$

The next section applies the stochastic sufficiency criterion to lifted-state systems, where a visible trace $V_{0:T}$ may or may not retain all target-relevant information carried by a hidden memory trace $M_{0:T}$.

8 Lifted-State Target Relevance and Visible-Trace Insufficiency

8.1 Purpose of the section

The previous sections developed reconstruction and sufficiency obstructions for a finite trace $S_{0:T}$ observed through a deterministic compressed observation $O = c(S_{0:T})$. This section specializes that framework to lifted-state algorithms. In such algorithms, the full state contains a visible component and a hidden memory component,

$$S_t = (V_t, M_t),$$

where V_t is the state retained by a compressed or visible algorithmic representation and M_t is an additional memory, signature, mode, label, phase, or latent algorithmic variable.

The main question is not merely whether the memory trace $M_{0:T}$ can be reconstructed from the visible trace $V_{0:T}$. The sharper CEOT IV question is whether $M_{0:T}$ carries information about the task target Y that is not already present in $V_{0:T}$. This distinction is essential. Hidden memory may be unreconstructible and yet irrelevant to the target. Conversely, hidden memory may be precisely the missing information that makes a visible-state algorithm insufficient.

The section therefore separates two quantities:

$$H(M_{0:T} \mid V_{0:T})$$

and

$$I(Y; M_{0:T} \mid V_{0:T}).$$

The first measures memory reconstructibility. The second measures memory relevance to the target under the stochastic sufficiency criterion of Section 6. The first can be positive while the second is zero. Only the second is a certificate of target relevance relative to the visible trace.

8.2 Lifted traces

Definition 8.1 (Lifted finite trace)**.** A lifted finite trace is a finite algorithmic trace

$$S_{0:T} = (S_0, \dots, S_T)$$

such that each state space factors as

$$\mathcal{S}_t = \mathcal{V}_t \times \mathcal{M}_t.$$

For each t, the random state has the coordinate representation

$$S_t = (V_t, M_t),$$

where

$$V_t := \pi_{\mathcal{V}_t}(S_t), \qquad M_t := \pi_{\mathcal{M}_t}(S_t).$$

The visible trace and memory trace are the random variables

$$V_{0:T} := (V_0, \dots, V_T), \qquad M_{0:T} := (M_0, \dots, M_T).$$

Thus the full lifted trace can be written as

$$S_{0:T} = (V_{0:T}, M_{0:T})$$

as a random variable taking values in

$$\mathcal{S}_{0:T} = \prod_{t=0}^{T} \mathcal{S}_t = \prod_{t=0}^{T} (\mathcal{V}_t \times \mathcal{M}_t).$$

The visible trace is a deterministic observation of the full trace:

$$O_V = V_{0:T} = \pi_V(S_{0:T}).$$

Thus the lifted-state setting is a special case of the deterministic observation framework developed earlier, with $O = O_V$.

Example 8.2 (Typical hidden memory variables)**.** The component M_t may represent different kinds of hidden algorithmic state:

(i) a dynamic-programming memory label or unresolved boundary condition;

(ii) a mode variable in a hybrid or switching process;

(iii) a signature, parity, accumulated algebraic label, or path-dependent state;

(iv) a feasibility certificate not stored in the visible state;

(v) a latent class determining the future transition law;

(vi) a retained sufficient statistic for a downstream task.

CEOT IV does not require these interpretations to be ontological. They are finite algorithmic variables inside a specified presentation.

8.3 Memory reconstruction obstruction

The first obstruction is ordinary memory reconstructibility.

Definition 8.3 (Memory reconstruction obstruction)**.** For a lifted trace $S_{0:T} = (V_{0:T}, M_{0:T})$, define

$$\mathrm{Ob}_{\mathrm{mem}}(V) := H(M_{0:T} \mid V_{0:T}).$$

This is a reconstruction obstruction. Its vanishing means that the memory trace is a support-relative function of the visible trace.

Proposition 8.4 (Memory zero criterion)**.** *One has*

$$H(M_{0:T} \mid V_{0:T}) = 0$$

if and only if there exists a map

$$R_M : \mathrm{supp}(V_{0:T}) \to \mathrm{supp}(M_{0:T})$$

such that

$$R_M(V_{0:T}) = M_{0:T} \quad a.s.$$

Proof. This is Theorem 4.12 applied to the pair

$$A = M_{0:T}, \qquad O = V_{0:T}.$$

Since all variables are finite, conditional entropy vanishes exactly when every positive-probability visible fiber contains a single memory value. □

A positive value of $\mathrm{Ob}_{\mathrm{mem}}(V)$ says that the visible trace does not reconstruct the memory trace. It does not yet say that the memory is needed for the target. That distinction is the point of the next definition.

8.4 Conditional target sufficiency for lifted traces

Let the target be generated by the support-restricted kernel

$$L_+ : \mathcal{T}_+ \rightsquigarrow \mathcal{Y}_+,$$

where

$$\mathcal{T}_+ \subseteq \prod_{t=0}^{T} (\mathcal{V}_t \times \mathcal{M}_t)$$

is the positive support of the lifted trace. Equivalently, along the coordinate decomposition of the trace, the target law may be written as

$$P(Y \in \cdot \mid V_{0:T}, M_{0:T}).$$

The visible trace is target-sufficient precisely when conditioning on the hidden memory adds no target-relevant information beyond $V_{0:T}$.

Definition 8.5 (Conditional target sufficiency of the visible trace)**.** The visible trace $V_{0:T}$ is conditionally target-sufficient for Y if

$$Y \perp M_{0:T} \mid V_{0:T}.$$

Equivalently,

$$P(Y \mid V_{0:T}, M_{0:T}) = P(Y \mid V_{0:T})$$

on support.

This definition is the lifted-state specialization of stochastic target sufficiency:

$$I(Y; S_{0:T} \mid O) = 0.$$

Indeed, when $O = V_{0:T}$ and $S_{0:T} = (V_{0:T}, M_{0:T})$,

$$I(Y; S_{0:T} \mid V_{0:T}) = I(Y; (V_{0:T}, M_{0:T}) \mid V_{0:T}) = I(Y; M_{0:T} \mid V_{0:T}).$$

Definition 8.6 (Memory-to-target obstruction)**.** The memory-to-target obstruction is

$$\mathrm{Ob}_{\mathrm{mem}\to Y} := I(Y; M_{0:T} \mid V_{0:T}).$$

Thus CEOT IV distinguishes

$$\mathrm{Ob}_{\mathrm{mem}}(V) = H(M_{0:T} \mid V_{0:T})$$

from

$$\mathrm{Ob}_{\mathrm{mem}\to Y} = I(Y; M_{0:T} \mid V_{0:T}).$$

The former is a memory reconstruction obstruction. The latter is a task-relevance obstruction.

Proposition 8.7 (Memory reconstruction bounds memory relevance)**.** *For every finite lifted trace*

$$S_{0:T} = (V_{0:T}, M_{0:T})$$

and every finite target Y*,*

$$0 \leq \mathrm{Ob}_{\mathrm{mem}\to Y} \leq \mathrm{Ob}_{\mathrm{mem}}(V).$$

Equivalently,

$$0 \leq I(Y; M_{0:T} \mid V_{0:T}) \leq H(M_{0:T} \mid V_{0:T}).$$

Proof. Nonnegativity of conditional mutual information gives

$$I(Y; M_{0:T} \mid V_{0:T}) \geq 0.$$

For the upper bound, use the identity

$$I(Y; M_{0:T} \mid V_{0:T}) = H(M_{0:T} \mid V_{0:T}) - H(M_{0:T} \mid V_{0:T}, Y).$$

The final conditional entropy is nonnegative, so

$$I(Y; M_{0:T} \mid V_{0:T}) \leq H(M_{0:T} \mid V_{0:T}).$$

□

Corollary 8.8 (Reconstructible memory is never target-obstructing)**.** *If*

$$\mathrm{Ob}_{\mathrm{mem}}(V) = 0,$$

then

$$\mathrm{Ob}_{\mathrm{mem}\to Y} = 0$$

for every finite target Y.

Proof. This follows immediately from Proposition 8.7. □

Corollary 8.9 (Target-irrelevant hidden memory)**.** *It is possible that*

$$\mathrm{Ob}_{\mathrm{mem}}(V) > 0$$

but

$$\mathrm{Ob}_{\mathrm{mem}\to Y} = 0.$$

Proof. Let V be constant, let $M \sim \mathrm{Bernoulli}(1/2)$, and let Y be constant or independent of M conditional on V. Then

$$H(M \mid V) = H(M) > 0,$$

whereas

$$I(Y; M \mid V) = 0.$$

By Lemma 2.10, this finite witness is realized by a finite CEOT IV trace-compression presentation.

□

Remark 8.10 (When memory relevance saturates reconstruction). The upper bound

$$I(Y; M_{0:T} \mid V_{0:T}) \leq H(M_{0:T} \mid V_{0:T})$$

is saturated exactly when

$$H(M_{0:T} \mid V_{0:T}, Y) = 0.$$

In that case the target reveals all memory uncertainty left unresolved by the visible trace. In general, only part of the unreconstructible memory may be target-relevant.

8.5 Lifted-state target-sufficiency theorem

Theorem 8.11 (Visible-trace target-sufficiency criterion)**.** *For a finite lifted trace with coordinate decomposition*

$$S_{0:T} = (V_{0:T}, M_{0:T})$$

and support-restricted target kernel

$$L_+ : \mathcal{T}_+ \rightsquigarrow \mathcal{Y}_+,$$

the following are equivalent:

(i) $V_{0:T}$ *is conditionally target-sufficient for* Y*;*

(ii) $Y \perp M_{0:T} \mid V_{0:T}$*;*

(iii) $I(Y; M_{0:T} \mid V_{0:T}) = 0$*;*

(iv) on every positive-probability visible fiber,

$$P(Y \mid V_{0:T}, M_{0:T}) = P(Y \mid V_{0:T}).$$

Consequently,

$$I(Y; M_{0:T} \mid V_{0:T}) > 0$$

is exactly the obstruction to conditional-independence target sufficiency of the visible trace.

Proof. The equivalence between conditional independence and zero conditional mutual information is the finite conditional-independence criterion used in Theorem 6.10. The fiberwise equality is the finite-support form of the same condition. Since $S_{0:T} = (V_{0:T}, M_{0:T})$, stochastic sufficiency of the observation $O = V_{0:T}$ is

$$I(Y; S_{0:T} \mid V_{0:T}) = 0.$$

By the chain rule and the fact that $V_{0:T}$ is already in the conditioning variable,

$$I(Y; S_{0:T} \mid V_{0:T}) = I(Y; (V_{0:T}, M_{0:T}) \mid V_{0:T}) = I(Y; M_{0:T} \mid V_{0:T}).$$

This proves the theorem. ☐

Definition 8.12 (Target-relevant memory quotient over the visible trace)**.** Fix a lifted trace support

$$\mathcal{T}_+ \subseteq \mathcal{V}_+ \times \mathcal{M}_+$$

and a support-restricted target kernel

$$L_+ : \mathcal{T}_+ \rightsquigarrow \mathcal{Y}_+.$$

For each visible trace $v \in \mathcal{V}_+$, define the positive memory fiber

$$\mathcal{M}_+(v) := \{m : (v, m) \in \mathcal{T}_+\}.$$

On $\mathcal{M}_+(v)$ define

$$m \sim_{Y|v} m' \quad \Longleftrightarrow \quad L_+(\cdot \mid v, m) = L_+(\cdot \mid v, m').$$

The target-relevant memory quotient over v is

$$q_{Y|v} : \mathcal{M}_+(v) \to \mathcal{M}_+(v)/{\sim_{Y|v}}.$$

The corresponding visible-relative target quotient is

$$Q_Y := (V_{0:T}, q_{Y|V_{0:T}}(M_{0:T})).$$

It retains exactly the visible trace together with the memory information that changes the target law on a positive visible fiber.

Definition 8.13 (Minimal target-sufficient refinement over the visible trace)**.** Let

$$S_{0:T} = (V_{0:T}, M_{0:T})$$

be a finite lifted trace with positive trace support $\mathcal{T}_+$ and target kernel L_+. On $\mathcal{T}_+$ define

$$\tau \sim_{V,Y} \tau'$$

if and only if

$$V_{0:T}(\tau) = V_{0:T}(\tau')$$

and

$$L_+(\cdot \mid \tau) = L_+(\cdot \mid \tau').$$

Let

$$q_{V,Y} : \mathcal{T}_+ \to \mathcal{T}_+/{\sim_{V,Y}}$$

be the corresponding quotient and define

$$Q_{V,Y} := q_{V,Y}(S_{0:T}).$$

We call $Q_{V,Y}$ the minimal target-sufficient refinement over the visible trace.

Theorem 8.14 (Visible insufficiency and minimal required refinement)**.** *For a finite lifted trace*

$$S_{0:T} = (V_{0:T}, M_{0:T}),$$

the following are equivalent:

(i)

$$I(Y; M_{0:T} \mid V_{0:T}) = 0.$$

(ii) The trace-conditioned target law is constant on every positive visible fiber:

$$V_{0:T}(\tau) = V_{0:T}(\tau') \quad \Longrightarrow \quad L_+(\cdot \mid \tau) = L_+(\cdot \mid \tau')$$

for all $\tau, \tau' \in \mathcal{T}_+$.

(iii) The minimal target-sufficient refinement $Q_{V,Y}$ factors through the visible trace $V_{0:T}$.

Consequently,

$$I(Y; M_{0:T} \mid V_{0:T}) > 0$$

means that the visible trace alone is not law-sufficient for Y. It does not by itself imply that the entire memory trace $M_{0:T}$ is minimal or uniquely necessary.

Proof. The equivalence between (i) and (ii) is the finite stochastic law-sufficiency criterion applied to the observation $V_{0:T}$ and the full lifted trace $(V_{0:T}, M_{0:T})$. Condition (ii) says exactly that inside each positive visible fiber, the target law is constant. This is equivalent to saying that the quotient $q_{V,Y}$ is already determined by $V_{0:T}$, i.e. $Q_{V,Y}$ factors through the visible trace.

The final statement is the contrapositive. If the conditional mutual information is positive, then the visible trace fails the target-sufficiency criterion. Since $Q_{V,Y}$ is the coarsest sufficient refinement over the visible trace, some additional feature beyond $V_{0:T}$ is required, but the full memory trace need not be minimal. ☐

Theorem 8.15 (Full memory is not the asserted necessity)**.** *If*

$$I(Y; M_{0:T} \mid V_{0:T}) > 0,$$

then the visible trace $V_{0:T}$ alone is not conditionally target-sufficient for Y. However, CEOT IV does not imply that the full memory trace $M_{0:T}$ is minimal, unique, or necessary as retained data. The visible-relative quotient

$$Q_Y = (V_{0:T}, q_{Y|V_{0:T}}(M_{0:T}))$$

is target-law sufficient:

$$I(Y; S_{0:T} \mid Q_Y) = 0.$$

Moreover, it is coarsest among deterministic visible-relative quotients that preserve the support-restricted target kernel on positive lifted-trace support.

Proof. By construction, $q_{Y|v}$ identifies precisely those memory values inside a fixed positive visible fiber that induce the same conditional target law. Therefore $L_+(\cdot \mid v, m)$ factors through Q_Y, so conditioning on Q_Y preserves the same target law as conditioning on the full lifted trace. This gives

$$I(Y; S_{0:T} \mid Q_Y) = 0.$$

Conversely, any deterministic quotient over the visible trace that preserves the target kernel must separate pairs (v, m) and (v, m') with different laws $L_+(\cdot \mid v, m)$ and $L_+(\cdot \mid v, m')$. Hence Q_Y is the coarsest such support-relative quotient. ☐

Remark 8.16 (Positive memory relevance is not full-memory minimality)*.* The inequality

$$I(Y; M_{0:T} \mid V_{0:T}) > 0$$

means exactly that the visible trace alone has lost target-relevant information. It does not say that every coordinate of $M_{0:T}$ is target-relevant, and it does not say that $M_{0:T}$ is the minimal sufficient augmentation. Minimality, when desired, is a visible-fiber quotient statement such as Theorem 8.15, not a statement about retaining the entire hidden memory trace.

Corollary 8.17 (Visible-trace insufficiency)**.** *If*

$$I(Y; M_{0:T} \mid V_{0:T}) > 0,$$

then the visible trace $V_{0:T}$ is not conditionally target-sufficient for Y.

Remark 8.18 (Why this is a target theorem, not merely a memory theorem)*.* The condition

$$H(M_{0:T} \mid V_{0:T}) > 0$$

only says that the memory trace is not reconstructible. It does not imply target failure. The target-relevance certificate is instead

$$I(Y; M_{0:T} \mid V_{0:T}) > 0.$$

Thus CEOT IV does not justify lifted states merely because they are hidden. It certifies visible-trace insufficiency when hidden coordinates carry target-relevant information not present in the visible trace.

8.6 Same-visible-fiber criterion

The preceding theorem can be written as an explicit finite support criterion.

Proposition 8.19 (Same-visible-fiber criterion)**.** *For a finite lifted trace,*

$$I(Y; M_{0:T} \mid V_{0:T}) > 0$$

if and only if there exists a visible trace $v_{0:T}$ with

$$P(V_{0:T} = v_{0:T}) > 0$$

and two memory traces $m_{0:T}$ and $m'_{0:T}$ with

$$P(M_{0:T} = m_{0:T} \mid V_{0:T} = v_{0:T}) > 0,$$

$$P(M_{0:T} = m'_{0:T} \mid V_{0:T} = v_{0:T}) > 0,$$

such that

$$P(Y \mid V_{0:T} = v_{0:T}, M_{0:T} = m_{0:T}) \neq P(Y \mid V_{0:T} = v_{0:T}, M_{0:T} = m'_{0:T}).$$

Proof. If the displayed pair of memory traces exists, then the conditional target law varies inside a positive-probability visible fiber. Hence Y is not conditionally independent of $M_{0:T}$ given $V_{0:T}$, so

$$I(Y; M_{0:T} \mid V_{0:T}) > 0.$$

Conversely, if the conditional mutual information is positive, then conditional independence fails on some positive-probability visible fiber. Therefore the conditional law of Y cannot be constant over all positive-probability memory traces inside that fiber. Hence two such memory traces with different target laws exist. □

This criterion is often the most operational form of visible-trace insufficiency. To prove that a visible trace is insufficient, it is enough to exhibit two hidden memory histories compatible with the same visible history but inducing different target laws.

8.7 Deterministic targets inside lifted traces

Suppose the target is deterministic on positive lifted-trace support:

$$Y = \varphi_+(V_{0:T}, M_{0:T}).$$

Then Theorem 6.19 gives

$$I(Y; M_{0:T} \mid V_{0:T}) = H(Y \mid V_{0:T}).$$

Thus the lifted-state sufficiency criterion reduces to deterministic target recovery from the visible trace.

Corollary 8.20 (Deterministic lifted target criterion)**.** *If $Y = \varphi_+(V_{0:T}, M_{0:T})$ is deterministic on positive lifted-trace support, then the following are equivalent:*

(i) *$V_{0:T}$ is conditionally target-sufficient for Y;*

(ii) *$I(Y; M_{0:T} \mid V_{0:T}) = 0$;*

(iii) *$H(Y \mid V_{0:T}) = 0$;*

(iv) *there exists a decoder d such that*

$$d(V_{0:T}) = Y \quad a.s.$$

Proof. Since Y is a support-relative function of $(V_{0:T}, M_{0:T})$, conditional on $V_{0:T}$ one has

$$I(Y; M_{0:T} \mid V_{0:T}) = H(Y \mid V_{0:T}).$$

The zero criterion for $H(Y \mid V_{0:T})$ is Theorem 5.5 applied to the observation $V_{0:T}$. □

This corollary is the precise CEOT IV form of a common algorithmic statement: if two hidden memory states compatible with the same visible trace force different deterministic outputs, then the visible trace alone is not enough to recover the output.

8.8 Two-axis hidden-memory classification

Memory non-reconstructibility and target relevance are different axes. The following theorem replaces the informal dichotomy by a support-relative classification.

Theorem 8.21 (Two-axis hidden-memory classification)**.** *For a finite lifted trace*

$$S_{0:T} = (V_{0:T}, M_{0:T})$$

and a finite target Y, define

$$\mathrm{Ob}_{\mathrm{mem}}(V) := H(M_{0:T} \mid V_{0:T})$$

and

$$\mathrm{Ob}_{\mathrm{mem}\to Y} := I(Y; M_{0:T} \mid V_{0:T}).$$

The exact-zero possibilities are exhausted by the following cases:

(i) ***Visible-reconstructible memory.***

$$H(M_{0:T} \mid V_{0:T}) = 0.$$

Then

$$I(Y; M_{0:T} \mid V_{0:T}) = 0$$

for every finite target Y.

(ii) ***Hidden but target-irrelevant memory.***

$$H(M_{0:T} \mid V_{0:T}) > 0, \qquad I(Y; M_{0:T} \mid V_{0:T}) = 0.$$

The memory is not reconstructible from the visible trace, but it carries no additional target-law information beyond the visible trace.

(iii) ***Hidden and target-relevant memory.***

$$H(M_{0:T} \mid V_{0:T}) > 0, \qquad I(Y; M_{0:T} \mid V_{0:T}) > 0.$$

The visible trace is not target-sufficient for Y.

The remaining formal case

$$H(M_{0:T} \mid V_{0:T}) = 0, \qquad I(Y; M_{0:T} \mid V_{0:T}) > 0$$

is impossible. Moreover, the hidden-but-target-irrelevant and hidden-and-target-relevant cases both occur in finite lifted systems.

Proof. The inequality

$$0 \le I(Y; M_{0:T} \mid V_{0:T}) \le H(M_{0:T} \mid V_{0:T})$$

shows that visible-reconstructible memory cannot carry additional target information beyond the visible trace. If the memory is genuinely hidden, i.e.

$$H(M_{0:T} \mid V_{0:T}) > 0,$$

then target relevance is decided by whether

$$I(Y; M_{0:T} \mid V_{0:T})$$

is zero or positive. Positivity is exactly the failure of visible-trace target sufficiency by Theorem 8.11.

For realizability, take V constant and $M \sim \text{Bernoulli}(1/2)$. If Y is independent of M given V, then

$$H(M \mid V) = \log_b 2, \qquad I(Y; M \mid V) = 0.$$

If instead $Y = M$, then

$$H(M \mid V) = \log_b 2, \qquad I(Y; M \mid V) = \log_b 2.$$

Embedding either one-time construction into a trace by taking all other coordinates trivial gives finite lifted trace examples. □

Corollary 8.22 (Target-relevant memory is necessarily hidden)**.** *If*

$$I(Y; M_{0:T} \mid V_{0:T}) > 0,$$

then

$$H(M_{0:T} \mid V_{0:T}) > 0.$$

Proof. This follows immediately from

$$I(Y; M_{0:T} \mid V_{0:T}) \le H(M_{0:T} \mid V_{0:T}).$$

□

This classification prevents an overclaim. CEOT IV does not say that all hidden memory must be retained. It says that a visible-only representation fails when the target-sufficiency obstruction detects target-relevant information beyond the visible trace; a smaller target-sufficient refinement than the full memory may still exist.

8.9 Local versus global visible-trace insufficiency

One may also consider local quantities

$$I(Y; M_t \mid V_t).$$

These can be useful diagnostic indicators, but they are not the main CEOT IV criterion for trace-level sufficiency.

The global criterion is

$$I(Y; M_{0:T} \mid V_{0:T}).$$

The difference matters. A local visible state V_t may fail to screen off Y from M_t, while the full visible trace $V_{0:T}$ contains later or earlier information that restores sufficiency. Conversely, every individual local quantity may fail to expose the relevant dependence if the target depends on a joint pattern of memory values across time. Therefore the flagship lifted-state theorem is global in the trace:

$$V_{0:T} \text{ is conditionally target-sufficient for } Y \iff I(Y; M_{0:T} \mid V_{0:T}) = 0.$$

Theorem 8.23 (Local memory indicators do not characterize global sufficiency)**.** *The local indicators*

$$I(Y; M_t \mid V_t)$$

do not characterize the global visible-trace sufficiency criterion

$$I(Y; M_{0:T} \mid V_{0:T}).$$

In finite lifted traces, both non-implications can occur:

$$I(Y; M_t \mid V_t) > 0 \textit{ for some } t \not\Longrightarrow I(Y; M_{0:T} \mid V_{0:T}) > 0,$$

and

$$I(Y; M_t \mid V_t) = 0 \textit{ for all } t \not\Longrightarrow I(Y; M_{0:T} \mid V_{0:T}) = 0.$$

Proof. First let $T = 1$. Let $B \sim \text{Bernoulli}(1/2)$ and set

$$Y = B, \qquad V_0 = *, \qquad M_0 = B, \qquad V_1 = B, \qquad M_1 = *.$$

Then

$$I(Y; M_0 \mid V_0) = I(B; B \mid *) = \log_b 2 > 0.$$

However, the full visible trace already contains B, since $V_{0:1} = (*, B)$. Therefore

$$I(Y; M_{0:1} \mid V_{0:1}) = I(B; (B, *) \mid *, B) = 0.$$

A positive local indicator can therefore disappear at the global trace level.

For the opposite non-implication, let B_0, B_1 be independent fair bits and set

$$V_0 = V_1 = *, \qquad M_0 = B_0, \qquad M_1 = B_1, \qquad Y = B_0 \oplus B_1.$$

For each time t, Y is independent of the individual bit B_t, since XOR with the other independent fair bit randomizes it. Hence

$$I(Y; M_0 \mid V_0) = 0, \qquad I(Y; M_1 \mid V_1) = 0.$$

But jointly Y is determined by (M_0, M_1), and therefore

$$I(Y; M_{0:1} \mid V_{0:1}) = I(Y; (B_0, B_1) \mid *, *) = H(Y) = \log_b 2 > 0.$$

Thus all local indicators can vanish while the global visible trace is target-insufficient. ☐

Principle 8.24 (Local indicators are diagnostics, not certificates). The quantities

$$I(Y; M_t \mid V_t)$$

may help locate where target-relevant memory appears, but they are not CEOT IV sufficiency certificates. The certificate is global:

$$I(Y; M_{0:T} \mid V_{0:T}) = 0.$$

Local indicators can both over-detect and under-detect target-relevant memory because visible information and memory synergy are temporal.

Principle 8.25 (Global before local). Local lifted-state obstructions are useful diagnostics. The CEOT IV sufficiency certificate is the global trace-level obstruction

$$I(Y; M_{0:T} \mid V_{0:T}).$$

8.10 Relation to state augmentation

Lifted-state algorithms introduce additional state variables. In ordinary algorithm design, such variables are often justified informally: they carry memory, encode a mode, preserve a constraint, or retain a signature. CEOT IV gives a finite information-theoretic certificate for when the visible trace alone is insufficient for the declared target.

If

$$I(Y; M_{0:T} \mid V_{0:T}) = 0,$$

then the visible trace already contains all target-law information carried by the full lifted trace. The memory may still help computational efficiency, implementation, or proof organization, but CEOT IV does not certify it as target-relevant relative to the visible trace.

If

$$I(Y; M_{0:T} \mid V_{0:T}) > 0,$$

then the visible trace loses target-relevant information. Any representation that keeps only $V_{0:T}$ fails the stochastic sufficiency criterion for Y. Equivalently, some nontrivial target-sufficient refinement of $V_{0:T}$ is required. The full memory trace $M_{0:T}$ is one possible source of such a refinement, but CEOT IV does not claim that it is minimal or unique.

8.11 Section summary

This section specialized CEOT IV to lifted traces

$$S_t = (V_t, M_t).$$

It introduced the memory reconstruction obstruction

$$\mathrm{Ob}_{\mathrm{mem}}(V) = H(M_{0:T} \mid V_{0:T})$$

and the memory-to-target obstruction

$$\mathrm{Ob}_{\mathrm{mem}\to Y} = I(Y; M_{0:T} \mid V_{0:T}).$$

The central theorem was the lifted-state sufficiency criterion:

$$V_{0:T} \text{ is conditionally target-sufficient for } Y \iff I(Y; M_{0:T} \mid V_{0:T}) = 0.$$

Equivalently,

$$I(Y; M_{0:T} \mid V_{0:T}) > 0$$

is the obstruction to conditional-independence target sufficiency of the visible trace. It also proved the hierarchy

$$0 \le \mathrm{Ob}_{\mathrm{mem}\to Y} \le \mathrm{Ob}_{\mathrm{mem}}(V),$$

a same-visible-fiber criterion, and a hidden-memory dichotomy, showing that memory can be hidden but irrelevant or hidden and target-relevant relative to the visible trace.

The next section develops separation and non-completeness principles. There the paper shows that isomorphic observed target laws can coexist with different hidden trace obstructions, and that Markov lumpability is logically distinct from CEOT IV trace reconstruction and task sufficiency.

9 Separation and Non-Completeness Principles

9.1 Purpose of the section

The preceding sections established reconstruction criteria, target criteria, monotonicity under deterministic observation refinement, profile-level invariance, and visible-trace insufficiency. This section records the main separation principles that prevent those criteria from collapsing into one another. The goal is to make explicit which implications CEOT IV proves and which implications it deliberately rejects.

There are three separations. The first two are proved by minimal finite witnesses and then illustrated by algorithmic trace examples. The minimal witnesses are intentionally small: they isolate the logical implication being refuted. The dynamic-programming and signature-path witnesses below show that the same separations occur in algorithmic trace-compression settings. By Lemma 2.10, all finite random-variable witnesses in this section can be realized as genuine CEOT IV trace-compression presentations with explicit kernels.

First, hidden trace loss is not the same as target loss. It may happen that

$$H(S_{\mathrm{int}} \mid O) > 0$$

while the deterministic target is perfectly recoverable:

$$H(Y \mid O) = 0.$$

For stochastic targets, it may similarly happen that

$$H(S_{\mathrm{int}} \mid O) > 0$$

while

$$I(Y; S_{0:T} \mid O) = 0.$$

Thus hidden computational trace loss is not, by itself, an algorithmic failure certificate.

Second, observed compressed target behavior does not determine hidden trace obstruction. Two finite trace-compression presentations may have isomorphic joint laws of the compressed observation and the target,

$$P^{\mathfrak{A}}_{O,Y} \cong_{\mathrm{law}} P^{\mathfrak{A}'}_{O',Y'},$$

but different hidden trace obstructions. This is the CEOT IV analogue of the CEOT III principle that observed behavior does not determine hidden completion obstruction.

Third, Markov lumpability is logically distinct from CEOT IV trace reconstruction and task sufficiency. Lumpability concerns preservation of visible transition laws. CEOT IV concerns reconstruction of hidden traces and sufficiency of compressed observations for a specified target. Neither notion subsumes the other.

9.2 Trace loss without deterministic target loss

The first separation is the simplest one: a compressed observation may forget hidden trace information that is irrelevant to the deterministic target.

Theorem 9.1 (Hidden trace loss without deterministic target loss)**.** *There exists a finite deterministic trace-compression presentation such that*

$$\mathrm{Ob}_{\mathrm{trace}}(O) = H(S_{\mathrm{int}} \mid O) > 0,$$

but

$$\mathrm{Ob}^{Y}_{\mathrm{det}}(O) = H(Y \mid O) = 0.$$

Proof. Take a one-interior-step system with time horizon $T = 2$ and trivial endpoints

$$S_0 = *, \qquad S_2 = *.$$

Let

$$Y \sim \mathrm{Bernoulli}(1/2), \qquad Z \sim \mathrm{Bernoulli}(1/2), \qquad Z \perp Y,$$

and set

$$S_1 = (Y, Z).$$

Let the compressed observation be

$$O = Y,$$

and let the deterministic target also be Y. Then

$$H(Y \mid O) = H(Y \mid Y) = 0.$$

However,

$$H(S_1 \mid O) = H(Y, Z \mid Y) = H(Z \mid Y) = \log_b 2 > 0.$$

Since $S_{\mathrm{int}} = S_1$, this gives

$$\mathrm{Ob}_{\mathrm{trace}}(O) > 0 \qquad \text{and} \qquad \mathrm{Ob}^{Y}_{\mathrm{det}}(O) = 0.$$

□

This theorem is the most elementary warning against overinterpreting trace obstruction. A positive trace obstruction says that the hidden internal history is not reconstructible from the compressed observation. It does not say that the chosen deterministic target has been lost.

Example 9.2 (Dynamic-programming tie witness)**.** Consider a finite dynamic-programming trace with two internal optimal policies

$$\pi_0, \qquad \pi_1$$

that both achieve the same final value $V^* = 0$. Let

$$S_0 = *, \qquad S_1 = \Pi, \qquad S_2 = V^*,$$

where $\Pi \sim \mathrm{Bernoulli}(1/2)$ selects one of the tied policies. Let the compressed observation retain only the final value,

$$O = V^*,$$

and let the deterministic target also be the final value,

$$Y = V^*.$$

Then

$$H(Y \mid O) = 0,$$

because the value target is observed exactly. But

$$H(S_{\mathrm{int}} \mid O) = H(\Pi \mid V^*) = \log_b 2.$$

Thus CEOT IV separates recovery of the final value from recovery of the internal policy path: the hidden trace records which optimal policy was used, while the task target records only the value.

9.3 Trace loss without stochastic target loss

The same separation exists for stochastic targets. The correct stochastic obstruction is not hidden trace entropy, but conditional mutual information between the target and the full trace after conditioning on the compressed observation.

Theorem 9.3 (Hidden trace loss without stochastic target loss)**.** *There exists a finite deterministic trace-compression presentation with a stochastic target such that*

$$H(S_{\mathrm{int}} \mid O) > 0,$$

but

$$I(Y; S_{0:T} \mid O) = 0.$$

Proof. Use the same trace as in Theorem 9.1:

$$S_0 = *, \qquad S_1 = (A, Z), \qquad S_2 = *,$$

where A and Z are independent Bernoulli variables. Let

$$O = A.$$

Let Y be generated from A by any non-degenerate finite kernel

$$L_A : A \rightsquigarrow Y,$$

and assume that, conditional on A, the target Y is independent of Z. Since $O = A$, the conditional law of Y given the full trace depends only on O:

$$P(Y \mid S_{0:T}) = P(Y \mid A) = P(Y \mid O).$$

Thus

$$Y \perp S_{0:T} \mid O,$$

and hence

$$I(Y; S_{0:T} \mid O) = 0.$$

However,

$$H(S_1 \mid O) = H(A, Z \mid A) = H(Z \mid A) = \log_b 2 > 0.$$

Therefore the hidden trace is not reconstructible, while the compressed observation is stochastically target-sufficient. $\square$

The theorem shows why Section 6 uses conditional mutual information rather than trace entropy as the stochastic target obstruction. A compressed observation can lose irrelevant trace coordinates while preserving the full target law.

Example 9.4 (Signature-dependent path witness)**.** Let a finite path algorithm choose an internal signature

$$\Sigma \in \{-1, +1\}$$

and an irrelevant tie variable

$$Z \in \{0, 1\},$$

independently and uniformly. Let the trace have trivial endpoints and internal state

$$S_0 = *, \qquad S_1 = (\Sigma, Z), \qquad S_2 = *.$$

Let the compressed observation retain only the signature,

$$O = \Sigma.$$

Let the stochastic target be generated by a kernel depending only on the signature:

$$P(Y = 1 \mid S_1 = (\Sigma, Z)) = p_\Sigma,$$

where $p_{+1}, p_{-1} \in (0, 1)$. Then the target law is already determined by O, so

$$Y \perp S_{0:2} \mid O$$

and therefore

$$I(Y; S_{0:2} \mid O) = 0.$$

However,

$$H(S_{\text{int}} \mid O) = H((\Sigma, Z) \mid \Sigma) = H(Z) = \log_b 2.$$

Thus the compressed observation preserves the full stochastic target law while forgetting a genuine hidden internal coordinate.

9.4 Trace-level non-completeness of compressed target behavior

We now prove the main non-completeness principle of CEOT IV. The observed joint law of the compressed observation and the target does not determine the hidden trace obstruction.

Definition 9.5 (Observed target law isomorphism)**.** Let (O, Y) and (O', Y') be finite observed-target pairs. Their observed target laws are isomorphic, written

$$P_{O,Y} \cong_{\text{law}} P_{O',Y'},$$

if there exist bijections

$$\beta : \text{supp}(O) \to \text{supp}(O'), \qquad \theta : \text{supp}(Y) \to \text{supp}(Y')$$

such that

$$P_{O',Y'}(\beta(o), \theta(y)) = P_{O,Y}(o, y)$$

for all $o \in \text{supp}(O)$ and $y \in \text{supp}(Y)$. Equivalently,

$$P_{O',Y'} = (\beta, \theta)_* P_{O,Y}.$$

If the finite supports have already been canonically identified, literal equality is a harmless abbreviation for this law isomorphism.

Remark 9.6 (Observed law isomorphism forgets hidden fibers). The relation

$$P_{O,Y} \cong_{\text{law}} P_{O',Y'}$$

identifies only the observable target interface. It does not include a bijection between hidden trace supports, does not identify posterior fibers such as

$$P_{S_{\text{int}}|O=o} \qquad \text{and} \qquad P_{S'_{\text{int}}|O'=\beta(o)},$$

and does not assert an isomorphism of coherent CEOT IV profiles. It is therefore strictly weaker than presentation or profile isomorphism. Hidden-trace obstruction is not expected to be invariant under this weaker relation.

Theorem 9.7 (Observed target law is not complete for hidden trace obstruction)**.** *There exist finite deterministic trace-compression presentations*

$$\mathfrak{A} \qquad \textit{and} \qquad \mathfrak{A}'$$

with observed target law isomorphism

$$P^{\mathfrak{A}}_{O,Y} \cong_{\text{law}} P^{\mathfrak{A}'}_{O',Y'}$$

in the sense of Definition 9.5, but with different hidden-trace posterior profiles and different hidden trace obstructions:

$$\mathsf{Post}^{\mathfrak{A}}_{S_{\text{int}}|O} \ncong \mathsf{Post}^{\mathfrak{A}'}_{S'_{\text{int}}|O'},$$

and

$$\mathrm{Ob}^{\mathfrak{A}}_{\text{trace}}(O) \neq \mathrm{Ob}^{\mathfrak{A}'}_{\text{trace}}(O').$$

Equivalently, the observed target law, even up to support relabeling, does not determine the hidden trace obstruction.

Proof. Let both systems have time horizon $T = 2$ and trivial endpoints.

For the first presentation $\mathfrak{A}$, let

$$Y \sim \text{Bernoulli}(1/2), \qquad Z \sim \text{Bernoulli}(1/2), \qquad Z \perp Y,$$

and set

$$S_0 = *, \qquad S_1 = (Y, Z), \qquad S_2 = *.$$

Let the compressed observation be

$$O = Y,$$

and let the target be Y. Then

$$P^{\mathfrak{A}}_{O,Y}$$

is the law of (Y, Y), where Y is Bernoulli$(1/2)$. The hidden trace obstruction is

$$\mathrm{Ob}^{\mathfrak{A}}_{\text{trace}}(O) = H(S_1 \mid O) = H(Y, Z \mid Y) = \log_b 2.$$

For the second presentation $\mathfrak{A}'$, let

$$S'_0 = *, \qquad S'_1 = Y', \qquad S'_2 = *,$$

with an isomorphic Bernoulli(1/2) target variable Y', and set

$$O' = Y'.$$

Then

$$P^{\mathfrak{A}'}_{O',Y'}$$

is again the law of a diagonal Bernoulli pair. After the evident identification of the two Bernoulli target supports, the two observed target laws are literally equal. Equivalently,

$$P^{\mathfrak{A}}_{O,Y} \cong_{\mathrm{law}} P^{\mathfrak{A}'}_{O',Y'}.$$

But

$$\mathrm{Ob}^{\mathfrak{A}'}_{\mathrm{trace}}(O') = H(S'_1 \mid O') = H(Y' \mid Y') = 0.$$

Thus the two presentations have isomorphic observed target laws but different hidden trace obstructions. □

The theorem is deliberately trace-level. It does not merely say that complete trace obstruction can vary behind the same observed behavior. It says that the roadmap obstruction

$$H(S_{\mathrm{int}} \mid O)$$

itself is not determined by the law of the observed compressed target pair (O, Y).

Corollary 9.8 (Observed target law is not a complete CEOT IV invariant)**.** *The assignment*

$$\mathfrak{A} \longmapsto [P^{\mathfrak{A}}_{O,Y}]_{\cong_{\mathrm{law}}}$$

does not determine the hidden-trace posterior profile

$$\mathsf{Post}_{S_{\mathrm{int}}|O}$$

and does not determine the scalar hidden-trace obstruction

$$H(S_{\mathrm{int}} \mid O).$$

Proof. The two presentations in Theorem 9.7 have isomorphic observed target laws but different values of $H(S_{\mathrm{int}} \mid O)$. Since the scalar entropy is computed from the posterior profile, the profile cannot be determined by the observed target law, even up to support relabeling, either. □

Corollary 9.9 (No contradiction with profile invariance)**.** *The non-completeness theorem does not contradict profile invariance. Profile isomorphism preserves the hidden posterior profile and therefore preserves hidden trace obstruction. Observed target law isomorphism is strictly weaker: it preserves only the task interface $P_{O,Y}$, not the hidden posterior fibers.*

Proof. Profile invariance assumes an isomorphism of the typed CEOT profile, including the trace support, observation map, target kernel, and induced posterior fibers. Theorem 9.7 assumes only an isomorphism of (O, Y). The latter deliberately forgets the data needed to compute $H(S_{\mathrm{int}} \mid O)$. □

This is the profile-first reason for not reducing CEOT IV to a theory of observed target laws. The target law records what the compressed algorithm appears to do at the task interface. The hidden trace profile records how much internal computational history remains unreconstructible behind that interface.

9.5 Complete-trace, hidden-trace, and target-obstruction non-equivalence

The examples above also clarify the relation among the main scalar obstructions. The hierarchy

$$\mathrm{Ob}_{\mathrm{bridge}}(O) \leq \mathrm{Ob}_{\mathrm{trace}}(O) \leq \mathrm{Ob}_{\mathrm{complete}}(O)$$

compares trace-reconstruction tasks. It does not compare trace reconstruction with target sufficiency.

For deterministic targets, Section 5 proved

$$\mathrm{Ob}_{\mathrm{complete}}(O) = 0 \Longrightarrow \mathrm{Ob}^{Y}_{\mathrm{det}}(O) = 0.$$

The converse fails by Theorem 9.1. Likewise, for stochastic targets,

$$\mathrm{Ob}^{Y}_{\mathrm{law}}(O) = I(Y; S_{0:T} \mid O)$$

can vanish while

$$\mathrm{Ob}_{\mathrm{trace}}(O) > 0.$$

Thus CEOT IV has two axes:

reconstructing hidden computational history

and

preserving target-relevant information.

The first axis is measured by complete, trace, and bridge obstructions. The second is measured by deterministic target entropy or stochastic target conditional mutual information.

Example 9.10 (Action-value separation witness)**.** Let a finite optimization routine have two hidden internal traces

$$S_{\mathrm{int}} \in \{s_0, s_1\}$$

with equal probability. Suppose both traces yield the same value but different selected actions:

$$V(s_0) = V(s_1) = 0, \qquad A(s_0) \neq A(s_1).$$

Let the compressed observation retain only the value,

$$O = V.$$

Then the value target is perfectly recovered:

$$H(V \mid O) = 0.$$

But the action target is not recovered:

$$H(A \mid O) > 0.$$

The same compressed observation is therefore complete for the value target and obstructed for the action target. Sufficiency is a property of the declared target, not of the compression map alone.

9.6 Markov lumpability and CEOT IV are different questions

We next separate CEOT IV from classical Markov lumpability. Lumpability is a transition-law condition. CEOT IV trace reconstruction and target sufficiency are reconstruction and information-sufficiency conditions.

Definition 9.11 (Finite Markov lumpability)**.** Let $K : S \rightsquigarrow S$ be a transition kernel on a finite state space, and let

$$q : S \to \bar{S}$$

be a finite compression. The compression q is *Markov-lumpable* for K, in the standard finite Markov-chain sense [15], if, for every pair of states $s, s' \in S$ with

$$q(s) = q(s'),$$

and every visible state $\bar{r} \in \bar{S}$, one has

$$\sum_{r \in q^{-1}(\bar{r})} K(r \mid s) = \sum_{r \in q^{-1}(\bar{r})} K(r \mid s').$$

Equivalently, the one-step law of $q(S_{t+1})$ given $S_t = s$ depends only on $q(s)$, not on the representative s inside its compression fiber.

Definition 9.12 (Support-relative one-step lumpability)**.** Let $K : S \rightsquigarrow S$ be a finite Markov kernel, let $q : S \to \bar{S}$ be a finite compression, and let μ_t be the law of S_t. The compression q is *support-relative lumpable at time t* if, for every pair $s, s' \in \operatorname{supp}(\mu_t)$ with

$$q(s) = q(s'),$$

and every $\bar{r} \in q(S)$, one has

$$\sum_{r \in q^{-1}(\bar{r})} K(r \mid s) = \sum_{r \in q^{-1}(\bar{r})} K(r \mid s').$$

Equivalently, the one-step law of $q(S_{t+1})$ given $S_t = s$ is constant on every positive current-state fiber of q.

Theorem 9.13 (Support-relative lumpability equals one-step CEOT sufficiency)**.** *Let $K : S \rightsquigarrow S$ be a finite Markov kernel, let $q : S \to \bar{S}$ be a finite compression, let μ_t be the law of S_t, and set*

$$\bar{S}_t = q(S_t), \qquad Y = \bar{S}_{t+1} = q(S_{t+1}).$$

Then

$$I(\bar{S}_{t+1}; S_t \mid \bar{S}_t) = 0$$

if and only if q is support-relative lumpable at time t.

Proof. For $s \in \operatorname{supp}(\mu_t)$ and $\bar{r} \in q(S)$,

$$P(\bar{S}_{t+1} = \bar{r} \mid S_t = s) = \sum_{r \in q^{-1}(\bar{r})} K(r \mid s).$$

By the finite conditional-independence criterion,

$$I(\bar{S}_{t+1}; S_t \mid \bar{S}_t) = 0$$

if and only if the conditional law of $\bar{S}_{t+1}$ given $S_t = s$ depends only on $\bar{S}_t = q(s)$ for positive-support states s. This is exactly support-relative lumpability at time t. □

Thus support-relative one-step lumpability is exactly CEOT IV law-sufficiency for the next visible-state target. Classical finite Markov lumpability implies this support-relative criterion whenever it holds on the positive current-state support. Conversely, CEOT IV one-step law-sufficiency implies classical all-states lumpability only under an additional full-support or all-states condition.

Remark 9.14 (What lumpability does not certify). Proposition 9.13 gives the positive connection between lumpability and CEOT IV: lumpability certifies target-law sufficiency for the next visible state. It still does not certify

$$H(S_t \mid q(S_t)) = 0,$$

nor does it certify sufficiency for arbitrary targets. The target must be the visible predictive law generated by the lumped Markov dynamics.

If q is lumpable, the visible process $q(S_t)$ has a well-defined Markov transition kernel. This is useful for reduced visible dynamics. It is not a reconstruction statement: it does not say that S_t is recoverable from $q(S_t)$.

9.7 Lumpable but not trace reconstructible

Proposition 9.15 (Lumpability does not imply hidden-state reconstruction)**.** *There exists a finite Markov chain and a lumpable compression $q : S \to \bar{S}$ such that*

$$H(S_t \mid q(S_t)) > 0.$$

Proof. Let

$$S = \{s_0, s_1\}, \qquad \bar{S} = \{\bar{s}\},$$

and define the compression by

$$q(s_0) = q(s_1) = \bar{s}.$$

Any transition kernel on S is lumpable with respect to this one-point visible compression, because the visible next state is always $\bar{s}$. In particular, the visible transition law is independent of whether the hidden representative is s_0 or s_1.

Now choose an initial law with

$$P(S_t = s_0) > 0, \qquad P(S_t = s_1) > 0$$

for the time of interest. Since $q(S_t) = \bar{s}$ is constant, conditioning on $q(S_t)$ gives no information distinguishing s_0 from s_1. Therefore

$$H(S_t \mid q(S_t)) = H(S_t) > 0.$$

Thus the compression is lumpable but not hidden-state reconstructive. □

The proposition shows that preserving visible transition laws is weaker than reconstructing hidden states. Lumpability may be useful for a reduced Markov description even when the hidden state is unrecoverable.

9.8 Not lumpable but target sufficient

The converse separation also holds: failure of lumpability does not imply failure of target sufficiency.

Proposition 9.16 (Non-lumpability does not imply target insufficiency)**.** *There exists a finite compression that is not Markov-lumpable but is sufficient for a deterministic target.*

Proof. Let $S = \{a, b, c, d\}$ and let the compression be

$$q(a) = q(b) = \bar{0}, \qquad q(c) = \bar{1}, \qquad q(d) = \bar{2}.$$

Define a deterministic transition kernel by

$$K(c \mid a) = 1, \qquad K(d \mid b) = 1.$$

Then a and b lie in the same visible fiber, but the visible next state is $\bar{1}$ from a and $\bar{2}$ from b. Hence q is not lumpable.

Now take the deterministic target to be constant:

$$Y \equiv 0.$$

For any observation O, including $O = q(S_t)$ or a compressed trace built from q, one has

$$H(Y \mid O) = 0.$$

Thus the compression is target-sufficient for this target even though it is not lumpable. □

This example is intentionally simple. It proves the logical point: non-lumpability is not, by itself, a CEOT IV target obstruction. A compression may fail to preserve visible Markov dynamics while still preserving a specified target.

9.9 Logical separation table

The preceding results can be summarized as follows.

Statement	Status	Witness or reason
Complete trace recovery implies deterministic target recovery	True	Theorem 5.7; $Y = f_+(S_{0:T})$
Deterministic target recovery implies complete trace recovery	False	Theorem 9.1
Hidden trace loss implies stochastic target loss	False	Theorem 9.3
Same observed target law up to relabeling determines hidden trace obstruction	False	Theorem 9.7
Lumpability implies hidden-state reconstruction	False	Proposition 9.15
Lumpability implies visible next-state law sufficiency	True	Proposition 9.13
Non-lumpability implies target insufficiency	False	Proposition 9.16

The table captures the main logical discipline of the paper. CEOT IV supplies exact finite obstruction criteria, but it does not identify all compression failures with one another.

9.10 Section summary

This section established the main separation and non-completeness principles of CEOT IV. It proved that hidden trace loss may occur without deterministic or stochastic target loss. It then proved the trace-level non-completeness theorem:

$$P^{\mathfrak{A}}_{O,Y} \cong_{\mathrm{law}} P^{\mathfrak{A}'}_{O',Y'} \not\Longrightarrow \mathrm{Ob}^{\mathfrak{A}}_{\mathrm{trace}}(O) = \mathrm{Ob}^{\mathfrak{A}'}_{\mathrm{trace}}(O').$$

Thus observed target behavior, even up to support relabeling, does not determine hidden trace obstruction. Finally, the section separated CEOT IV from Markov lumpability: support-relative lumpability is exactly next-visible-state sufficiency, but lumpability does not imply reconstruction, and non-lumpability does not imply task insufficiency for arbitrary targets.

The next section turns these principles into concrete finite case studies. The first case study concerns lifted dynamic programming, where memory variables may or may not be target-relevant. The second concerns signature-dependent paths, where a vertex-only path may fail to preserve signature-dependent targets.

10 Case Studies: Dynamic Programming and Signature-Dependent Paths

10.1 Purpose of the section

The preceding sections were deliberately abstract: they treated finite traces, deterministic observations, target kernels, posterior profiles, reconstruction obstructions, target-sufficiency obstructions, refinement monotonicity, profile invariance, visible-trace insufficiency, and separation principles. This section instantiates those definitions in two small algorithmic environments.

The first family of examples comes from lifted dynamic programming [16, 17]. A lifted dynamic program often augments a visible state variable by a memory, mode, budget, signature, parity, constraint, or certificate coordinate. CEOT IV asks a precise question about such augmentation:

does the hidden coordinate carry target-relevant information, or only hidden trace information?

The second family comes from signature-dependent paths. There the visible vertex path may be fixed while a hidden signature state changes the target. This is the finite path analogue of the obstruction treated in Section 8: a vertex-only trace is sufficient exactly when the target is conditionally independent of the hidden signature given the visible path.

The examples are intentionally finite and minimal. They are not meant to be efficient algorithms, nor benchmark instances, nor complexity-theoretic lower bounds. Their purpose is to show that the obstruction quantities defined above are computable certificates attached to finite presentations.

Convention 10.1 (Status of the case studies)**.** The case studies in this section are finite instantiations of the CEOT IV criteria. They are not independent completeness theorems for the full classes of dynamic programming algorithms, shortest-path algorithms, or signature-dependent path systems. Each case study fixes a finite trace-target presentation and evaluates the reconstruction, target-law, memory, or loss obstruction within that declared presentation.

Principle 10.2 (From CEOT criterion to case-study conclusion)*.* A case-study conclusion in CEOT IV has the following logical form:

$$\text{finite presentation data} \quad + \quad \text{declared target or loss} \quad + \quad \text{CEOT IV criterion} \quad \Longrightarrow \quad \text{obstruction statement insi}$$

It should not be read as a claim about all algorithms of the same informal type unless a separate class-level theorem is stated and proved.

10.2 A lifted dynamic-programming presentation

A lifted dynamic-programming trace will be written as

$$S_t = (V_t, M_t),$$

where V_t is the visible problem state and M_t is a memory coordinate. The memory coordinate may encode a constraint status, a parity bit, a resource counter, a signature label, a previously chosen mode, or a certificate flag. The vertex-only compressed observation is

$$O = V_{0:T}.$$

For a task target Y, the CEOT IV memory-to-target obstruction is

$$\mathrm{Ob}_{\mathrm{mem}\to Y} = I(Y; M_{0:T} \mid V_{0:T}).$$

By Theorem 8.11, the visible trace is conditionally target-sufficient precisely when this quantity is zero.

In dynamic programming language, Y may be an optimal action, an optimal value, a policy certificate, a feasibility certificate, or a finite future-law parameter. CEOT IV treats all of these as finite targets. The same lifted trace can be relevant for one target and irrelevant for another target; the obstruction is target-relative.

Principle 10.3 (Dynamic-programming target relativity)*.* A lifted dynamic-programming coordinate is not intrinsically relevant or irrelevant. Relative to a target Y, the visible-only trace is insufficient precisely when

$$I(Y; M_{0:T} \mid V_{0:T}) > 0.$$

In that case, some target-sufficient refinement of $V_{0:T}$ is required. The full memory trace $M_{0:T}$ may provide such a refinement, but CEOT IV does not claim that it is minimal.

The memory is hidden but target-irrelevant when

$$H(M_{0:T} \mid V_{0:T}) > 0, \qquad I(Y; M_{0:T} \mid V_{0:T}) = 0.$$

This is the operational content of the two-axis hidden-memory classification from Theorem 8.21.

10.3 Relevant-memory model: the hidden mode determines the optimal action

Consider the following one-stage lifted dynamic-programming instance. The visible state space is a singleton

$$V = \{v\},$$

and the hidden memory space is

$$M = \{0, 1\}.$$

Let

$$P(M = 0) = P(M = 1) = \frac{1}{2}.$$

There are two admissible actions,

$$A = \{a_0, a_1\}.$$

The reward is deterministic conditional on the hidden mode:

$$r(a_i, M) = \mathbf{1}_{\{i=M\}}.$$

Thus the unique optimal action is

$$Y = a_M.$$

The compressed observation is the visible state alone:

$$O = V = v.$$

Proposition 10.4 (Relevant hidden memory in lifted dynamic programming)**.** *For the above finite lifted dynamic-programming presentation,*

$$H(M \mid V) = \log_b 2,$$

and

$$I(Y; M \mid V) = \log_b 2.$$

Consequently the visible state V is not conditionally target-sufficient for the optimal-action target Y.

Proof. Since V is constant and $M \sim \text{Bernoulli}(1/2)$,

$$H(M \mid V) = H(M) = \log_b 2.$$

The target $Y = a_M$ is a deterministic bijective function of M. Therefore

$$H(Y \mid V) = H(Y) = \log_b 2,$$

while

$$H(Y \mid V, M) = 0.$$

Hence

$$I(Y; M \mid V) = H(Y \mid V) - H(Y \mid V, M) = \log_b 2.$$

By Corollary 8.17, any representation retaining only V fails the stochastic target-sufficiency criterion for Y. □

This example gives the smallest possible CEOT IV certificate for a non-cosmetic lifted state. The visible problem state contains no information about which action is optimal. The memory coordinate is not merely a hidden part of the trace; it is target-relevant.

10.4 Irrelevant-memory model: hidden memory without target relevance

Now keep the same visible and hidden spaces,

$$V = \{v\}, \qquad M = \{0, 1\}, \qquad M \sim \text{Bernoulli}(1/2),$$

but change the target. Let the two actions have the same certified value independently of the hidden mode, and let the target be the canonical tie-broken action

$$Y = a_0.$$

Equivalently, Y is constant.

Proposition 10.5 (Hidden but irrelevant memory in lifted dynamic programming)**.** *For the above finite lifted dynamic-programming presentation,*

$$H(M \mid V) = \log_b 2,$$

but

$$I(Y; M \mid V) = 0.$$

Consequently the memory coordinate is hidden from the visible state but target-irrelevant for the target Y.

Proof. Again V is constant and $M \sim \text{Bernoulli}(1/2)$, so

$$H(M \mid V) = \log_b 2.$$

However $Y = a_0$ is constant, hence

$$H(Y \mid V) = 0.$$

Since conditional mutual information is bounded above by $H(Y \mid V)$,

$$I(Y; M \mid V) \le H(Y \mid V) = 0.$$

Therefore $I(Y; M \mid V) = 0$. □

The two dynamic-programming examples have the same hidden-memory reconstruction obstruction,

$$H(M \mid V) = \log_b 2,$$

but different target obstructions. In the first example the hidden mode determines the optimal action. In the second example the hidden mode is invisible but irrelevant. This is the dynamic-programming realization of the separation

$$H(M \mid V) > 0 \not\Longrightarrow I(Y; M \mid V) > 0.$$

10.5 Value targets versus action targets

The same finite lifted dynamic program can behave differently depending on whether the target is an action, a value, or a certificate.

Return to the relevant-memory reward model of Subsection 10.3:

$$r(a_i, M) = \mathbf{1}_{\{i=M\}}.$$

The optimal action target is

$$Y_{\text{act}} = a_M,$$

and Proposition 10.4 gives

$$I(Y_{\text{act}}; M \mid V) = \log_b 2.$$

However the optimal value is

$$Y_{\text{val}} = \max_{i \in \{0,1\}} r(a_i, M) = 1,$$

which is constant. Therefore

$$I(Y_{\text{val}}; M \mid V) = 0.$$

Thus the same hidden mode is target-relevant for recovering the optimal action, but target-irrelevant for recovering the optimal value.

Corollary 10.6 (Action/value target separation)**.** *There exists a finite lifted dynamic-programming presentation for which*

$$I(Y_{\mathrm{act}}; M \mid V) > 0, \qquad I(Y_{\mathrm{val}}; M \mid V) = 0.$$

Hence target relevance is not a property of the lifted state alone; it is a property of the pair consisting of the lifted state and the declared target.

This corollary is important for algorithmic compression. A compressed representation may be valid for computing values but invalid for extracting policies; valid for feasibility but invalid for witnesses; valid for costs but invalid for certificates. CEOT IV forces the target to be named before the obstruction is interpreted.

10.6 Case Study II: signature-dependent paths

The second case study concerns path problems in which the visible vertex path does not contain all target-relevant information. Let the visible graph have the deterministic path

$$a \longrightarrow b \longrightarrow c.$$

Let the lifted state space factor as

$$\mathcal{S}_t = \mathcal{V}_t \times \Sigma_t,$$

and write the state random variable as

$$S_t = (V_t, \Sigma_t),$$

where V_t is the visible vertex random variable and Σ_t is a hidden signature mode. Assume the visible trace is always

$$V_{0:2} = (a, b, c) \quad \text{a.s.},$$

while the final signature is balanced:

$$P(\Sigma_2 = +) = P(\Sigma_2 = -) = \frac{1}{2}.$$

Let the target be

$$Y = \mathbf{1}_{\{\Sigma_2 = +\}}.$$

The compressed observation is the vertex trace

$$O = V_{0:2}.$$

Proposition 10.7 (Signature-dependent path obstruction)**.** *In the above finite signature-dependent path presentation,*

$$I(Y; \Sigma_{0:2} \mid V_{0:2}) = H(Y \mid V_{0:2}) = \log_b 2.$$

Consequently the visible vertex path is not conditionally target-sufficient for Y.

Proof. The visible trace $V_{0:2}$ is almost surely constant. Since Σ_2 is balanced and $Y = \mathbf{1}_{\{\Sigma_2=+\}}$,

$$Y \sim \text{Bernoulli}(1/2)$$

conditional on the visible trace. Therefore

$$H(Y \mid V_{0:2}) = \log_b 2.$$

On the other hand, Y is a deterministic function of the signature trace, since it is determined by Σ_2. Hence

$$H(Y \mid V_{0:2}, \Sigma_{0:2}) = 0.$$

Thus

$$I(Y; \Sigma_{0:2} \mid V_{0:2}) = H(Y \mid V_{0:2}) - H(Y \mid V_{0:2}, \Sigma_{0:2}) = \log_b 2.$$

□

The conclusion is not that every signature coordinate must be retained in every path algorithm. The conclusion is conditional: if the declared target distinguishes signature modes inside a visible path fiber, then the vertex-only trace is not sufficient for that target.

10.7 General signature-fiber criterion

The previous example is a two-signature special case of the same-visible-fiber criterion from Proposition 8.19. Let a lifted path trace have coordinate decomposition

$$S_{0:T} = (V_{0:T}, \Sigma_{0:T}),$$

where $V_{0:T}$ is the visible vertex trace and $\Sigma_{0:T}$ is a hidden signature, mode, parity, or algebraic-label trace. Let

$$O = V_{0:T}$$

be the vertex-only observation.

Proposition 10.8 (General signature-fiber obstruction)**.** *Suppose there exists a visible path value $v_{0:T}$ of positive probability and two signature trace values $\sigma_{0:T}$ and $\sigma'_{0:T}$ with positive conditional probability given $V_{0:T} = v_{0:T}$ such that*

$$L_+(\cdot \mid v_{0:T}, \sigma_{0:T}) \neq L_+(\cdot \mid v_{0:T}, \sigma'_{0:T}).$$

Then

$$I(Y; \Sigma_{0:T} \mid V_{0:T}) > 0.$$

In particular, the vertex-only observation is not conditionally target-sufficient.

Proof. The assumption says that, on a positive-probability visible fiber, the support-restricted target kernel is not constant as a function of the hidden signature trace. Hence the conditional target law is not measurable with respect to the visible trace alone. By the finite conditional-independence criterion for stochastic target sufficiency, equivalently Theorem 6.10 applied to the lifted decomposition,

$$Y \not\perp \Sigma_{0:T} \mid V_{0:T},$$

and therefore

$$I(Y; \Sigma_{0:T} \mid V_{0:T}) > 0.$$

□

The proposition gives a direct test for signature-state target relevance. It does not require reconstructing the entire hidden signature trace. It only requires detecting whether the target law varies inside a visible-path fiber.

10.8 Signature state irrelevant to a visible target

For completeness, consider the opposite case. Let the lifted path state again have coordinate decomposition

$$S_{0:T} = (V_{0:T}, \Sigma_{0:T}),$$

but let the target depend only on the visible path:

$$Y = \psi_+(V_{0:T}).$$

Then Y is already a deterministic function of the vertex-only observation. Hence

$$H(Y \mid V_{0:T}) = 0,$$

and therefore

$$I(Y; \Sigma_{0:T} \mid V_{0:T}) = 0.$$

Corollary 10.9 (Visible path targets do not require hidden signatures)**.** *If a target Y is a function of the visible path $V_{0:T}$ alone, then the vertex-only observation is sufficient for Y, regardless of whether*

$$H(\Sigma_{0:T} \mid V_{0:T})$$

is positive.

Thus signature reconstruction and signature target relevance are distinct. A signature trace may be hidden from the visible path while still being irrelevant to a visible target.

10.9 Relation to stochastic shortest paths

Signature-dependent stochastic shortest paths fit exactly into the finite CEOT IV template. The lifted state has the coordinate form

$$S_t = (V_t, \Sigma_t),$$

where V_t is the current graph vertex and Σ_t is a hidden or accumulated label. The target may be any finite object attached to the path, for example:

$$\begin{gathered}\text{selected action}, \quad \text{selected path}, \quad \text{terminal mode}, \quad \text{cost bucket}, \\ \text{feasibility certificate}, \quad \text{future-law parameter}.\end{gathered}$$

The vertex-only representation is CEOT IV sufficient exactly when

$$I(Y; \Sigma_{0:T} \mid V_{0:T}) = 0.$$

If this quantity is positive, then a vertex-only algorithm may still compute some coarser statistic, but it cannot be certified as sufficient for the declared target Y.

This gives a clean obstruction-theoretic version of the usual state-augmentation intuition. One should not retain a signature merely because it exists, nor discard it merely because the visible graph is simple. The decision depends on the target law inside the visible-path fibers.

10.10 Case-study summary table

The case studies can be summarized as follows.

Case	Obstruction value	Interpretation
Relevant DP memory	$I(Y; M \mid V) = \log_b 2$	visible state does not determine the optimal-action target
Irrelevant DP memory	$I(Y; M \mid V) = 0$	memory is hidden but not target-relevant
Action/value separation	$I(Y_{\text{act}}; M \mid V) > 0$, $I(Y_{\text{val}}; M \mid V) = 0$	the same hidden mode may be relevant for policies but not values
Balanced signature path	$I(Y; \Sigma_{0:2} \mid V_{0:2}) = \log_b 2$	vertex-only path loses a signature-dependent target
Visible path target	$I(Y; \Sigma_{0:T} \mid V_{0:T}) = 0$	hidden signature reconstruction is irrelevant for visible targets

The table illustrates the main thesis of CEOT IV within the declared finite presentations. Hidden state, hidden memory, and hidden signature information become obstructions only after a target has been declared. Conditional entropy measures reconstruction loss. Conditional mutual information measures target-relevant loss.

10.11 Section summary

This section instantiated the abstract CEOT IV criteria in finite algorithmic examples. Within the declared lifted dynamic-programming presentations, a hidden memory coordinate can be target-relevant, target-irrelevant, or relevant for one target and irrelevant for another. Within the declared signature-dependent path presentations, a vertex-only path has a nonzero target-law obstruction exactly when the target law varies across hidden signature traces inside a visible-path fiber.

The next section relates these conclusions back to CEOT I–III and explains why a later CEOT V should move from finite presentations to finite factors, measurable reconstruction, and possibly randomized observation channels.

11 Relation to CEOT I–III and Outlook to CEOT V

11.1 Purpose of the section

The preceding sections developed CEOT IV as a finite theory of algorithmic compression. The objects are no longer only hidden variables, endpoints, or marked coordinates in a static diagram. They are finite computational traces

$$S_{0:T} = (S_0, S_1, \dots, S_T),$$

deterministic compressed observations

$$O = c(S_{0:T}),$$

and task targets generated by deterministic maps or stochastic kernels. The central distinction is therefore

$$\text{trace reconstruction loss} \qquad \text{versus} \qquad \text{target-relevant information loss.}$$

This final section locates that distinction inside the CEOT sequence. CEOT I supplies the finite support-relative reconstruction principle. CEOT II supplies the endpoint/interior and bridge viewpoint. CEOT III supplies the marked-presentation, profile-first, and non-completeness discipline. CEOT IV imports all three layers and adds algorithmic time, compression maps, target kernels, deterministic and stochastic sufficiency criteria, and visible-trace insufficiency.

The section also explains why the next natural step should not be a larger finite case study, but a change of setting. CEOT V should pass from finite presentations to measurable systems through finite factors and increasing observation sigma-algebras. That extension must preserve the support-relative and profile-first discipline of the finite theory while avoiding false claims about differential entropy, infinite conditional entropy, or unrestricted randomized channels.

11.2 Relation to CEOT I: finite support-relative reconstruction

CEOT I established the basic finite obstruction principle: for finite random variables U and O,

$$H(U \mid O) = 0 \quad \Longleftrightarrow \quad U = r(O) \quad \text{a.s.}$$

for some support-relative decoder $r : \mathrm{supp}(O) \to \mathrm{supp}(U)$. CEOT IV uses exactly the same finite principle, but applies it to several different variables extracted from an algorithmic presentation:

$$U = S_{0:T}, \qquad U = S_{\mathrm{int}}, \qquad U = (S_{\mathrm{int}} \mid S_0, S_T), \qquad U = Y.$$

Thus the CEOT I obstruction is not replaced. It is replicated at different semantic levels of the trace-target presentation.

Proposition 11.1 (CEOT I as the reconstruction core of CEOT IV)**.** *Let A be any finite random variable defined on a CEOT IV presentation and let $O = c(S_{0:T})$ be its deterministic observation. Then*

$$H(A \mid O) = 0 \quad \Longleftrightarrow \quad A = r(O) \quad \textit{a.s.}$$

for a support-relative decoder r. In particular, the complete-trace, hidden-trace, deterministic-target, and memory-reconstruction zero criteria are all CEOT I reconstruction criteria applied to different choices of A.

Proof. This is Lemma 4.1 applied to the chosen finite variable A. The complete-trace criterion follows by taking $A = S_{0:T}$, the hidden-trace criterion by taking $A = S_{\mathrm{int}}$, deterministic target recovery by taking $A = Y$, and memory reconstruction by taking $A = M_{0:T}$ in a lifted trace. □

The conceptual change in CEOT IV is not the zero-entropy lemma. The change is the classification of which finite variable should be reconstructed. A compression may fail to reconstruct the trace while still reconstructing the declared target. Conversely, a hidden coordinate may be unreconstructible and yet irrelevant to the target law. CEOT IV therefore keeps the CEOT I obstruction calculus but refuses to use one reconstruction demand as a proxy for all others.

Principle 11.2 (CEOT I principle inside CEOT IV). CEOT IV is CEOT I applied repeatedly to algorithmically meaningful random variables. The mathematical criterion remains conditional entropy zero; the new content is the separation between complete trace, hidden trace, bridge trace, deterministic target, stochastic target, and lifted memory.

11.3 Relation to CEOT II: endpoint/interior and bridge separation

CEOT II introduced the endpoint/interior separation viewpoint. In CEOT IV, this becomes the distinction between complete trace reconstruction,

$$\mathrm{Ob}_{\mathrm{complete}}(O) = H(S_{0:T} \mid O),$$

hidden internal trace reconstruction,

$$\mathrm{Ob}_{\mathrm{trace}}(O) = H(S_{\mathrm{int}} \mid O),$$

and endpoint-conditioned bridge reconstruction,

$$\mathrm{Ob}_{\mathrm{bridge}}(O) = H(S_{\mathrm{int}} \mid O, S_0, S_T).$$

The bridge obstruction is the direct algorithmic analogue of the CEOT II endpoint-conditioned interior obstruction. It asks what internal computational history remains uncertain after both the compressed observation and the endpoints are known.

The hierarchy proved in Theorem 4.24,

$$\mathrm{Ob}_{\mathrm{bridge}}(O) \le \mathrm{Ob}_{\mathrm{trace}}(O) \le \mathrm{Ob}_{\mathrm{complete}}(O),$$

should be read as the CEOT IV version of endpoint/interior discipline. Conditioning on endpoints can only reduce hidden interior uncertainty. Reconstructing the full trace is at least as strong as reconstructing the internal trace. Neither implication should be reversed without additional structural assumptions.

Proposition 11.3 (CEOT II bridge obstruction inside CEOT IV)**.** *For a finite trace $S_{0:T}$ and deterministic observation $O = c(S_{0:T})$, the bridge obstruction*

$$H(S_{\mathrm{int}} \mid O, S_0, S_T)$$

is the endpoint-conditioned CEOT II obstruction for the algorithmic interior S_{int} relative to the observed data (O, S_0, S_T).

Proof. The CEOT II endpoint/interior pattern conditions on endpoint information and measures remaining uncertainty of the interior. In CEOT IV the endpoints are S_0 and S_T, the observed compression is O, and the interior variable is S_{int}. Substitution gives exactly the displayed conditional entropy. □

The bridge viewpoint is important because many algorithms care about endpoints, outputs, or selected terminal states, not about the entire internal path. A zero complete-trace obstruction is sufficient for all endpoint and target questions, but it is often much too strong. CEOT IV keeps the bridge obstruction as an intermediate reconstruction demand between full history recovery and pure target sufficiency.

11.4 Relation to CEOT III: marked presentations and profile-first obstruction

CEOT III studied finite marked sampleable stochastic presentations. Its typical collapse had the form

$$\mathbf{X}_V \longmapsto \mathbf{X}_{V_{\mathrm{obs}}},$$

where $V_{\text{obs}} \subseteq V$ marked the observed coordinates. The CEOT III obstruction measured the hidden uncertainty left by that marked projection. CEOT IV replaces the static marked-coordinate collapse by a finite trace-law compression,

$$S_{0:T} \longmapsto O = c(S_{0:T}).$$

The formal analogy is

$$H(\mathbf{X}_{V \setminus V_{\text{obs}}} \mid \mathbf{X}_{V_{\text{obs}}}) \quad \rightsquigarrow \quad H(S_{\text{int}} \mid O),$$

or, when complete trace reconstruction is demanded,

$$H(\mathbf{X}_V \mid \mathbf{X}_{V_{\text{obs}}}) \quad \rightsquigarrow \quad H(S_{0:T} \mid O).$$

The arrow is not an identity of categories. It is a transfer of obstruction logic from finite diagram collapse to finite ordered trace compression.

The shared discipline is that the posterior profile is primary. The scalar entropy is only a numerical shadow of the posterior fibers:

$$\mathsf{Post}_{A|O} \quad \longmapsto \quad H(A \mid O).$$

This is why Section 3 introduced posterior reconstruction profiles before assigning scalar obstruction values, and why Section 7 proved groupoid-level profile invariance before extracting the scalar factorization.

Principle 11.4 (Profile-first continuity from CEOT III to CEOT IV). CEOT III and CEOT IV share the same invariant discipline:

$$\text{presentation} \longrightarrow \text{posterior profile} \longrightarrow \text{scalar entropy obstruction}.$$

The scalar obstruction is invariant under relabeling only because the posterior profile is invariant under the corresponding support-relative finite presentation isomorphism.

The non-completeness theorem of Section 9 is the algorithmic counterpart of the CEOT III non-completeness phenomenon. In CEOT III, equal endpoint composites or equal observed behavior need not determine the hidden completion obstruction. In CEOT IV, isomorphic compressed target behavior need not determine the hidden trace obstruction:

$$P^{\mathfrak{A}}_{O,Y} \cong_{\text{law}} P^{\mathfrak{A}'}_{O',Y'} \not\Longrightarrow \text{Ob}^{\mathfrak{A}}_{\text{trace}}(O) = \text{Ob}^{\mathfrak{A}'}_{\text{trace}}(O').$$

This is not a defect. It is the correct invariance boundary. The hidden obstruction is attached to a specified presentation and its posterior trace fibers, not merely to the marginal law of the compressed target pair.

11.5 What CEOT IV adds

CEOT IV adds five structural ingredients to the earlier finite theory.

First, it introduces ordered traces. The object of obstruction is no longer only a hidden variable or a hidden diagram coordinate, but a finite algorithmic history

$$S_0 \to S_1 \to \cdots \to S_T.$$

This makes it possible to distinguish endpoints, internal states, memory coordinates, visible coordinates, and hidden signatures.

Second, it separates reconstruction from task sufficiency. Conditional entropy answers reconstruction questions such as

$$H(S_{\mathrm{int}} \mid O) = 0.$$

Conditional mutual information answers stochastic sufficiency questions such as

$$I(Y; S_{0:T} \mid O) = 0.$$

The two quantities have different meanings and should not be merged.

Third, it includes deterministic targets as a special but important case. If

$$Y = f_{+}(S_{0:T}),$$

then stochastic sufficiency reduces to deterministic recovery because

$$I(Y; S_{0:T} \mid O) = H(Y \mid O).$$

This explains why ordinary exact-output preservation is compatible with the broader stochastic target framework.

Fourth, it provides a finite visible-trace insufficiency theorem. If

$$S_t = (V_t, M_t),$$

then visible trace sufficiency for Y is exactly the conditional-independence statement

$$Y \perp M_{0:T} \mid V_{0:T}.$$

Equivalently, the obstruction is

$$I(Y; M_{0:T} \mid V_{0:T}).$$

Hidden memory is therefore not automatically required. It is target-relevant precisely when it changes the target law inside a visible-trace fiber; even then, a smaller sufficient refinement than the full memory may exist.

Fifth, CEOT IV separates observation refinement from presentation isomorphism. Observation refinement is a poset relation inside one trace law:

$$O \preceq O' \iff \sigma(O) \subseteq \sigma(O') \subseteq \sigma(S_{0:T}).$$

Support-relative presentation isomorphism is a groupoid-level relabeling relation between finite positive-support presentations. Monotonicity belongs to the first structure; invariance belongs to the second.

11.6 The CEOT IV theorem package

The results of the paper can be grouped into five clusters.

Cluster	Main content
Reconstruction criteria	Complete-trace, hidden-trace, bridge, static, and memory zero criteria; see Theorems 4.12, 4.15, 4.18, and Proposition 4.20.
Target sufficiency	Deterministic recovery, complete-trace sufficiency for deterministic targets, stochastic target sufficiency, and deterministic recovery as a special case; see Theorems 5.5, 5.7, 6.10, and 6.19.
Monotonicity and invariance	Observation-refinement monotonicity and finite profile factorization; see Theorems 7.8, 7.10, 7.12, and 7.27.
Non-completeness	Isomorphic observed target law does not determine hidden trace obstruction; see Theorem 9.7.
Visible-trace insufficiency	Visible trace sufficiency is exactly conditional independence from the hidden memory trace; see Theorem 8.11 and Proposition 8.19.

The package is deliberately finite. Every theorem is stated for finite state spaces, deterministic observations of the trace, and finite targets or finite target-law parameters. This restriction is not cosmetic. It ensures that conditional entropy, conditional mutual information, posterior fibers, support-relative reconstruction, and groupoid-level profile invariance have unambiguous meanings.

11.7 Boundary with lumpability and classical state reduction

Classical Markov lumpability and CEOT IV target sufficiency answer different questions. Lumpability asks whether a visible process has closed Markov transition dynamics. CEOT IV target sufficiency asks whether a compressed observation preserves the target law. These conditions are independent in both directions, as shown by Propositions 9.15 and 9.16.

The practical lesson is that no single compression certificate should be used for all algorithmic purposes. A visible state may be a valid Markov lumping and still fail to reconstruct hidden computational history. A visible state may fail Markov lumpability and still be sufficient for a constant or otherwise visible target. A lifted state may be irrelevant for one target and relevant for another. The correct CEOT IV question is always:

Which variable or target law is the compression required to preserve?

Once that question is fixed, the relevant obstruction is one of

$$H(S_{0:T} \mid O), \quad H(S_{\text{int}} \mid O), \quad H(S_{\text{int}} \mid O, S_0, S_T), \quad H(Y \mid O), \quad I(Y; S_{0:T} \mid O).$$

Without the target declaration, the compression problem is under-specified.

11.8 Limits of the present finite theory

CEOT IV proves exact finite obstruction criteria, but it does not claim more than those criteria supply. The following limits are structural.

(i) The observation O is deterministic: $O = c(S_{0:T})$. Randomized observation channels are deferred.

(ii) The state spaces and targets are finite. No claim is made about differential entropy or unrestricted infinite-state Shannon entropy.

(iii) Zero obstruction is an exact reconstructibility or sufficiency statement. It is not a computational complexity bound.

(iv) Positive obstruction proves failure of the declared exact criterion. It does not by itself quantify approximate algorithmic performance.

(v) Groupoid-level invariance is finite relabeling invariance. It is not a universal categorical entropy functor on arbitrary stochastic systems.

These restrictions are what make the results sharp. They also identify the correct entry points for future extensions: randomized observations, measurable state spaces, finite-factor approximations, and approximate target sufficiency.

Remark 11.5 (Finite scope of CEOT IV). All reconstruction, sufficiency, profile, and functoriality theorems in CEOT IV are finite and support-relative. This paper does not prove a measurable reconstruction theorem, does not use differential entropy as a reconstruction obstruction, and does not assert convergence of finite obstruction profiles to an infinite-state invariant.

11.9 Outlook to CEOT V: finite factors and measurable reconstruction

The next stage should move from finite presentations to measurable systems without abandoning finite certificates. This extension is not automatic. The finite equivalence

$$H(A \mid O) = 0 \iff A = r(O) \quad \text{a.s.}$$

must be replaced by an appropriate measurable statement involving regular conditional laws and measurable decoders.

A CEOT V measurable theory should address at least four separate problems:

(i) define admissible measurable trace presentations and observation sigma-algebras;

(ii) formulate zero reconstruction through almost-sure measurable decoders rather than ambient pointwise decoders;

(iii) replace finite entropy obstructions by measure-theoretic conditional information quantities where they are finite, or by finite-factor approximations where they are not;

(iv) prove which finite CEOT IV profile invariants survive under projective limits, refinements of finite partitions, or other limiting procedures.

No such measurable theorem is asserted in CEOT IV.

A plausible CEOT V setting is a standard Borel trace or hidden variable U, an observed measurable variable O, and families of finite-valued measurable factors

$$\pi : U \to F, \qquad \gamma : O \to G.$$

For each finite pair (π, γ) one obtains a finite obstruction

$$\mathrm{Ob}_{\pi,\gamma}(U, O) := H(\pi(U) \mid \gamma(O)).$$

Thus measurable reconstruction would be tested through finite shadows rather than through a single potentially ill-behaved infinite entropy.

For traces, the analogous finite-factor obstruction would use measurable factors of the whole trace,

$$\pi(S_{0:T}),$$

and finite factors of the observation,

$$\gamma(O).$$

The finite CEOT IV obstruction then becomes

$$H(\pi(S_{0:T}) \mid \gamma(O)),$$

or, for hidden internal trace questions,

$$H(\pi(S_{\mathrm{int}}) \mid \gamma(O)).$$

Task-relative versions would use finite target factors or finite target-law parameters.

A careful CEOT V should also use increasing observation sigma-algebras. If

$$\mathcal{G}_1 \subseteq \mathcal{G}_2 \subseteq \cdots \subseteq \sigma(O),$$

then finite-factor monotonicity should mirror the CEOT IV refinement theorem:

$$H(\pi(U) \mid \mathcal{G}_{n+1}) \leq H(\pi(U) \mid \mathcal{G}_n)$$

whenever the displayed finite conditional entropies are represented through finite factors. The measurable theory should then ask whether the limiting finite-factor obstruction vanishes for a declared class of hidden factors π.

Principle 11.6 (Finite factors before measurable entropy). A measurable CEOT should not begin with a global infinite-state entropy. It should begin with finite hidden factors, finite observed factors, and monotone families of finite certificates. The finite theory is the certificate layer; the measurable theory is the organization of these finite certificates.

Remark 11.7 (Differential entropy is not a reconstruction obstruction). A measurable or continuous CEOT theory should not replace finite conditional entropy by differential entropy without additional structure. Differential entropy is coordinate-dependent, may be negative, and does not by itself characterize almost-sure reconstruction. Any CEOT V extension must instead work with measurable conditional laws, mutual information defined measure-theoretically, finite partitions, or another declared invariant with a proved reconstruction interpretation.

Problem 11.8 (Finite-factor approximation problem for CEOT V)**.** Let (A, O) be measurable random variables and let

$$\mathcal{P}_1 \preceq \mathcal{P}_2 \preceq \cdots$$

be an increasing sequence of finite partitions or finite observation factors. Determine conditions under which the finite CEOT IV obstructions of the factorized variables

$$A_n = \pi_n(A), \qquad O_n = \rho_n(O)$$

converge to a meaningful measurable obstruction, and determine whether zero limiting obstruction implies the existence of an almost-sure measurable decoder.

Randomized observation channels are a second CEOT V direction. If the observation is generated by a channel

$$C : S_{0:T} \rightsquigarrow O$$

rather than a deterministic map, then the finite compatibility and refinement statements must be restated. One must specify whether the target channel and observation channel are conditionally independent given the trace, and deterministic refinement should be replaced by a stochastic comparison order for observation channels. CEOT IV deliberately avoids this issue so that the deterministic finite theory is closed and exact.

The role of CEOT IV is therefore foundational: it provides the finite support-relative obstruction calculus and the typed profile language. CEOT V must decide which of these finite statements admit measurable analogues and which fail without additional regularity assumptions.

11.10 Final summary

CEOT IV completes the finite algorithmic-compression layer of the CEOT sequence. Its main point is the separation

$$\boxed{\text{hidden computational trace loss} \neq \text{task-relevant information loss}}.$$

Conditional entropy measures reconstruction obstruction. Conditional mutual information measures stochastic target-law sufficiency obstruction, and its approximate form controls conditional total variation and bounded-loss degradation. Complete trace recovery, hidden trace recovery, bridge recovery, deterministic target recovery, stochastic law-sufficiency, loss-relative sufficiency, and visible-trace insufficiency are therefore distinct finite questions.

The paper also fixes the correct invariance boundary. The obstruction is invariant under support-relative presentation isomorphism and monotone under deterministic observation refinement, but it is not determined by compressed target behavior alone. This is the algorithmic counterpart of the non-completeness phenomena already present in CEOT III.

The resulting theory is intentionally modest and exact: finite traces, deterministic compressed observations, finite posterior profiles, support-relative reconstruction, and task-relative targets. That finite core is the correct platform for CEOT V, where the same obstruction logic should be rebuilt through finite factors inside measurable systems.

References

[1] C. E. Shannon. A mathematical theory of communication. *Bell System Technical Journal*, 27:379–423, 623–656, 1948.

[2] T. M. Cover and J. A. Thomas. *Elements of Information Theory*. Wiley-Interscience, second edition, 2006.

[3] I. Csiszár and J. Körner. *Information Theory: Coding Theorems for Discrete Memoryless Systems*. Cambridge University Press, second edition, 2011.

[4] P. R. Halmos and L. J. Savage. Application of the Radon–Nikodym theorem to the theory of sufficient statistics. *The Annals of Mathematical Statistics*, 20(2):225–241, 1949.

[5] R. R. Bahadur. Sufficiency and statistical decision functions. *The Annals of Mathematical Statistics*, 25(3):423–462, 1954.

[6] E. L. Lehmann and G. Casella. *Theory of Point Estimation.* Springer, second edition, 1998.

[7] A. W. van der Vaart. *Asymptotic Statistics.* Cambridge University Press, 1998.

[8] E. Torgersen. *Comparison of Statistical Experiments.* Cambridge University Press, 1991.

[9] J. R. Norris. *Markov Chains.* Cambridge University Press, 1998.

[10] Duy Hoang Nguyen. *Finite Categorical Entropy Obstruction Calculus I: Stochastic Factorizations and Bridge Reconstruction.* Zenodo, 2026. doi:10.5281/zenodo.20500996.

[11] Duy Hoang Nguyen. *Categorical Entropy Obstruction Theory II: Scaling Bridge Obstructions and Endpoint/Interior Reconstruction Separation.* Zenodo, 2026. doi:10.5281/zenodo.20506069.

[12] Duy Hoang Nguyen. *Categorical Entropy Obstruction Theory III: Diagram-Collapse Obstructions for Finite Marked Sampleable Stochastic Presentations.* Zenodo, 2026. doi:10.5281/zenodo.20518578.

[13] S. Mac Lane. *Categories for the Working Mathematician.* Springer, second edition, 1998.

[14] D. Blackwell. Equivalent comparisons of experiments. *The Annals of Mathematical Statistics*, 24(2):265–272, 1953.

[15] J. G. Kemeny and J. L. Snell. *Finite Markov Chains.* Van Nostrand, 1960.

[16] R. Bellman. *Dynamic Programming.* Princeton University Press, 1957.

[17] D. P. Bertsekas. *Dynamic Programming and Optimal Control.* Athena Scientific, fourth edition, 2012.