WHEN ARE NON-INVASIVE MEDICAL DIGITAL TWINS IDENTIFIABLE? TRACING CLINICAL IDENTIFIABILITY VIA A CASCADED LENS

Anonymous authors
Paper under double-blind review

ABSTRACT

Non-invasive medical digital twins seek to infer patient-specific physiology from indirect clinical observations. However, distinct latent physiological states and model parameters may produce indistinguishable measurements, leaving their interpretation ambiguous even when a twin accurately reproduces the observed data. We study clinical identifiability: whether physiological states, parameters, or clinically prescribed summaries are uniquely determined by the available observations. Specifically, we formulate non-invasive digital twins as cascaded dynamics that separate upstream physiology from downstream observation processes. This structure yields a low-order sufficient criterion for local observability and structural identifiability, certifying recoverability from a small collection of Lie-derivative blocks without constructing the full observability-identifiability matrix. We further combine admissible reduction with functional identifiability to assess when clinically meaningful physiological summaries remain recoverable despite ambiguity in the full latent state. Experiments on synthetic cascades and a cerebral hemodynamics digital twin constructed from perfusion MRI evaluate low-order certification, computational cost, and the stability of reduced physiological targets across observationally similar solutions. Together, these analyses provide a principled basis for matching model resolution, observation pathways, and clinical inferential targets to the information supported by non-invasive measurements.

1 INTRODUCTION

Medical digital twins are emerging as a promising route to patient-specific physiological modeling in healthcare (Laubenbacher et al., 2024; Sun et al., 2023). Rather than treating clinical data as inputs to a purely predictive system, this paradigm seeks to couple patient-specific observations with executable models of underlying physiology. Such models provide a basis for in silico experimentation, intervention analysis, and decision support, while also shifting the goal from prediction alone to patient-specific modeling of latent physiology (Moingeon et al., 2023; Laubenbacher et al., 2021).

Realizing this goal is especially challenging in the *non-invasive* regime. Although routine clinical acquisition may retain signatures of the underlying physiological states to be encoded by a digital twin, these latent quantities do not appear in the observed measurements. This naturally forms an inverse problem where latent physiology must be inferred from downstream observations rather than directly measured. Estimation from non-invasive data therefore takes the form of a composite inverse problem (Kuang et al., 2024), where a physiological model is coupled with a downstream observation process, mapping latent physiological states to observed measurements.

Essentially, this composite inverse problem has an inherently dynamical structure: the latent physiological variables of interest evolve according to governing state equations, whereas the available data arise only through downstream observation of that evolution. Before proceeding to downstream tasks such as decision-making and treatment planning, it is therefore essential to determine whether the observations uniquely support the quantities required for its intended use. Distinct latent trajectories or parameterizations may reproduce the same measurements while yielding different physiological interpretations, eventually leading to ambiguity in inference and control. However, recovering every latent state and parameter may be unnecessary: clinically prescribed summaries can remain

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

identifiable even when finer physiological details are ambiguous. This motivates assessing clinical identifiability at the level of the intended physiological target.

Therefore, we ask a fundamental question before estimation: *can the available observation process uniquely determine the latent state trajectory and model parameters encoded by the twin, or at least the clinically meaningful functionals derived from them?* In this paper, we study this question through observability and functional identifiability for medical digital twins characterized by physiological dynamical systems in the non-invasive latent-to-observation regime, providing a principled basis for well-posed inference, model design, parameterization, and learning. Without such guarantees, clinically plausible modeling choices may yield multiple latent explanations for the same data, undermining physiological interpretability and reliability in clinical use.

Contribution. We provide a cascaded lens for understanding not only *whether* a non-invasive digital twin is identifiable, but also which clinically prescribed quantities are *supported by the observations*, and at what model resolution. In this view, identifiability is a design principle: it guides the choice of model resolution, observation pathway, and inferential target before estimation is performed.

- We formulate non-invasive medical digital twins as *cascaded dynamical systems* that separate upstream latent physiological dynamics from downstream observation dynamics.
- We derive a *low-order block criterion* for local observability and structural identifiability in cascaded systems. We show that identifiability can be certified from a small collection of low-order Lie-derivative blocks rather than the full observability-identifiability matrix, with theory and experiments given in Theorem 3.4 and Figure 3.
- We study *admissible reduction* and *functional identifiability*, which provide a principled way to replace unsupported full-state recovery with reduced targets of clinical interest, with theory and experiments given in Theorem 3.9 and Figure 4.
- We evaluate both the structural and functional perspectives in controlled synthetic cascaded systems and in a cerebral hemodynamics digital twin, showing how modeling choices affect identifiability and when reduced physiological summaries remain stable despite ambiguity in full-state recovery.

Notation. Df denotes the Jacobian of a differentiable map f. The row space of a matrix A is denoted by $\mathrm{row}(A)$. Uppercase (X, Z) denotes augmented states, lowercase (x, z) denotes generic states.

2 Background and Preliminaries

We first formulate non-invasive medical digital twins as composite inverse problems in Section 2.1. Then, we recall observability and structural identifiability for nonlinear dynamical systems in Section 2.2. These preliminaries provide the dynamical-systems foundation for our analysis in Section 3.

2.1 Non-Invasive Medical Digital Twins

We formalize a medical digital twin as a physics-informed dynamical model $\mathcal{M} : \Theta \to \mathcal{X}$, where Θ denotes physiological parameters and $\mathcal{X}$ denotes the corresponding latent physiological states or trajectories (Kuang et al., 2024). The "twin" here is not the observed clinical data themselves, but the underlying physiological process that gives rise to them. Clinical measurements are represented by an observed data variable $y \in \mathcal{Y}$, where $\mathcal{Y}$ denotes the observation space. These measurements arise through a separate data formation map $\mathcal{D} : \mathcal{X} \to \mathcal{Y}$ so that the forward relation takes the form $y = (\mathcal{D} \circ \mathcal{M})(\theta)$. This decomposition makes explicit the structural gap between the latent physiology and its clinical observation, clarifying how invasive and non-invasive data acquisition regimes differ.

The importance of the non-invasive regime comes from a basic clinical mismatch. Specifically, the physiological quantities most relevant for constructing a medical digital twin are often not directly available under routine acquisition, whereas the measurements that are routinely available arise only through downstream sensing or imaging processes. For instance, in cerebral hemodynamics digital twin modeling, one may wish to incorporate physiological signals such as arterial blood pressure and intracranial pressure (Briggs et al., 2026). Yet these bedside measurements are not standardized in routine clinical practice. In cardiology, a similar gap appears in pressure-volume (PV) analysis: PV loops are highly informative but traditionally require invasive catheterization (Kass

et al., 1986), making non-invasive identification of a patient-specific cardiac digital twin especially valuable (Kuang et al., 2024).

Accordingly, our structural identifiability analysis focuses on the non-invasive regime, where latent physiology is coupled to observations through a downstream acquisition process with its own dynamics. This encompasses a broad class of physics-informed digital twins in which clinically relevant parameters are inferred only indirectly from routine data.

2.2 Observability and Structural Identifiability for Dynamical Systems

To analyze the non-invasive medical digital twin formulation systematically, we recall the notions of observability (Hermann & Krener, 1977) and structural identifiability (Villaverde et al., 2019; 2016) for nonlinear dynamical systems. In essence, structural identifiability asks whether the governing dynamics together with the observation process uniquely determine the underlying latent states and parameters from the measured trajectory. Failure of structural identifiability implies that distinct latent explanations remain observationally indistinguishable, and therefore any subsequent estimation effort is intrinsically ambiguous and may lead to unreliable estimates, wasted resources, and misleading conclusions (Heinemann & Raue, 2016).

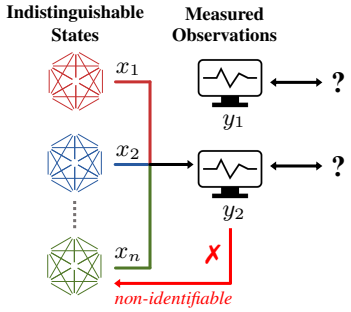


Figure 1: Distinguishability.

Dynamical Systems Setup. Consider a parametrized dynamical system (possibly nonlinear) $\mathbf{\Sigma}$:

$$\mathbf{\Sigma} : \dot{x}(t) = f(x(t), u(t), \theta), \quad y(t) = g(x(t)), \quad x(t_0) = x_0, \quad t \in [0, T],$$

where $x(t)$ denotes the state with initial condition x_0. $u(t)$ is a known input vector, and $y(t)$ denotes an observable vector obtained through an analytic vector-valued map $g(\cdot)$.

Distinguishability. As illustrated in Figure 1, even if the system model and the input are fixed, different internal states may produce the same observed output trajectory. In that case, the observer cannot tell those states apart from data alone. Formally, we define the distinguishability:

Definition 2.1 (Distinguishable State)**.** Given a system $\mathbf{\Sigma}$, two states $x_i \neq x_j$ are distinguishable under the prescribed input, if there exists $t \in [0, T]$ such that $y(t, x_i, u(t)) \neq y(t, x_j, u(t))$, where $y(\cdot, x', \cdot)$ is the output trajectory of the system for a given input $u(t)$ and an initial state $x(t_0) = x'$.

Following (Vidyasagar, 2002), we study local structural identifiability as local observability of the augmented system in which the parameters are treated as constant states.

Definition 2.2 (Local Observability and Structural Identifiability)**.** $\mathbf{\Sigma}$ is locally observable at x_0 if there exists a neighbourhood $\mathcal{N}$ of x_0 such that every other state $x \in \mathcal{N}$, $x \neq x_0$ is distinguishable. Further, $\mathbf{\Sigma}$ is locally structurally identifiable at (x_0, p) if it becomes locally observable after augmenting the parameter vector p as additional state variables satisfying $\dot{p} = 0$.

For nonlinear systems, observability and structural identifiability can be investigated through successive Lie derivatives of the output along the system dynamics (Hermann & Krener, 1977; Vidyasagar, 2002), revealing how latent physiological states and model parameters are reflected in the observed trajectory $y(t)$. Known exogenous inputs are treated as fixed analytic functions when computing Lie derivatives of the augmented system. Formally, the Lie derivative is defined as follows:

Definition 2.3 (Lie Derivative)**.** Let $x \in \mathbb{R}^n$, $\phi : \mathbb{R}^n \to \mathbb{R}^m$ and $f : \mathbb{R}^n \to \mathbb{R}^n$. We define the k-th order Lie derivative of ϕ with respect to f by the following recursion:

$$L_f^0 \phi(x) = \phi(x), \quad L_f^k \phi(x) = \frac{\partial (L_f^{k-1} \phi)}{\partial x} \big[f(x) \big] \in \mathbb{R}^m \quad \text{for} \quad k \geq 1.$$

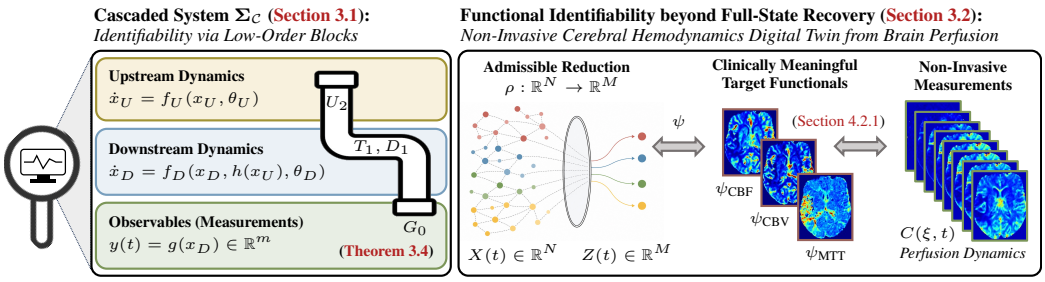


Figure 2: **Cascaded digital twin overview. (a) Cascaded System**. Identifiability in non-invasive digital twins arises through structured propagation across the cascade, captured by low-order blocks $(G_0, T1, D_1, U_2)$. **(b) Functional Identifiability**. Admissible reduction further reduces latent representations and recovers identifiable, clinically meaningful functionals beyond full-state recovery.

3 Structural and Functional Identifiability of Clinical Targets for Non-Invasive Medical Digital Twins

We start from the structural characterization of local observability and structural identifiability in Section 3.1, and extend the analysis to reduced representations and functional identifiability of clinical targets and summaries in Section 3.2.

3.1 Structural Identifiability under a Cascaded Dynamical System.

In the non-invasive regime, latent physiological states are not observed directly, but evolve through upstream physiological dynamics and affect measurements only through downstream observations. This induces a natural *cascade: an upstream subsystem encoding latent physiology, followed by a downstream subsystem encoding acquisition or measurement formation.*

As illustrated in Figure 2 (left), different components of the augmented state become visible at different derivative orders. Downstream states appear directly in the output, upstream states appear only after first-order propagation across the cascade, and upstream parameters become visible only through higher-order propagation. This structure is what makes a low-order identifiability characterization possible, and induces a structured system with upstream latent dynamics and downstream measurement dynamics[1]:

Definition 3.1 (Cascaded Dynamical System)**.** Consider an upstream subsystem with state $x_U \in \mathbb{R}^{n_U}$ and a downstream subsystem with state $x_D \in \mathbb{R}^{n_D}$, parametrized by $\theta_U \in \mathbb{R}^{q_U}$ and $\theta_D \in \mathbb{R}^{q_D}$, respectively. A cascaded dynamical system $\mathbf{\Sigma}_\mathcal{C}$ is governed by the ODE system:

$$\dot{x}_U = f_U(x_U, \theta_U), \quad \dot{x}_D = f_D(x_D, h(x_U), \theta_D), \quad y(t) = g(x_D) \in \mathbb{R}^m,$$

for all $t \in [0, T]$, where y is an observable vector, h is a coupling map transmitting the upstream state x_U to the downstream dynamics, and f_U, f_D, h, g are analytic vector-valued maps.

Definition 3.2 (Observability-Identifiability Matrix of Cascaded Dynamical System)**.** Given $\mathbf{\Sigma}_\mathcal{C}$, let[a] $X := (x_D^\top, x_U^\top, \theta_D^\top, \theta_U^\top)^\top \in \mathbb{R}^N$, $F(X,t) := (f_D^\top, f_U^\top, 0_{q_D}^\top, 0_{q_U}^\top)^\top \in \mathbb{R}^N$, $N := n_U + n_D + q_U + q_D$. We define the observability-identifiability matrix: $O_I(X) = \left[\frac{\partial L_F^k y}{\partial X}(X)\right]_{k=0}^{N-1}$.

[a] We order the state X in this way so that the directly observed downstream block appears first in $O_I(X)$.

[1] For notational simplicity, we suppress known exogenous inputs, as known analytic inputs can be incorporated symbolically in the Lie derivative computation (Villaverde et al., 2016; Anguelova, 2004). Accordingly, the resulting rank conditions are evaluated conditional on the prescribed input trajectory. A failure of the low-order conditions may therefore reflect insufficient excitation under that input rather than intrinsic non-identifiability.

Definition 3.3 (Low-Order Blocks of Observability-Identifiability Matrix)**.** For a given point $X_0 := X(t_0)$ at some $t_0 > 0$, we define low-order blocks of the observability-identifiability:

$$G_0(X_0) := \frac{\partial(L_F^0 y)}{\partial x_D}(X_0) \in \mathbb{R}^{m \times n_D}, \quad T_1(X_0) := \frac{\partial(L_F^1 y)}{\partial x_U}(X_0) \in \mathbb{R}^{m \times n_U},$$
$$D_1(X_0) := \frac{\partial(L_F^1 y)}{\partial \theta_D}(X_0) \in \mathbb{R}^{m \times q_D}, \quad U_2(X_0) := \frac{\partial(L_F^2 y)}{\partial \theta_U}(X_0) \in \mathbb{R}^{m \times q_U}.$$

The low-order blocks in Definition 3.3 have a direct structural meaning. G_0 tests direct visibility of downstream states from the observable; T_1 and D_1 test the first-order appearance of upstream states and downstream parameters; and U_2 tests the second-order appearance of upstream parameters after propagation through the upstream dynamics and across the cascade. Thus, identifiability in cascaded systems appears progressively: downstream quantities are visible immediately, while upstream quantities require higher-order propagation. (See numerical experiments in Section 4).

Theorem 3.4 (Observability and Identifiability of Cascaded Dynamical Systems)**.** Consider the cascaded dynamical system $\mathbf{\Sigma}_\mathcal{C}$ in Definition 3.1, and let G_0, T_1, D_1 and U_2 be low-order blocks defined in Definition 3.3. If $m \geq \max\{n_D, n_U + q_D, q_U\}$, and the following rank conditions hold:

$$\operatorname{rank}(G_0(X_0)) = n_D, \quad \operatorname{rank}\begin{pmatrix} T_1(X_0) & D_1(X_0) \end{pmatrix} = n_U + q_D, \quad \operatorname{rank}(U_2(X_0)) = q_U, \quad (3.1)$$

then, $\mathbf{\Sigma}_\mathcal{C}$ with the augmented state X is locally observable and identifiable around X_0. Please refer to Appendix D.1 for a detailed proof.

Theorem 3.4 provides a sufficient low-order criterion for local observability and structural identifiability without constructing the full observability-identifiability matrix. This motivates an order-adaptive analysis that terminates when the low-order conditions certify identifiability and otherwise proceeds to higher-order Lie derivatives. Appendix F.2 shows that this staged procedure preserves higher-order recovery while yielding a $14.8\times$ overall speedup over always constructing the higher-order reference.

3.2 Functional Identifiability of Clinical Targets beyond Full State Recovery

In many realistic settings, the latent physiological model contains substantially more states and parameters than can be supported by the available observations. The identifiability of the full augmented state in Section 3.1 is often too strong a target, especially for high-dimensional medical digital twins. As a result, *full-state recovery may be structurally impossible or unnecessarily complex, even when lower-dimensional or clinically meaningful summaries remain identifiable*. This motivates a reduced perspective where the inferential object is not necessarily the full latent twin, but an effective physiological quantity, summary index, or clinically relevant target functional (Figure 2(right)).

Let $X := (x_D^\top, x_U^\top, \theta_D^\top, \theta_U^\top)^\top \in \mathbb{R}^N$ denote the augmented state of a cascaded dynamical system $\mathbf{\Sigma}_\mathcal{C}$ (Definition 3.1). We consider a reduced augmented state $Z := (z_D, z_U, \alpha_D, \alpha_U) \in \mathbb{R}^M$, where $z_D \in \mathbb{R}^{m_D}, z_U \in \mathbb{R}^{m_U}, \alpha_D \in \mathbb{R}^{p_D}, \alpha_U \in \mathbb{R}^{p_U}$ and $M = m_D + m_U + p_D + p_U$. Here, z_D, z_U are reduced downstream and upstream states. α_D, α_U are reduced downstream and upstream parameters.

Definition 3.5 (Admissible Reduction)**.** Let $\mathbf{\Sigma}_\mathcal{C}$ be a cascaded dynamical system with augmented state $X(t) \in \mathbb{R}^N$. A map $\rho : \mathbb{R}^N \to \mathbb{R}^M$ is an admissible reduction if there exist maps $f_{U,\mathrm{red}}, f_{D,\mathrm{red}}, h_{\mathrm{red}}, \gamma$ and a fixed observation readout $\chi_{\mathrm{red}} : y \to \mathbb{R}^r$, such that for every trajectory in $\mathbf{\Sigma}_\mathcal{C}$ and for all $t \in [0, T]$, the reduced augmented state $Z(t) = \rho(X)$ has components satisfying:

$$\dot{z}_U(t) = f_{U,\mathrm{red}}(z_U, \alpha_U), \quad \dot{z}_D(t) = f_{D,\mathrm{red}}(z_D, h_{\mathrm{red}}(z_U), \alpha_D), \quad \chi_{\mathrm{red}}(y) = \gamma(z_D).$$

We denote by F_{red} the corresponding augmented vector field on Z.

The maps $\rho(\cdot)$ and $\chi_{\mathrm{red}}(\cdot)$ reduce the original twin to a lower-dimensional twin with the same cascaded architecture and extract the reduced observation from raw measurements. When a high-resolution twin is overparameterized relative to the available measurements, an admissible reduction therefore converts a fine but poorly posed representation into a coarser, potentially identifiable one, while reducing the identifiability analysis from dimension N to M.

Definition 3.6 (Compatibility of Reduced Observable)**.** For a given admissible reduction $\rho(\cdot)$, the reduced readout $\gamma(\cdot)$ is ρ-compatible, if, for any two augmented states $X, \widetilde{X} \in \mathbb{R}^N$,

$$\rho(X) = \rho(\widetilde{X}) \quad \Longrightarrow \quad \chi_{\mathrm{red}}(g(x_D)) = \chi_{\mathrm{red}}(g(\widetilde{x}_D)).$$

In practice, the inferential target is often not the entire reduced state, but a coarser yet clinically meaningful quantity. For example, a regional flow summary, an aggregate transport index, or an effective parameter combination is sufficient for downstream interpretation. This motivates formulating observability and structural identifiability directly at the level of a target functional $\psi(\cdot)$. Rather than insisting on recovery of every latent degree of freedom, functional identifiability asks whether the quantity of interest is uniquely determined by the observation process.

Definition 3.7 (Functional Observability and Identifiability)**.** Let $\psi : \mathbb{R}^M \to \mathbb{R}^d$ be a target functional on the reduced augmented state Z. A cascaded system obtained from an admissible reduction is locally ψ-observable and locally ψ-identifiable at $Z_0 := Z(t_0)$, if there exists a neighborhood $\mathcal{U}$ of Z_0, such that for any $Z, \tilde{Z} \in \mathcal{U}$, $\gamma(t; Z) = \gamma(t; \tilde{Z})$ implies $\psi(Z) = \psi(\tilde{Z})$.

Remark 3.8. Definition 3.7 weakens the full-state identifiability by requiring only that the target functional $\psi(\cdot)$ be uniquely determined. Thus, even when the full latent state remains non-identifiable, lower-dimensional or clinically meaningful summaries may remain identifiable.

Theorem 3.9 (ψ-Identifiability for Reduced Cascaded System)**.** Let $\mathbf{\Sigma}_{\mathcal{R}}$ be a cascaded system obtained from an admissible reduction $\rho(\cdot)$, with reduced augmented state $Z(t) \in \mathbb{R}^M$, reduced readout $\gamma(\cdot) \in \mathbb{R}^r$, and reduced augmented vector field F_{red}. Let $\psi : \mathbb{R}^M \to \mathbb{R}^d$ be a target functional on $Z(t)$. Suppose that the reduced readout $\gamma(\cdot)$ is ρ-compatible and, in a neighborhood of $Z_0 := Z(t_0)$, $\mathrm{row}(D\psi(Z)) \subseteq \mathrm{row}(D\mathcal{J}(Z))$, where $\mathcal{J}(Z) := \left((L^0_{F_{\mathrm{red}}}\gamma(Z))^\top \ (L^1_{F_{\mathrm{red}}}\gamma(Z))^\top \ (L^2_{F_{\mathrm{red}}}\gamma(Z))^\top\right)^\top \in \mathbb{R}^{3r}$ has local constant rank. Then, $\mathbf{\Sigma}_{\mathcal{R}}$ is locally ψ-observable and ψ-identifiable around Z_0. Appendix D.2 provides a detailed proof.

The condition $\mathrm{row}(D\psi) \subseteq \mathrm{row}(D\mathcal{J})$ means that every infinitesimal direction invisible to the reduced observable is also invisible to the target functional; hence ψ is constant along local observational equivalence classes. Theorem 3.9 shows that full recovery of the reduced augmented state is unnecessary when the target functional ψ is determined by the differential information in the reduced readout. In medical digital twins, clinically meaningful targets may remain identifiable even when fine-scale latent states are not. Therefore, admissible reduction provides a principled route from overparameterized twins to identifiable physiological targets better matched to the available data.

4 Experiments

In controlled synthetic cascades (Section 4.1), we test whether the low-order blocks localize failures of certifiable identifiability and reduce the observability-identifiability construction cost. In a cerebral perfusion digital twin (Section 4.2), we further test whether full voxelwise recovery is unsupported while clinically meaningful regional functionals remain stable under admissible reduction.

4.1 Controlled Synthetic Cascade Systems

We examine the diagnostics of Theorem 3.4 on controlled smooth cascade systems (linear and trigonometric) with known upstream-downstream structure. (1) When observables depend on the upstream subsystem only through a downstream bottleneck, distinct upstream trajectories can yield nearly identical downstream observations; low-order blocks should reveal when upstream-sensitive directions are not certified as identifiable at the corresponding derivative orders. (2) Constructing only the low-order blocks should be substantially cheaper than forming the full observability-identifiability matrix, especially for nonlinear cascades. Full model construction and fitting details are in Appendix E.1.

Latent Ambiguity and Block Diagnostics. We first test whether the sampled low-order diagnostics reveal unresolved upstream-sensitive directions from a downstream scalar readout. For both cascades, multi-initialization fitting yields near-optimal minimizers with nearly indistinguishable downstream

Table 1: Low-order block columns indicate whether the stacked sampled block reaches its target rank. G_0 is full rank while $[T_1\ D_1]$ and U_2 are rank deficient, indicating that downstream states are visible but upstream-sensitive directions are not certified despite small downstream reconstruction errors.

Family	# minimizers	G_0	$[T_1\ D_1]$	U_2	mean error $\pm$ std	max error
Linear	5	✓	×	×	$3.30 \times 10^{-6} \pm 3.40 \times 10^{-6}$	1.28×10^{-5}
Trigonometric	6	✓	×	×	$5.13 \times 10^{-4} \pm 4.07 \times 10^{-4}$	1.86×10^{-3}

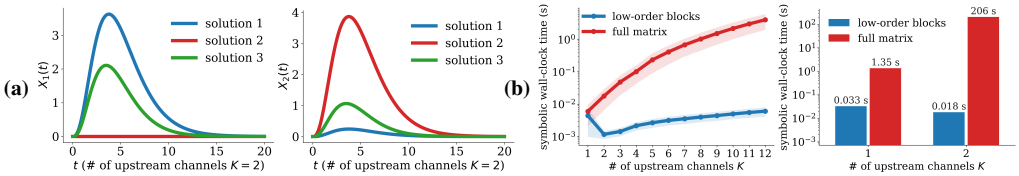


Figure 3: **Experiments on synthetic cascades.** **(a)** Near-optimal minimizers produce indistinguishable downstream observations but distinct upstream trajectories. **(b)** Low-order blocks yield substantially lower wall-clock cost than reconstructing the full observability-identifiability matrix.

trajectories but visibly different upstream trajectories, as shown in Figure 3(a). Table 1 connects this ambiguity to the proposed diagnostics: G_0 reaches its target rank, whereas $[T_1\ D_1]$ and U_2 are rank deficient, indicating that low-order propagation does not certify all upstream-sensitive directions. These rank patterns localize which directions are certifiably resolved at low order and which require higher-order analysis, as further examined in Appendix F.2. The small downstream reconstruction errors further confirm that this ambiguity persists despite nearly identical observations.

Runtime Scaling. We next compare the wall-clock cost of our low-order construction against the full observability-identifiability matrix. Figure 3(b) shows that the low-order construction remains substantially cheaper as the number of upstream channels (K) increases. The gap is already visible for the linear cascade and becomes much larger for the trigonometric cascade, where repeated symbolic differentiation through nonlinear terms causes rapid expression growth in the full matrix (> 200s for $K = 2$, versus only 0.018s for the low-order construction). In contrast, the cascaded diagnostic only constructs the blocks needed to probe the first appearances of downstream states, upstream states, downstream parameters, and upstream parameters. Thus, the scaling trend in Figure 3(b) supports the practical motivation for Theorem 3.4: constructing the selected low-order blocks reduces computational cost, with higher-order analysis reserved for cases where these blocks do not certify identifiability. Appendix F.1 extends this comparison to separable nonlinear, pairwise polynomial, and mixed couplings, showing that the computational advantage persists beyond additive coupling.

Input Dependence. In Appendix F.4, we further compare prescribed input profiles with known-input derivatives explicitly included, showing that the analysis remains applicable under both constant and time-varying inputs. The resulting diagnostics capture input-dependent changes in numerical certification while retaining the structural ambiguity predicted by the exact aggregate dynamics.

4.2 Cerebral Hemodynamics Digital Twin from Brain Perfusion

We next study cerebral perfusion imaging as a real-data medical digital twin problem. In perfusion magnetic resonance imaging (MRI), an injected tracer is transported by blood flow, producing a 4D tracer concentration time series (Fieselmann et al., 2011). The measurements are tracer curves, while the target functionals are latent hemodynamic summaries such as blood flow and blood volume (Kermode et al., 1990; Desai et al., 2007). This naturally forms a cascade: upstream hemodynamics drive downstream tracer transport, which generates non-invasive image observations. We use this setting to test our functional-identifiability claim: voxelwise full-state recovery may be unsupported, while clinically meaningful regional functionals remain stable after admissible reduction. Experiments use ISLES2017 (Winzeck et al., 2018), a public acute ischemic stroke dataset with 4D perfusion MRI from 44 patients. Dataset and processing details are in Appendix E.2.

4.2.1 HEMODYNAMICS DIGITAL TWIN CONSTRUCTION

We model cerebral tracer transport as a reduced cascade: upstream layer captures effective hemodynamic exchange; downstream layer maps this exchange to measured tracer concentration signal.

Upstream Hemodynamics with Windkessel Closure. High-dimensional cerebrovascular models require boundary conditions and terminal vascular assumptions that are not available from routine perfusion MRI. We therefore use a compartmental upstream model (Olufsen et al., 2002) over a partition $\{\Omega_r\}_{r=1}^{R}$ of the brain domain $\Omega \subset \mathbb{R}^3$. Each region is treated as an effective hemodynamic compartment whose exchange is summarized by inflow, storage, and outflow. Using a two-element Windkessel closure (Westerhof et al., 2009; Kuang et al., 2024), the upstream outflow state evolves as

$$\frac{\mathrm{d}}{\mathrm{d}t} z^r_{U,\mathrm{out}}(t) = \frac{\alpha^r_{U,F}(t) - z^r_{U,\mathrm{out}}(t)}{\alpha^r_{U,\tau}}, \quad z^r_{U,\mathrm{out}}(0) = \alpha^r_{U,0} \quad \text{for} \quad r = 1, \ldots, R. \tag{4.1}$$

Here, $\alpha^r_{U,F}(t)$ denotes the effective regional inflow, $z^r_{U,\mathrm{out}}(t)$ denotes the reduced upstream outflow state, and $\alpha^r_{U,0}$ is a nuisance initialization parameter. The parameter $\alpha^r_{U,\tau} > 0$ refers to the relaxation time, which is the product of effective vascular compliance and resistance (Korte et al., 2024).

Downstream Tracer Dynamics via Advection. Given the upstream hemodynamic response, the downstream layer maps blood-flow-driven transport to tracer concentration curves. At the field level, tracer concentration $C(\xi, t)$ evolves according to advection (Liu et al., 2021; Lee & Liu, 2026),

$$\frac{\partial C(\xi,t)}{\partial t} + \nabla_\xi \cdot (v(\xi,t) C(\xi,t)) = 0, \quad \xi \in \Omega. \tag{4.2}$$

Here, the divergence term represents the net advective flux of tracer induced by blood motion. Mechanistically, local tracer concentration increases when this flux enters a tissue neighborhood and decreases when it leaves. This equation describes how the latent hemodynamic drive affects the observed MRI signal only after being propagated through tracer transport. For the reduced regional cascade, we integrate (4.2) over each tissue support Ω_r and represent the boundary fluxes by effective upstream inflow and outflow terms, yielding the following tracer transport ODE,

$$\frac{\mathrm{d}}{\mathrm{d}t}(\alpha^r_{D,V} \bar{z}^r_D(t)) = \alpha^r_{U,F}(t) z^r_{D,\mathrm{in}}(t) - z^r_{U,\mathrm{out}}(t) \bar{z}^r_D(t), \quad r = 1, \ldots, R, \tag{4.3}$$

where $\bar{z}^r_D(t)$ is the regional mean downstream tracer concentration over Ω_r, $z^r_{D,\mathrm{in}}(t)$ is a prescribed exogenous tracer input, and $\alpha^r_{D,V}$ is the downstream vascular volume scale. We defer the detailed descriptions of the setup (4.1) in Appendix C.2, and derivation from (4.2) to (4.3) to Appendix C.3.

Observation Scale and Reduced Cascade. Given a partition, the field-level tracer dynamics in (4.2) and its regional compartmental form in (4.3) define a reduced cascade with state $Z(t)$. Spatial averaging over each tissue support instantiates the admissible reduction in Definition 3.5: voxelwise fields are replaced by regional exchange coordinates while preserving the upstream-downstream cascade. The reduced readout returns the regional tracer curves measured by perfusion MRI. We fit the reduced twin by matching predicted and observed regional tracer trajectories using a weighted normalized reconstruction discrepancy with weak log-scale priors on inflow amplitude, effective volume, flow scale, and relaxation parameters. Full implementation details are given in Appendix E.3.

Target Functionals. Cerebral blood flow (CBF), cerebral blood volume (CBV), and mean transit time (MTT) summarize flow delivery, effective blood volume, and tracer transit, and are commonly used in cerebral perfusion assessment (Kermode et al., 1990; Desai et al., 2007). Here, we test their functional identifiability and evaluate whether these clinically meaningful summaries remain identifiable from reduced tracer observations even when voxelwise full-state recovery is unsupported. Specifically, we denote the target functional $\psi := (\psi_{\mathrm{CBF}}(Z), \psi_{\mathrm{CBV}}(Z), \psi_{\mathrm{MTT}}(Z))$, where $\psi_{\mathrm{MTT}}(Z) = 60 \cdot \psi_{\mathrm{CBV}}(Z) \,/\, \psi_{\mathrm{CBF}}(Z)$ (Zierler, 2000). Precise definitions and computations are in Appendix C.1.

4.2.2 FUNCTIONAL IDENTIFIABILITY DIAGNOSTICS

At the voxel scale, full latent recovery is underdetermined: each voxel has flow, volume, input, and response variables but only a tracer-concentration trajectory is observed. We assess clinical functional stability across observationally similar dynamics, using retained-minimizer stability as empirical evidence and the row-space residual to probe the sufficient differential condition in Theorem 3.9.

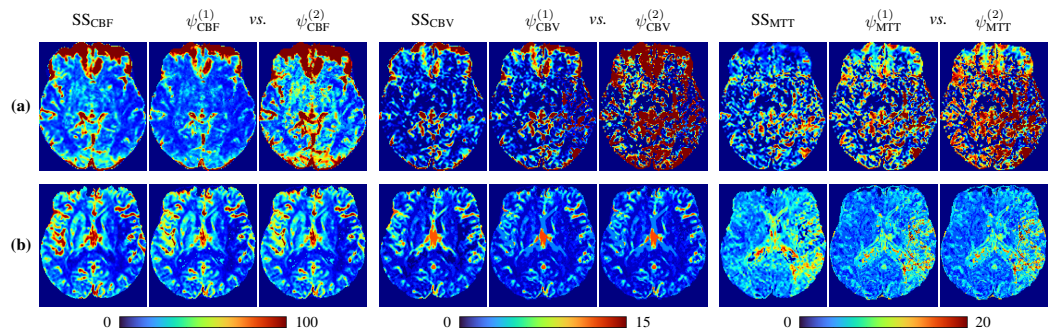


Figure 4: **Functional identifiability diagnostics for cerebral hemodynamics. (a)** Voxelwise fitting remains ambiguous: retained solutions match tracer reconstructions but yield different CBF, CBV, and MTT maps. **(b)** Admissible regional reduction produces much more stable maps across retained minimizers. Silver-standard references (SS) are ISLES-provided estimates used only for visualization.

Table 2: **Quantitative functional-identifiability diagnostics across all patients in ISLES2017.**

Functional	Mean functional diff. (%)	Max functional diff. (%)	Row-space residual (%)
ψ_{CBF}	0.133 ± 0.095	0.185 ± 0.133	$2.45 \times 10^{-6} \pm 1.6 \times 10^{-5}$
ψ_{CBV}	0.162 ± 0.244	0.236 ± 0.358	$3.89 \times 10^{-10} \pm 9.1 \times 10^{-10}$
ψ_{MTT}	0.604 ± 1.008	0.975 ± 1.700	$5.41 \times 10^{-6} \pm 3.5 \times 10^{-5}$

Empirical Functional Discrepancy. Multi-initialization yields retained solutions $Z^{(i)}$ with near-equivalent regional reconstructions and hence near-equivalent observables $\gamma(z_D^{(i)})$. Identifiability predicts agreement of ψ; for retained solutions i and j, we measure the normalized discrepancy

$$d_{i,j}(\psi_k) := \|\psi_k(Z^{(i)}) - \psi_k(Z^{(j)})\|_2 \,/\, (\|\psi_k(Z^{(i)})\|_2 + 10^{-12}) \text{ for } k \in \{\text{CBF}, \text{CBV}, \text{MTT}\}. \quad (4.4)$$

Figure 4(a) shows voxelwise retained solutions with matching tracer reconstructions but discrepant CBF, CBV, and MTT maps. Admissible regional reduction in Figure 4(b) yields stable maps matching ISLES2017's silver-standard references (visualization-only estimates, not ground truth). Reduction shifts inference from an unsupported fine-scale field to summaries better matched to observations.

Row-Space Residual. As a numerical diagnostic of the clinical fits, let $W_1 := DJ(Z)$ and $W_2 := D\psi_k(Z)$ for each retained solution and target ψ_k[2]. Row inclusion requires the rows of W_2 to lie in the row space of W_1. Let $\widehat{W}_2^\top$ be the least-squares projection of $W_2^\top$ onto $\text{col}(W_1^\top)$. We report

$$\text{Res}(W_1, W_2) := \|W_2^\top - \widehat{W}_2^\top\|_F \,/\, (\|W_2^\top\|_F + 10^{-12}), \quad (4.5)$$

Table 2 reports small row-space residuals and retained-minimizer discrepancies, supporting stable regional CBF, CBV, and MTT inference. Appendix F.3 provides further evidence of numerical robustness across SVD thresholds and physiologically motivated boundary subsets in all 44 patients.

5 Conclusion

We studied a central question for non-invasive medical digital twins: which patient-specific physiological quantities are supported by routine clinical observations? Because latent physiology is observed only indirectly, reproducing measured data does not necessarily imply a unique physiological interpretation. We formulate this problem through cascaded dynamics, with low-order diagnostics tracing observability and structural identifiability through the observation pathway. Admissible reduction and functional identifiability further assess when clinically meaningful summaries remain recoverable despite fine-scale ambiguity. Synthetic cascades and cerebral hemodynamics digital twin experiments illustrate that fine-scale recovery may remain ambiguous, while reduced regional functionals can remain stable across observationally similar solutions. These findings motivate matching model resolution and inferential targets to the information available in non-invasive measurements.

[2]Here local Jacobians are stacked over time, assuming static flow/volume readouts with fixed input amplitude.

REFERENCES

Milena Anguelova. *Nonlinear Observability and Identifiability: General Theory and a Case Study of a Kinetic Model for S. cerevisiae*. Chalmers Tekniska Hogskola (Sweden), 2004.

Mark Asch. *A toolbox for digital twins: from model-based to data-driven*. SIAM, 2022.

Jean-Pierre Barbot, Driss Boutat, and Thierry Floquet. An observation algorithm for nonlinear systems with unknown inputs. *Automatica*, 45(8):1970–1974, 2009.

R. Bellman and Karl Johan Åström. On structural identifiability. *Mathematical biosciences*, 7(3-4): 329–339, 1970.

Richard N Bergman, Y Ziya Ider, Charles R Bowden, and Claudio Cobelli. Quantitative estimation of insulin sensitivity. *American Journal of Physiology-Endocrinology And Metabolism*, 236(6): E667, 1979.

Jennifer K Briggs, JN Stroh, Soojin Park, Brandon Foreman, Michael M Tymko, Jay Carr, Melike Sirlanci, Philip N Ainslie, Richard KP Benninger, Tellen D Bennett, et al. Towards a physics informed digital twin to predict cerebral blood flow and cerebral vascular regulation. *NPJ Digital Medicine*, 2026.

Stanislas Demuth, Jérôme De Sèze, Gilles Edan, Tjalf Ziemssen, Françoise Simon, and Pierre-Antoine Gourraud. Digital representation of patients as medical digital twins: Data-centric viewpoint. *JMIR Medical Informatics*, 13(1):e53542, 2025.

Brinda S Desai, Angela J Monahan, Paul M Carvey, and Bill Hendey. Blood–brain barrier pathology in Alzheimer's and Parkinson's disease: implications for drug therapy. *Cell transplantation*, 16(3): 285–299, 2007.

Andreas Fieselmann, Markus Kowarschik, Arundhuti Ganguly, Joachim Hornegger, and Rebecca Fahrig. Deconvolution-based CT and MR brain perfusion measurement: theoretical model revisited and practical implementation details. *International Journal of Biomedical Imaging*, 2011(1): 467563, 2011.

Bruce Fischl. FreeSurfer. *Neuroimage*, 2012.

Matthew F Glasser, Timothy S Coalson, Emma C Robinson, Carl D Hacker, John Harwell, Essa Yacoub, Kamil Ugurbil, Jesper Andersson, Christian F Beckmann, Mark Jenkinson, et al. A multi-modal parcellation of human cerebral cortex. *Nature*, 536(7615):171–178, 2016.

Tim Heinemann and Andreas Raue. Model calibration and uncertainty analysis in signaling networks. *Current opinion in biotechnology*, 39:143–149, 2016.

Robert Hermann and Arthur Krener. Nonlinear controllability and observability. *IEEE Transactions on automatic control*, 22(5):728–740, 1977.

David A Kass, Toji Yamazaki, Daniel Burkhoff, W Lowell Maughan, and Kiichi Sagawa. Determination of left ventricular end-systolic pressure-volume relationships by the conductance (volume) catheter technique. *Circulation*, 73(3):586–595, 1986.

AG Kermode, AJ Thompson, P Tofts, DG MacManus, BE Kendall, DPE Kingsley, IF Moseley, P Rudge, and WI McDonald. Breakdown of the blood-brain barrier precedes symptoms and other MRI signs of new lesions in multiple sclerosis: pathogenetic and clinical implications. *Brain*, 113 (5):1477–1489, 1990.

Bas M Kessels, Rob HB Fey, and Nathan van de Wouw. Real-time parameter updating for nonlinear digital twins using inverse mapping models and transient-based features. *Nonlinear dynamics*, 111 (11):10255–10285, 2023.

Jana Korte, Ehlar Sophie Klopp, and Philipp Berg. Multi-Dimensional Modeling of Cerebral Hemodynamics: A Systematic Review. *Bioengineering*, 11(1):72, 2024.

Keying Kuang, Frances Dean, Jack B Jedlicki, David Ouyang, Anthony Philippakis, David Sontag, and Ahmed M Alaa. Med-real2sim: Non-invasive medical digital twins using physics-informed self-supervised learning. *Advances in Neural Information Processing Systems*, 37:5757–5788, 2024.

Reinhard Laubenbacher, James P Sluka, and James A Glazier. Using digital twins in viral infection. *Science*, 371(6534):1105–1106, 2021.

Reinhard Laubenbacher, Borna Mehrad, Ilya Shmulevich, and Natalia Trayanova. Digital twins in medicine. *Nature computational science*, 4(3):184–191, 2024.

Yi-Chen Lee and Peirong Liu. HemoPIC: A physics-informed cerebral hemodynamics digital twin for brain perfusion. *arXiv preprint arXiv:2607.08799*, 2026.

Genyuan Li and Herschel Rabitz. A general analysis of exact lumping in chemical kinetics. *Chemical engineering science*, 44(6):1413–1430, 1989.

Peirong Liu, Yueh Z Lee, Stephen R Aylward, and Marc Niethammer. Perfusion imaging: an advection diffusion approach. *IEEE Transactions on Medical Imaging*, 40(12):3424–3435, 2021.

YW Lui, ER Tang, AM Allmendinger, and V Spektor. Evaluation of ct perfusion in the setting of cerebral ischemia: patterns and pitfalls. *American journal of neuroradiology*, 31(9):1552–1563, 2010.

Agostino Martinelli. Nonlinear unknown input observability: Extension of the observability rank condition. *IEEE Transactions on Automatic Control*, 64(1):222–237, 2018.

Philippe Moingeon, Marylore Chenel, Cécile Rousseau, Emmanuelle Voisin, and Mickael Guedj. Virtual patients, digital twins and causal disease models: Paving the ground for in silico clinical trials. *Drug discovery today*, 28(7):103605, 2023.

Susan A Murphy and Aad W Van der Vaart. On profile likelihood. *Journal of the American Statistical Association*, 95(450):449–465, 2000.

Ivan A Nestorov, Leon J Aarons, Philip A Arundel, and Malcolm Rowland. Lumping of whole-body physiologically based pharmacokinetic models. *Journal of pharmacokinetics and biopharmaceutics*, 26(1):21–46, 1998.

Mette S Olufsen, Ali Nadim, and Lewis A Lipsitz. Dynamics of cerebral blood flow regulation explained using a lumped parameter model. *American Journal of Physiology-Regulatory, Integrative and Comparative Physiology*, 282(2):R611–R622, 2002.

Leif Østergaard. Principles of cerebral perfusion imaging by bolus tracking. *Journal of Magnetic Resonance Imaging: An official journal of the international society for magnetic resonance in medicine*, 22(6):710–717, 2005.

Leif Østergaard, Robert M Weisskoff, David A Chesler, Carsten Gyldensted, and Bruce R Rosen. High resolution measurement of cerebral blood flow using intravascular tracer bolus passages. part i: Mathematical approach and statistical analysis. *Magnetic resonance in medicine*, 36(5):715–725, 1996.

Andreas Raue, Clemens Kreutz, Thomas Maiwald, Julie Bachmann, Marcel Schilling, Ursula Klingmüller, and Jens Timmer. Structural and practical identifiability analysis of partially observed dynamical models by exploiting the profile likelihood. *Bioinformatics*, 25(15):1923–1929, 2009.

Mahmoud Shams Falavarjani, Adriana González Vázquez, and Alejandro F Villaverde. Assessment of Computational Tools for Analysing the Observability and Accessibility of Nonlinear Models. *Computation*, 13(12):281, 2025.

Tianze Sun, Xiwang He, and Zhonghai Li. Digital twin in healthcare: Recent updates and challenges. *Digital health*, 9:20552076221149651, 2023.

E Tunali and Tzyh-Jong Tarn. New results for identifiability of nonlinear systems. *IEEE Transactions on Automatic Control*, 32(2):146–154, 1987.

Mathukumalli Vidyasagar. *Nonlinear systems analysis*. SIAM, 2002.

Alejandro F Villaverde. Observability and structural identifiability of nonlinear biological systems. *Complexity*, 2019(1):8497093, 2019.

Alejandro F Villaverde, Antonio Barreiro, and Antonis Papachristodoulou. Structural identifiability of dynamic systems biology models. *PLoS computational biology*, 12(10):e1005153, 2016.

Alejandro F Villaverde, Nikolaos Tsiantis, and Julio R Banga. Full observability and estimation of unknown inputs, states and parameters of nonlinear biological models. *Journal of the Royal Society Interface*, 16(156), 2019.

Danhong Wang, Randy L Buckner, Michael D Fox, Daphne J Holt, Avram J Holmes, Sophia Stoecklein, Georg Langs, Ruiqi Pan, Tianyi Qian, Kuncheng Li, et al. Parcellating cortical functional networks in individuals. *Nature neuroscience*, 18(12):1853–1860, 2015.

Nico Westerhof, Jan-Willem Lankhaar, and Berend E Westerhof. The arterial windkessel. *Medical & biological engineering & computing*, 47(2):131–141, 2009.

Franz-Georg Wieland, Adrian L Hauber, Marcus Rosenblatt, Christian Tönsing, and Jens Timmer. On structural and practical identifiability. *Current opinion in systems biology*, 25:60–69, 2021.

Stefan Winzeck, Arsany Hakim, Richard McKinley, José AADSR Pinto, Victor Alves, Carlos Silva, Maxim Pisov, Egor Krivov, Mikhail Belyaev, Miguel Monteiro, et al. ISLES 2016 and 2017-benchmarking ischemic stroke lesion outcome prediction based on multispectral MRI. *Frontiers in neurology*, 9:679, 2018.

Dingling Yao, Caroline Muller, and Francesco Locatello. Marrying causal representation learning with dynamical systems for science. *Advances in Neural Information Processing Systems*, 37: 71705–71736, 2024.

Kenneth Zierler. Indicator dilution methods for measuring blood flow, volume, and other properties of biological systems: a brief history and memoir. *Annals of biomedical engineering*, 28(8):836–848, 2000.

Appendix

A Related Work **14**

B Limitations and Future Work **15**

C Supplementary Backgrounds **15**
C.1 Regional Target Functionals for Cerebral Hemodynamics 15
C.2 From Windkessel Closure to Upstream Outflow Dynamics 16
C.3 From Field Level Advection to Compartmental Tracer Transport 16

D Proofs of Main Text **17**
D.1 Proof of Theorem 3.4 18
D.2 Proof of Theorem 3.9 19

E Experimental Details **19**
E.1 Synthetic Cascade Setup and Diagnostics 19
E.2 Dataset, Preprocessing, and Compute Environment 21
E.3 Tracer Reconstruction for Reduced Cascade 22

F Additional Numerical Results **23**
F.1 Nonlinear-Coupling Scalability 24
F.2 Low-Order Certification and Higher-Order Recovery 24
F.3 Stability Near Degenerate Regimes 25
F.4 Input-Dependent Identifiability 26

Impact Statement

This work is methodological in nature and studies identifiability in non-invasive medical digital twins rather than proposing a deployable clinical decision system. We therefore do not expect immediate negative social impact from the formal results or diagnostic experiments presented in this paper.

AI use statement

We used large language models (LLMs) to assist with writing and editing, including improving clarity, grammar, and conciseness. We also used LLMs for literature retrieval and discovery, such as identifying potentially missing related work. All technical content and results are original contributions of the authors, and the authors take full responsibility for the final content of this work.

Ethics statement

This work is methodological and involves no prospective human-subject recruitment or clinical intervention. The cerebral perfusion experiments use publicly available imaging data for research purposes. The proposed models are not intended for clinical diagnosis or treatment decision-making

Reproducibility statement

Complete proofs and assumptions are provided in Appendix D. Experimental details including architectures, hyperparameters, data processing, and evaluation procedures are detailed in Appendix E. The source code for reproducing the reported experiments is submitted together with the manuscript.

A Related Work

Non-Invasive Medical Digital Twins and Parameter Learning. Medical digital twins have recently emerged as a framework for patient-specific physiological modeling, with applications ranging from virtual patients and in silico trials to personalized clinical decision support (Laubenbacher et al., 2024; Moingeon et al., 2023; Demuth et al., 2025). A closely related line of work studies how to construct such twins from non-invasive measurements. In particular, Kuang et al. (2024) formulate non-invasive digital twin construction as a composite inverse problem, where a known physiological simulator is coupled with an unknown measurement map from latent states to non-invasive observations. Their method learns a differentiable simulator representation through a physics-informed pretext task and then trains a non-invasive-to-parameter map using measurable physiological summaries. This provides an important learning pipeline for patient-specific twin construction from routine data. However, the primary focus is estimation and representation learning, rather than a general structural analysis of when the latent states, physiological parameters, or clinically meaningful functionals are uniquely determined by the non-invasive observation process. Our work complements this direction by treating the observation mechanism as part of a cascaded dynamical system and studying identifiability before estimation is performed.

Identifiability in Dynamical Systems and Representation Learning. Structural identifiability has a long history in dynamical systems, systems biology, and control, where it is commonly analyzed through observability, augmented states, and Lie derivatives (Asch, 2022; Bellman & Åström, 1970; Hermann & Krener, 1977; Raue et al., 2009; Villaverde et al., 2016; Wieland et al., 2021; Tunali & Tarn, 1987). These methods provide the theoretical basis for asking whether latent states and parameters can be recovered from output trajectories, but generic observability-identifiability constructions often require high-order Lie derivatives or full matrix rank tests. Classical formulations also primarily focus on the recoverability of the full augmented state (Villaverde, 2019; Villaverde et al., 2019; Shams Falavarjani et al., 2025). Our setting is motivated by non-invasive medical digital twins, where latent physiology is observed only indirectly through downstream sensing, transport, or imaging processes, and where full-state recovery may be unsupported or unnecessary. By exploiting this cascaded latent-to-observation structure, we derive low-order block conditions aligned with the propagation structure of the cascade, and further introduce admissible reduction together with functional identifiability, which shifts the inferential target from full latent recovery to clinically meaningful physiological summaries. Recent work connects parameter estimation in dynamical systems with causal representation learning (Kessels et al., 2023). For example, Yao et al. (2024) formulate trajectory-specific physical parameters as latent variables and uses causal representation learning to obtain identifiable neural emulators. This perspective is complementary to ours, but it assumes existence, uniqueness, and structural identifiability as underlying conditions, and focuses on importing representation-identifiability results into dynamical system parameter learning. In contrast, our work studies how to verify structural identifiability in non-invasive medical digital twin setting.

Lumping and Clinically Meaningful Model Reduction. Lumping and model reduction have a long history in dynamical and physiological modeling. Classical work on exact lumping studies when lower-dimensional coordinates preserve closed dynamics of the original system (Li & Rabitz, 1989), while physiological applications have used reduced compartmental models to retain interpretable kinetic quantities, including insulin sensitivity in glucose regulation and effective kinetics in pharmacokinetic models (Bergman et al., 1979; Nestorov et al., 1998). Our setting is different in an important respect. Imaging-based digital twins often begin from spatially resolved dynamical measurements rather than a small collection of compartment states. For example, perfusion MRI acquires dynamic tracer measurements over a voxelized field and subsequently derives physiological quantities such as cerebral blood flow from these measurements (Østergaard, 2005; Østergaard et al., 1996). Similarly, functional MRI is commonly organized into anatomically or functionally meaningful regional representations rather than interpreted solely at the native voxel level (Glasser et al., 2016; Wang et al., 2015). Accordingly, reduction in our setting is not used only to decrease state dimension. Instead, it connects high-dimensional spatial dynamics to lower-dimensional physiological coordinates whose scale and interpretation are clinically meaningful. Consequently, our admissible-reduction formulation asks whether such a reduction preserves closed reduced dynamics and whether the resulting coordinates retain sufficient information for identifying the target physiological functionals.

B Limitations and Future Work

Scope of Identifiability Guarantees. This work focuses on identifiability as a prerequisite for reliable non-invasive medical digital twins. Our analysis asks whether latent physiological states, parameters, or reduced functionals are uniquely determined by the available observation process under the assumed dynamical model. We do not directly study downstream tasks such as treatment simulation, intervention optimization, or clinical decision policies. Rather, our goal is to establish the identifiability conditions that such downstream uses would require to be physiologically well posed.

We also remark that our results are local and model-based, relying on analytic dynamical systems and structural observability arguments. Within the current setting, we further examine the numerical stability near degenerate regimes in Appendix F.3. Extending this structural analysis to practical identifiability under finite, noisy, and potentially misspecified clinical data remains an important direction for future work, where established tools such as profile likelihood could be incorporated (Raue et al., 2009; Murphy & Van der Vaart, 2000). Such an extension would require systematically quantifying how observation noise, sampling density, model mismatch, and optimization uncertainty affect identifiable functionals beyond the empirical stability diagnostics used here.

Approximate Reductions in Empirical Models. In the cerebral hemodynamics study, the admissible regional cascade is obtained through spatial averaging and compartmental closure of the field-level tracer dynamics. This reduction is an analytic model approximation rather than an exact equivalence to the full voxelwise dynamics. Consequently, the structural guarantees apply to the reduced cascade under this closure, while approximation error from spatial aggregation may affect how directly these guarantees transfer to the original high-resolution system. Quantifying this reduction error and extending the framework to practical identifiability under approximate reductions are important directions for future work. More broadly, similar analysis may apply to other spatially resolved dynamical measurements, such as perfusion CT (Fieselmann et al., 2011; Lui et al., 2010) or fMRI, when a clinically meaningful reduced representation and compatible reduced dynamics can be defined.

Unknown and Partially Known Inputs. Our analysis treats exogenous inputs as known analytic functions when computing Lie derivatives, but many physiology-informed neural networks and surrogate digital twins rely on inputs that are unknown, partially observed, or learned from data. Extending the cascaded block construction to unknown-input observability (Barbot et al., 2009; Martinelli, 2018) is therefore an important direction for future work. Such inputs could be represented as latent variables, nuisance parameters, or constrained input dynamics within the surrogate model, yielding modified blocks that distinguish identifiable physiological directions from input-induced ambiguities. Likewise, ψ-identifiability could be adapted by requiring $D\psi$ to lie in the observable row space after quotienting out directions attributable to unknown or learned inputs.

C Supplementary Backgrounds

In this section, we provide supplementary background for the cerebral hemodynamics model. Appendix C.1 presents detailed definitions for the regional target functionals. Appendix C.2 derives the upstream Windkessel dynamics. Appendix C.3 derives the downstream tracer-transport dynamics.

C.1 Regional Target Functionals for Cerebral Hemodynamics

Here, we define the target functionals used in the cerebral hemodynamics experiment from the reduced regional state (Section 4.2.1). These functionals correspond to standard perfusion summaries in clinical brain perfusion analysis: cerebral blood flow (CBF), cerebral blood volume (CBV), and mean transit time (MTT). CBF measures the rate of blood delivery to tissue, CBV measures the effective blood volume contained in the tissue compartment, and MTT measures the average time required for tracer to pass through the tissue. In acute stroke imaging, CBF, CBV, and MTT maps are commonly interpreted together to characterize perfusion deficits and tissue at risk (Desai et al., 2007; Kermode et al., 1990). For self-containment, we state the region-wise definitions used in our identifiability diagnostics here; detailed clinical interpretation of CBF, CBV, and MTT in perfusion imaging, together with standard deconvolution-based perfusion modeling and estimation procedures underlying the silver-standard reference in Figure 4, are discussed in (Fieselmann et al., 2011; Winzeck et al., 2018).

For each region Ω_r, the reduced cascade contains an upstream hemodynamic component and a downstream tracer-transport component. The upstream component determines regional blood inflow and outflow, while the downstream component describes the regional tracer concentration trajectory. In our implementation, the effective inflow entering the reduced dynamics is represented by a fitted regional flow amplitude. When the inflow is written as a time-varying function, we parameterize it as

$$\alpha^r_{U,F}(t) = a^r_F \, \bar{u}_F(t),$$

where $\bar{u}_F(t)$ is a fixed normalized temporal profile and $a^r_F > 0$ is the fitted regional inflow amplitude. The CBF target is defined by this regional flow amplitude, since it controls the scale of blood delivery into Ω_r. The CBV target is defined by the downstream vascular volume scale $\alpha^r_{D,V}$, which represents the effective regional tracer distribution volume. MTT is defined through the central volume relation

$$\mathrm{MTT} = 60 \cdot \frac{\mathrm{CBV}}{\mathrm{CBF}},$$

where the factor 60 converts minutes to seconds under conventional perfusion units.

Accordingly, the regional target functionals are

$$\psi^r_{\mathrm{CBF}}(Z) := a^r_F, \qquad \psi^r_{\mathrm{CBV}}(Z) := \alpha^r_{D,V}, \qquad \psi^r_{\mathrm{MTT}}(Z) := 60 \cdot \frac{\psi^r_{\mathrm{CBV}}(Z)}{\psi^r_{\mathrm{CBF}}(Z)}.$$

The full target functional is the region-wise collection

$$\psi(Z) = \big(\psi_{\mathrm{CBF}}(Z), \psi_{\mathrm{CBV}}(Z), \psi_{\mathrm{MTT}}(Z)\big),$$

where each component stacks the corresponding regional quantities over all regions $r = 1, \dots, R$. These definitions make the row-space residual diagnostic explicit: the Jacobians $D\psi_k(Z)$ are computed with respect to the reduced augmented state and test whether the reduced tracer observable contains the differential information needed to identify regional CBF, CBV, and MTT.

C.2 From Windkessel Closure to Upstream Outflow Dynamics

We derive the upstream outflow dynamics (4.1) from the two-element Windkessel closure used in the cerebral hemodynamics cascade (Section 4.2.1). A Windkessel representation provides a compact lumped description of vascular response by separating storage and discharge effects. In each tissue support Ω_r, the compliance component represents vascular storage under pressure variation, while the resistance component represents pressure-driven outflow discharge. This gives the regional closure,

$$\alpha^r_{U,W} \frac{\mathrm{d}}{\mathrm{d}t}\left(P^r(t) - P^r_v\right) = z^r_{U,\mathrm{in}}(t) - z^r_{U,\mathrm{out}}(t), \quad z^r_{U,\mathrm{out}}(t) = \frac{P^r(t) - P^r_v}{\alpha^r_{U,R}},$$

where $P^r(t)$ is an effective microvascular pressure state, P^r_v is a constant downstream reference pressure, $\alpha^r_{U,W} > 0$ is the regional compliance parameter, and $\alpha^r_{U,R} > 0$ is the regional resistance parameter (Korte et al., 2024). For the upstream dynamics used in the main text, the compliance and resistance components enter through their product, $\alpha^r_{U,\tau} = \alpha^r_{U,W} \cdot \alpha^r_{U,R}$. This product defines the characteristic relaxation time of the regional outflow response: larger $\alpha^r_{U,\tau}$ produces slower adaptation of outflow to changes in inflow, while smaller $\alpha^r_{U,\tau}$ produces faster discharge. Thus, the reduced upstream model identifies the effective Windkessel response through the resistance-compliance product rather than separating resistance and compliance individually. Writing $z^r_{U,\mathrm{in}}(t) = \alpha^r_{U,F}(t)$ and $z^r_{U,\mathrm{out}}(0) = \alpha^r_{U,0}$, we obtain the upstream outflow equation (4.1) presented in Section 4.2.1,

$$\frac{\mathrm{d}}{\mathrm{d}t} z^r_{U,\mathrm{out}}(t) = \frac{\alpha^r_{U,F}(t) - z^r_{U,\mathrm{out}}(t)}{\alpha^r_{U,\tau}}, \qquad z^r_{U,\mathrm{out}}(0) = \alpha^r_{U,0}, \qquad r = 1, \dots, R.$$

C.3 From Field Level Advection to Compartmental Tracer Transport

We derive the downstream compartmental tracer equation (4.3) as the tracer component of the admissible reduction used in the cerebral hemodynamics cascade (Section 4.2.1). Given a tissue partition $\{\Omega_r\}_{r=1}^R$ with sufficiently regular boundaries, let $\partial\Omega_r$ denote the boundary of Ω_r, $n_r(\xi)$ denote the outward unit normal vector at $\xi \in \partial\Omega_r$, $\mathrm{d}\xi$ denote the volume measure, and $\mathrm{d}S$ denote the

surface measure on $\partial\Omega_r$. The reduction aggregates the field level tracer concentration and advective fluxes into compartmental coordinates. In our implementation, the downstream state is obtained by spatial averaging:

$$\bar{z}_D^r(t) = \frac{1}{|\Omega_r|}\int_{\Omega_r} C(\xi,t)\,\mathrm{d}\xi, \qquad r = 1,\dots,R.$$

Recall (4.2). At the field level, tracer transport dynamics follows the advection equation:

$$\frac{\partial C(\xi,t)}{\partial t} + \nabla_\xi \cdot (v(\xi,t)C(\xi,t)) = 0, \qquad \xi \in \Omega, \tag{C.1}$$

where $C(\xi,t)$ is the field level tracer concentration and $v(\xi,t)$ is the blood velocity field. Integrating (C.1) over Ω_r and applying the divergence theorem gives

$$\frac{\mathrm{d}}{\mathrm{d}t}\int_{\Omega_r} C(\xi,t)\,\mathrm{d}\xi = -\int_{\partial\Omega_r} (v(\xi,t)\cdot n_r(\xi))\, C(\xi,t)\,\mathrm{d}S. \tag{C.2}$$

The normal velocity $v(\xi,t)\cdot n_r(\xi)$ is positive when blood leaves Ω_r and negative when blood enters Ω_r. Then, we decompose the boundary into inflow and outflow parts according to the sign of $n_r(\cdot)$:

$$\Gamma_{\mathrm{in}}^r(t) = \{\xi \in \partial\Omega_r : v(\xi,t)\cdot n_r(\xi) < 0\}, \quad \Gamma_{\mathrm{out}}^r(t) = \{\xi \in \partial\Omega_r : v(\xi,t)\cdot n_r(\xi) > 0\}.$$

The corresponding aggregated inflow and outflow rates are

$$z_{U,\mathrm{in}}^r(t) = \int_{\Gamma_{\mathrm{in}}^r(t)} -(v(\xi,t)\cdot n_r(\xi))\,\mathrm{d}S, \quad z_{U,\mathrm{out}}^r(t) = \int_{\Gamma_{\mathrm{out}}^r(t)} (v(\xi,t)\cdot n_r(\xi))\,\mathrm{d}S, \tag{C.3}$$

and the associated inflow and outflow tracer concentrations are defined as flux averaged boundary concentrations:

$$z_{D,\mathrm{in}}^r(t) = \frac{\int_{\Gamma_{\mathrm{in}}^r(t)} -(v(\xi,t)\cdot n_r(\xi))\, C(\xi,t)\,\mathrm{d}S}{z_{U,\mathrm{in}}^r(t)}, \quad z_{D,\mathrm{out}}^r(t) = \frac{\int_{\Gamma_{\mathrm{out}}^r(t)} (v(\xi,t)\cdot n_r(\xi))\, C(\xi,t)\,\mathrm{d}S}{z_{U,\mathrm{out}}^r(t)}. \tag{C.4}$$

Substituting (C.3) and (C.4) into (C.2) yields the compartmental tracer balance

$$\frac{\mathrm{d}}{\mathrm{d}t}\int_{\Omega_r} C(\xi,t)\,\mathrm{d}\xi = z_{U,\mathrm{in}}^r(t) z_{D,\mathrm{in}}^r(t) - z_{U,\mathrm{out}}^r(t) z_{D,\mathrm{out}}^r(t).$$

The reduced downstream model represents the regional tracer content by an effective distribution volume multiplying the downstream tracer state:

$$\int_{\Omega_r} C(\xi,t)\,\mathrm{d}\xi \approx \alpha_{D,V}^r \bar{z}_D^r(t).$$

Therefore, the downstream compartmental tracer transport equation is

$$\frac{\mathrm{d}}{\mathrm{d}t}\left(\alpha_{D,V}^r \bar{z}_D^r(t)\right) = z_{U,\mathrm{in}}^r(t) z_{D,\mathrm{in}}^r(t) - z_{U,\mathrm{out}}^r(t) z_{D,\mathrm{out}}^r(t), \qquad r = 1,\dots,R.$$

For the reduced cascade used in the main text, the upstream inflow is represented by the effective flow input and the outflow tracer concentration is represented by the regional downstream tracer state:

$$z_{U,\mathrm{in}}^r(t) = \alpha_{U,F}^r(t), \qquad z_{D,\mathrm{out}}^r(t) = \bar{z}_D^r(t).$$

This gives the downstream tracer ODE as shown in (4.3),

$$\frac{\mathrm{d}}{\mathrm{d}t}\left(\alpha_{D,V}^r \bar{z}_D^r(t)\right) = \alpha_{U,F}^r(t) z_{D,\mathrm{in}}^r(t) - z_{U,\mathrm{out}}^r(t)\bar{z}_D^r(t), \qquad r = 1,\dots,R.$$

D Proofs of Main Text

In this section, we provide the proofs of the sufficient conditions stated in Section 3. Specifically, Appendix D.1 proves that the low-order block conditions imply local observability and structural identifiability for non-invasive medical digital twins formulated as cascaded dynamical systems. Further, Appendix D.2 proves a functional identifiability result for reduced cascaded systems.

D.1 Proof of Theorem 3.4

Before proving the theorem, we revisit a standard rank based criterion for local observability and structural identifiability. This criterion is an established result for general nonlinear dynamical systems, and we state here the corresponding form specialized to our augmented cascaded setting.

Lemma D.1 ((Villaverde et al., 2019; 2016; Hermann & Krener, 1977)). Consider a cascaded dynamical system $\mathbf{\Sigma}_{\mathcal{C}}$ in Definition 3.1. Then, $X := (x_D^\top, x_U^\top, \theta_D^\top, \theta_U^\top)^\top \in \mathbb{R}^N$ (Definition 3.2) is both locally observable and identifiable at X_0 if $\operatorname{rank}(O_I(X_0)) = N$.

We now turn to the low-order structure induced by the cascade and prove Theorem 3.4.

Theorem D.2 (Theorem 3.4 Restated, Observability and Identifiability of Cascaded Dynamical Systems). Consider the cascaded dynamical system $\mathbf{\Sigma}_{\mathcal{C}}$ in Definition 3.1, and let G_0, T_1, D_1 and U_2 be low-order blocks defined in Definition 3.3. If $m \geq \max\{n_D, n_U + q_D, q_U\}$, and the following rank conditions hold:

$$\operatorname{rank}(G_0) = n_D, \quad \operatorname{rank}\begin{pmatrix} T_1 & D_1 \end{pmatrix} = n_U + q_D, \quad \operatorname{rank}(U_2) = q_U, \tag{D.1}$$

then, $\mathbf{\Sigma}_{\mathcal{C}}$ with the augmented state X is locally observable and identifiable around X_0.

Proof. We first decompose the observability-identifiability matrix at $X = X_0$:

$$O_I(X_0) = \begin{pmatrix} \frac{\partial(L_F^0 y)}{\partial x_D} & \frac{\partial(L_F^0 y)}{\partial x_U} & \frac{\partial(L_F^0 y)}{\partial \theta_D} & \frac{\partial(L_F^0 y)}{\partial \theta_U} \\ \frac{\partial(L_F^1 y)}{\partial x_D} & \frac{\partial(L_F^1 y)}{\partial x_U} & \frac{\partial(L_F^1 y)}{\partial \theta_D} & \frac{\partial(L_F^1 y)}{\partial \theta_U} \\ \vdots & \vdots & \vdots & \vdots \\ \frac{\partial(L_F^{N-1} y)}{\partial x_D} & \frac{\partial(L_F^{N-1} y)}{\partial x_U} & \frac{\partial(L_F^{N-1} y)}{\partial \theta_D} & \frac{\partial(L_F^{N-1} y)}{\partial \theta_U} \end{pmatrix} = \begin{pmatrix} J_0 \\ J_1 \\ \vdots \\ J_{N-1}, \end{pmatrix},$$

(By Definition 3.2)

where J_k denotes the k-th row block of $O_I(X)$ defined by

$$J_k := \begin{pmatrix} \frac{\partial(L_F^k y)}{\partial x_D}(X_0) & \frac{\partial(L_F^k y)}{\partial x_U}(X_0) & \frac{\partial(L_F^k y)}{\partial \theta_D}(X_0) & \frac{\partial(L_F^k y)}{\partial \theta_U}(X_0) \end{pmatrix} \in \mathbb{R}^{m \times N}.$$

Since the observables vector y depends only on the downstream state x_D, the zeroth-order block takes the form:

$$J_0 = \begin{pmatrix} G_0 & 0_{m \times n_U} & 0_{m \times q_D} & 0_{m \times q_U} \end{pmatrix}.$$

(By Definition 3.1)

By direct computation of the first and second order Lie derivatives (Definition 2.3), we have

$$J_1 = \begin{pmatrix} A_1 & T_1 & D_1 & 0_{m \times q_U} \end{pmatrix}, \quad J_2 = \begin{pmatrix} A_2 & A_3 & A_4 & U_2 \end{pmatrix},$$

where

$$A_1 = \left(\frac{\partial(L_F^1 y)}{\partial x_D}(X_0)\right) \in \mathbb{R}^{m \times n_D}, \quad A_2 = \left(\frac{\partial(L_F^2 y)}{\partial x_D}(X_0)\right) \in \mathbb{R}^{m \times n_D},$$
$$A_3 = \left(\frac{\partial(L_F^2 y)}{\partial x_U}(X_0)\right) \in \mathbb{R}^{m \times n_U}, \quad A_4 = \left(\frac{\partial(L_F^2 y)}{\partial \theta_D}(X_0)\right) \in \mathbb{R}^{m \times q_D}$$

By (D.1), there exists row selection matrices $R_0 \in \mathbb{R}^{n_D \times m}$, $R_1 \in \mathbb{R}^{(n_U + q_D) \times m}$, and $R_2 \in \mathbb{R}^{q_U \times m}$ such that

$$\hat{G}_0 := R_0 G_0, \quad \hat{W}_1 := R_1 \begin{pmatrix} T_1 & D_1 \end{pmatrix}, \quad \hat{U}_2 := R_2 U_2$$

are invertible $n_D \times n_D$, $(n_U + q_D) \times (n_U + q_D)$ and $q_U \times q_U$ matrices, respectively. We then apply R_0, R_1, R_2 to J_0, J_1, J_2, and group the second and third column blocks for simplicity:

$$M := \begin{pmatrix} R_0 J_0 \\ R_1 J_1 \\ R_2 J_2 \end{pmatrix} = \begin{pmatrix} \hat{G}_0 & 0 & 0 \\ R_1 A_1 & \hat{W}_1 & 0 \\ R_2 A_2 & R_2 \begin{pmatrix} A_3 & A_4 \end{pmatrix} & \hat{U}_2 \end{pmatrix} \in \mathbb{R}^{N \times N}.$$

Notice that M is invertible since it is block lower triangular with invertible diagonal blocks $\hat{G}_0, \hat{W}_1$ and $\hat{U}_2$. Along with the fact that M is a $N \times N$ submatrix of $O_I(X_0)$, we conclude that $\mathrm{rank}(O_I(X_0)) = N$, which implies the local observability and identifiability by Lemma D.1.

This completes the proof. □

D.2 Proof of Theorem 3.9

We now give the main proof of Theorem 3.9.

Theorem D.3 (Theorem 3.9 Restated, ψ-Identifiability for Reduced Cascaded System)**.** Let $\mathbf{\Sigma}_{\mathcal{R}}$ be a cascaded system obtained from an admissible reduction $\rho(\cdot)$, with reduced augmented state $Z(t) \in \mathbb{R}^M$ (Definition 3.5), reduced observable $\gamma(\cdot) \in \mathbb{R}^r$, and reduced augmented vector field F_{red}. Let $\psi : \mathbb{R}^M \to \mathbb{R}^d$ be a target functional on $Z(t)$. Suppose that the reduced observable $\gamma(\cdot)$ is ρ-compatible, and in a neighborhood of a fixed point $Z_0 := Z(t_0)$,

$$\mathrm{row}(D\psi(Z)) \subseteq \mathrm{row}(D\mathcal{J}(Z)),\ \mathcal{J}(Z) := ((L^0_{F_{\mathrm{red}}}\gamma(Z))^\top (L^1_{F_{\mathrm{red}}}\gamma(Z))^\top (L^2_{F_{\mathrm{red}}}\gamma(Z))^\top)^\top \in \mathbb{R}^{3r}$$

has local constant rank. Then, $\mathbf{\Sigma}_{\mathcal{R}}$ is locally ψ-observable and ψ-identifiable around Z_0.

Proof. Let $Z, \widetilde{Z}$ lie in a sufficiently small neighborhood $\mathcal{U}$ of Z_0, and suppose that their reduced observation trajectories coincide. Then, by direct computation of the Lie derivatives along F_{red} ((Definition 2.3), the reduced readout and all of its time derivatives agree at t_0. This implies,

$$\mathcal{J}(Z) = \mathcal{J}(\widetilde{Z}).$$

We now show that this equality forces agreement of the target functional ψ.

Since J has locally constant rank on $\mathcal{U}$, the constant-rank theorem implies that, after possibly shrinking $\mathcal{U}$, the level sets of J form a smooth local foliation. Moreover, for any $Z \in \mathcal{U}$, $\ker D\mathcal{J}(Z)$ is the tangent space to the local level set of J through Z. The row-inclusion condition

$$\mathrm{row}(D\psi(Z)) \subseteq \mathrm{row}(D\mathcal{J}(Z))$$

is equivalent to

$$\ker D\mathcal{J}(Z) \subseteq \ker D\psi(Z).$$

Thus, every tangent direction along a local level set of $\mathcal{J}$ is annihilated by $D\psi$. Hence ψ is locally constant on each connected component of every level set of $\mathcal{J}$ in $\mathcal{U}$.

Because $\mathcal{J}(Z) = \mathcal{J}(\widetilde{Z})$, the two states belong to the same local level set of $\mathcal{J}$. Taking $\mathcal{U}$ small enough so that the relevant level-set component is connected, we obtain

$$\psi(Z) = \psi(\widetilde{Z}).$$

Therefore, any two locally indistinguishable reduced states agree on the target functional, and the reduced cascaded system is locally ψ-observable and locally ψ-identifiable around Z_0. This completes the proof.

□

E Experimental Details

E.1 Synthetic Cascade Setup and Diagnostics

This section provides the synthetic model specification and numerical diagnostic protocol used in Section 4.1. We reiterate that, the purpose of the synthetic study is not to benchmark prediction accuracy, but to (1) test whether the proposed cascaded diagnostics reveal upstream nonidentifiability when downstream observations remain nearly unchanged, and (2) to show the computational advantage of the criterion in Theorem 3.4 over constructing and testing the full observability-identifiability matrix.

Model Family and Observation Setup. We consider a smooth cascaded system $\mathbf{\Sigma}_{\mathcal{S}}$ observed on a time grid $\{0 \le t_1 \le t_2 < \cdots < t_T\}$ with upstream state $x_U \in \mathbb{R}^{n_u}$, upstream parameters $\theta_U \in \mathbb{R}^{q_U}$, coupling map $h(\cdot) \in \mathbb{R}^{d_h}$, downstream state $x_D \in \mathbb{R}^{n_D}$, downstream parameters $\theta_D \in \mathbb{R}^{q_D}$, known input $u(t) \in \mathbb{R}^{d_U}$[3] and observed samples $y(t_\ell)$. Explicitly, the system takes the following form:

$$\mathbf{\Sigma}_{\mathcal{S}} : \begin{cases} \dot{x}_U(t) = K_U(\theta_U)\, x_U(t) + Q_U(\theta_U)\, u(t) + \Phi_U(x_U, u), \\ \dot{x}_D(t) = K_D(\theta_D)\, x_D(t) + Q_D(\theta_D)\, h(x_U) + \Phi_D(x_D, h), \\ y(t_\ell) = g(x_D(t_\ell)) + \epsilon_\ell, \quad \text{for} \quad \ell = 1, \ldots, T, \end{cases}$$

where $K_U : \mathbb{R}^{q_U} \to \mathbb{R}^{n_u \times n_u}, Q_U : \mathbb{R}^{q_U} \to \mathbb{R}^{n_u \times d_u}, K_D : \mathbb{R}^{q_D} \to \mathbb{R}^{n_D \times n_D}$ and $Q_D : \mathbb{R}^{q_D} \to \mathbb{R}^{n_D \times d_h}$ define the parameterized linear backbone of the cascade, while $\Phi_U : \mathbb{R}^{n_u} \times \mathbb{R}^{d_u} \times \mathbb{R}^{s_U} \to \mathbb{R}^{n_u}$ and $\Phi_D : \mathbb{R}^{n_D} \times \mathbb{R}^{d_h} \times \mathbb{R}^{s_D} \to \mathbb{R}^{n_D}$ define smooth nonlinear corrections. Here, the fixed downstream readout $g : \mathbb{R}^{n_D} \to \mathbb{R}^m$ is the noiseless signal that defines the structural observation map for identifiability analysis, and the perturbation ϵ_ℓ represents finite sample observation noise.

Next, we instantiate this family with K upstream channels and a scalar downstream readout. In the following two cases, we consider $x_U := (X_1(t), \ldots, X_K(t))^\top, x_D := D(t)$, coupling map $h(x_U) := \sum_{i=1}^K X_i(t)$ and $g(D(t)) = D(t)$. Further, the known input is $u(t) := At^2 \exp(-\gamma t)$ with $A = 1.8, \gamma = 0.9$, and the fitted parameters are $\theta_U := (a, b_1, \ldots, b_K) \in \mathbb{R}^{K+1}$ and $\theta_D := \lambda \in \mathbb{R}$.

Linear Cascade. The linear instantiation removes the nonlinear correction terms by setting $\Phi_U, \Phi_D = 0$, and $K_U = -aI_K, Q_U = (b_1, \ldots, b_K)^\top, K_D(\theta_D) = -\lambda, Q_D(\theta_D) = 1$, giving:

$$\mathbf{\Sigma}_{\mathcal{L}} : \begin{cases} \dot{X}_i = -aX_i(t) + b_i u(t) \quad \text{for} \quad i = 1, \ldots, K \\ \dot{D}(t) = -\lambda D(t) + \sum_{i=1}^K X_i(t), \\ y(t) = D(t) \end{cases}$$

Trigonometric Cascade. The nonlinear instantiation keeps the same linear backbone and adds smooth trigonometric corrections. For each upstream channel,

$$\Phi_U(X_i) = \eta_U \sin(\omega_U X_i) + \beta_U \big(\cos(\omega_U X_i) - 1\big), \tag{E.1}$$

and for the downstream subsystem, with $S := \sum_{i=1}^K X_i(t)$, we set

$$\Phi_D(t) = \eta_D \sin(\omega_D S) + \beta_D \big(\cos(\omega_D S) - 1\big). \tag{E.2}$$

Together, this gives the following nonlinear cascade:

$$\mathbf{\Sigma}_{\mathcal{T}} : \begin{cases} \dot{X}_i = -aX_i(t) + b_i u(t) + \Phi_U(X_i, u) \quad \text{for} \quad i = 1, \ldots, K \\ \dot{D}(t) = -\lambda D(t) + \sum_{i=1}^K X_i(t) + \Phi_D(t), \\ y(t) = D(t), \end{cases}$$

where $\Phi_U(X_i, u)$ and $\Phi_D(t)$ follows (E.1) and (E.2), respectively. The nonlinear coefficients are fixed as $\eta_U = 0.10$, $\beta_U = 0.04$, $\omega_U = 1.20$, $\eta_D = 0.16$, $\beta_D = 0.06$, and $\omega_D = 1.00$.

Fitting Protocol and Retained Minimizers. Synthetic observations are generated from the ground truth trajectory and perturbed by independent Gaussian observation noise with standard deviation 0.003. For each synthetic family, we fit the unknown parameter $\theta = (\theta_U, \theta_D)$ from multiple random initializations by minimizing the empirical downstream reconstruction loss over a sampled time grid,

$$\mathcal{L}(\theta) = \frac{1}{T} \sum_{\ell=1}^{T} |y_\theta(t_\ell) - y_{\text{obs}}(t_\ell)|^2,$$

here, $y_\theta(t_\ell)$ denotes the downstream readout generated by the candidate parameter vector θ at time t_ℓ, and $y_{\text{obs}}(t_\ell)$ denotes the noisy synthetic observation at the same time point. Let $\mathcal{L}_{\min}$ denote the smallest loss across all initialization runs. We retain a fitted solution as near optimal if its loss satisfies

$$\mathcal{L}(\theta) \le \max\big\{\mathcal{L}_{\min} + 2 \times 10^{-5}, (1 + 5 \times 10^{-3})\mathcal{L}_{\min}\big\}.$$

[3]Here, we include u(t) to generate nontrivial transient dynamics without adding an unknown inference target.

This rule combines an absolute tolerance and a relative tolerance around the best-fitted loss. The absolute tolerance accounts for finite sample noise and numerical optimization error at the scale of the empirical reconstruction loss, while the relative tolerance prevents the retained set from depending too sensitively on the magnitude of $\mathcal{L}_{\min}$. After applying the loss tolerance, we remove near-duplicate fitted parameter vectors using the Euclidean parameter distance threshold 5×10^{-2}. The remaining solutions are ordered by loss, and at most six retained minimizers are kept for subsequent analysis.

This retained set is used only to define a reproducible collection of near-optimal fitted explanations. The tolerance rule is not a criterion for identifiability. In particular, the conclusion does not rely on the number of retained minimizers alone. Instead, we use the retained solutions to test whether multiple fitted parameter vectors produce nearly indistinguishable downstream trajectories, and then compare this empirical non-uniqueness with the sampled low-order rank diagnostics. This separation is important because local optimization can produce repeated or nearly identical fits. The parameter distance filter removes such duplicates before reporting the number of retained minimizers.

For output discrepancy, we take the best-fitted retained solution as the reference trajectory. For every other retained minimizer indexed by j, we compare its downstream trajectory to this reference trajectory at each sampled time point. The reported mean error and standard deviation summarize $|\delta y_j(t_\ell)|$ across retained non reference minimizers and sampled time points, and the reported maximum error records the largest absolute downstream discrepancy over the same collection. These output errors quantify how much the downstream observation changes across retained fitted explanations. In the synthetic experiments, these errors remain small even though the upstream sensitive diagnostic blocks are rank deficient, supporting the interpretation that the downstream readout does not resolve the upstream ambiguity.

The synthetic diagnostics in Figure 3(a) and Table 1 use $K = 2$ upstream channels for both the linear and trigonometric cascades. Table 1 reports the mean $\pm$ standard deviation and the maximum of $|\delta y_j(t_\ell)|$ over retained non-reference minimizers and sampled time points. The runtime scaling study in Figure 3(b) varies K to assess construction cost.

Sampled Low-Order Diagnostics. For each retained family, we compute the cascaded low-order blocks G_0, $[T_1\ D_1]$, and U_2 along sampled trajectory points. Since the synthetic readout has a single downstream observation channel, the pointwise rows from one sampled time point are not sufficient to represent all parameter directions. We therefore stack the sampled blocks across time and report whether each stacked block reaches its target rank. This diagnostic should be read as a sampled trajectory level rank probe, not as a pointwise verification of the sufficient condition in Theorem 3.4.

In the reported setting, the downstream block G_0 tests whether the downstream state is visible through the readout. The block $[T_1\ D_1]$ tests whether the downstream first-order information separates upstream state directions and downstream parameter directions. The block U_2 tests whether the upstream parameter directions become visible through the cascaded propagation. Thus, rank deficiency in $[T_1\ D_1]$ or U_2 identifies the part of the cascade where upstream ambiguity enters the downstream observation pathway. Table 1 reports these rank checks together with the retained minimizer count and downstream output discrepancy, so that the empirical non-uniqueness and the sampled rank deficiency can be interpreted jointly. Figure 5 provides a visual counterpart.

E.2 Dataset, Preprocessing, and Compute Environment

The cerebral hemodynamics experiments use ISLES2017, a publicly available acute ischemic stroke dataset (Winzeck et al., 2018). ISLES2017 contains 44 patients with 4D DSC-MRI perfusion images, deconvolution-derived perfusion maps (CBF, CBV, MTT), and gold-standard expert-annotated lesion masks. MR signals are converted to voxelwise tracer concentration (Fieselmann et al., 2011).

To obtain anatomically meaningful partitions $\{\Omega_r\}_{r=1}^{N}$, we apply FreeSurfer (Fischl, 2012) to segment grey matter (GM), white matter (WM), cerebrospinal fluid (CSF), and brain stem. GM and WM are further divided using k-means on tracer time-curve features (8 clusters for GM, 6 for WM), neighboring clusters with fewer than 400 voxels are merged. For each Ω_r, parameters are fitted from the regional mean concentration curve.

All experiments were run on CPU without GPU acceleration (Apple M1 MacBook Pro). The implementation uses Python with standard scientific computing libraries.

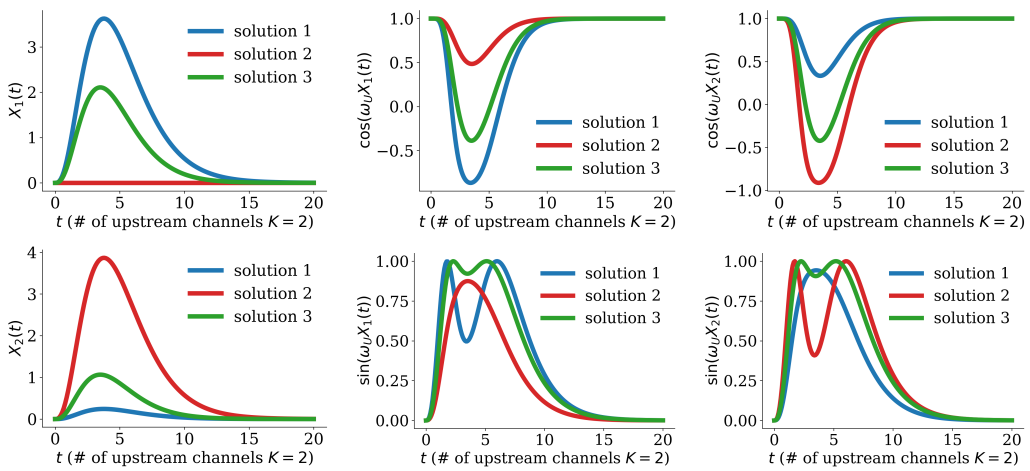


Figure 5: Indistinguishable state visualization, shown with $K = 2$ for display. Left: linear cascade states X_1 and X_2. Middle: trigonometric visualization for X_1, with $\cos(X_1)$ on top and $\sin(X_1)$ on bottom. Right: trigonometric visualization for X_2, with $\cos(X_2)$ on top and $\sin(X_2)$ on bottom.

E.3 Tracer Reconstruction for Reduced Cascade

In this section, we detail the construction of the reduction maps, forward prediction, reconstruction objective, and minimizer selection criterion used to solve the composite inverse problem for the cerebral hemodynamics digital twin under the reduced cascaded system introduced in Section 4.2.1.

Reduction and Observation Scale. The reduced inverse problem begins with a choice of scale. The observed tracer data are available as fine-grained rasterized voxel-wise curves, while the reduced cascade is evaluated on tissue supports that summarize effective regional exchange. Thus, the admissible reduction $\rho(\cdot)$ specifies how voxel-level observations and latent concentration fields are represented at the support level. Although Definition 3.5 allows more general admissible reductions, our experiments adopt spatial aggregation over a partition $\{\Omega_r\}_{r=1}^R$ of the brain domain. Specifically, the reduced cascade is represented by support level coordinates $Z(t) = (z_D, z_U, \alpha_D, \alpha_U)$, with

$$z_D(t) = (\bar{z}_D^1, \dots, \bar{z}_D^R), \quad \text{and} \quad \bar{z}_D^r = \frac{1}{|\Omega_r|} \sum_{v \in \Omega_r} C_v(t), \tag{E.3}$$

where $C_v(t)$ denotes the voxel level tracer concentration field before reduction, and upstream state, parameters and the downstream parameters are represented at the same support scale.

With this aggregation-based choice of admissible reduction ρ fixed, the reduced observation channel is specified on the same tissue supports. For voxel-level tracer observations $y_v(t)$, the fixed readout $\chi_{\text{red}}(y)^r$ is defined by the spatial aggregation as in (E.3). In our setting, in particular, the model readout extracts the regional concentration: $\gamma^r(z_D^r) = \bar{z}_D^r$ for all tissue support $r = 1, \dots, R$. Therefore, this compatibility specifies the empirical reconstruction pair in the inverse problem: the model-predicted readout $\gamma^r(\hat{z}_D(t_\ell))$ is compared with the reduced observation $\chi_{\text{red}}(y)^r(t_\ell)$.

Regional Reconstruction Objective. Given the reduced observation channel and compatible readout above, parameter estimation is posed at the support-level. For each tissue support Ω_r, let (α_D^r, α_U^r) denote the regional downstream and upstream parameters. For a candidate (α_D^r, α_U^r), the predicted downstream trajectory $\hat{z}_D(t; \alpha_D^r, \alpha_U^r)$ is generated by the composite forward solve: the upstream closure in (4.1) determines $\hat{z}_{U,\text{out}}^r(t)$, and the tracer compartment in (4.3) maps this outflow, the exogenous tracer input, and the downstream volume scale into the predicted concentration coordinate $\hat{\bar{z}}_D^r(t; \alpha_D^r, \alpha_U^r)$. In our implementation, the exogenous tracer input is written as

$$z_{D,\text{in}}^r(t) = a_{\text{in}}^r \widetilde{z}_{D,\text{in}}(t),$$

where $\widetilde{z}_{D,\mathrm{in}}(t)$ is the shared unit-peak input curve and $a^r_{\mathrm{in}} > 0$ is the fitted regional amplitude. The reconstruction term for support Ω_r is

$$\mathcal{L}^r_{\mathrm{rec}}(\alpha^r_D, \alpha^r_U) = \frac{\sum_{\ell=1}^T \omega_\ell \left(\gamma^r(\hat{z}_D(t_\ell; \alpha^r_D, \alpha^r_U)) - \chi_{\mathrm{red}}(y)^r(t_\ell)\right)^2}{\sum_{\ell=1}^T \omega_\ell \left(\chi_{\mathrm{red}}(y)^r(t_\ell)\right)^2 + \epsilon}, \tag{E.4}$$

$r = 1, \ldots, R$. Here $\omega_\ell = 0.2 + 0.8 z_{D,\mathrm{in}}(t_\ell)$ is a positive time weight that emphasizes the tracer passage window. Together, the fitted regional parameters minimize the following regularized objective:

$$\mathcal{L}_{\mathrm{fit}} = \sum_{r=1}^R \mathcal{L}^r_{\mathrm{rec}}(\alpha^r_D, \alpha^r_U) + \lambda_{\mathrm{in}}\mathcal{L}_{\mathrm{in}} + \lambda_V \mathcal{L}_V + \lambda_F \mathcal{L}_F + \lambda_{\mathcal{R}} \mathcal{L}_{\mathcal{R}}, \tag{E.5}$$

where the log-scale prior terms are

$$\mathcal{L}_{\mathrm{in}} = \sum_{r=1}^R \left(\log a^r_{\mathrm{in}} - \log a^r_{\mathrm{in},0}\right)^2, \qquad \mathcal{L}_V = \sum_{r=1}^R \left(\log \alpha^r_{D,V} - \log \alpha^r_{D,V,0}\right)^2,$$

$$\mathcal{L}_F = \sum_{r=1}^R \left(\log \alpha^r_{U,F} - \log \alpha^r_{U,F,0}\right)^2, \qquad \mathcal{L}_{\mathcal{R}} = \sum_{r=1}^R \left(\log \alpha^r_{U,\mathcal{R}} - \log \alpha^r_{U,\mathcal{R},0}\right)^2.$$

The priors $a^r_{\mathrm{in},0}$, $\alpha^r_{D,V,0}$, and $\alpha^r_{U,F,0}$ are derived from the regional tracer curve and baseline perfusion summaries, while $\alpha_{U,\mathcal{R},0}$ is a fixed reference. All optimized quantities are constrained to be positive by optimizing in log coordinates. We set $\lambda_{\mathrm{in}} = 0.06$, $\lambda_V = \lambda_F = 0.04$, $\lambda_{\mathcal{R}} = 0.002$ and $\alpha_{U,\mathcal{R},0} = 1$.

Optimization and Retained Minimizers. Parameter estimation is carried out at the support-level. Since the regional objective in (E.5) decomposes over tissue supports, each regional parameter tuple (α^r_D, α^r_U) is fitted independently. Each regional subproblem is solved by bounded Powell optimization in log coordinates. The optimized positive quantities include the Windkessel resistance scale, relaxation time, downstream volume, effective flow scale, and tracer input amplitude. The resistance scale determines the initial outflow in (4.1) through the fixed pressure reference p_0, so the resistance prior in (E.5) regularizes this upstream initialization while preserving positivity.

For each patient, we repeat the regional fitting procedure from randomized restarts around data-informed initial values. The initial effective flow and downstream volume are centered at regional CBF and CBV summaries, the tracer input amplitude is centered at the regional tracer peak, and the relaxation time is initialized within the acquisition window. Each restart applies seeded random jitter to these regional initial values, with the perturbation scale increasing across restart index. Thus, the restarts explore nearby local solutions around physiologically informed centers rather than drawing parameters from an unstructured random distribution. Each restart produces a full regional solution across all tissue supports. We rank restarts by the mean reconstruction loss across supports, computed from (E.4), rather than by the full regularized objective. This choice isolates parameterizations with comparable tracer readouts, since the subsequent empirical identifiability diagnostics are intended to compare solutions that are nearly indistinguishable at the reduced observation level.

Finally, the retained minimizers are the restarts whose mean reconstruction loss lies within a relative tolerance $\delta_{\mathrm{keep}} = 0.02$ of the best restart. If fewer than three restarts satisfy this criterion, we retain the three restarts with the smallest mean reconstruction losses. Each retained restart defines a reduced tuple $\hat{Z}_s = (\hat{z}_{D,s}, \hat{z}_{U,s}, \hat{\alpha}_{D,s}, \hat{\alpha}_{U,s})$. The functional discrepancy and row space residual diagnostics reported in the main text are computed across these retained tuples.

F Additional Numerical Results

In this section, we provide additional numerical experiments that complement the main synthetic and cerebral perfusion studies presented in Section 4. Specifically, Appendix F.1 tests whether the computational advantage of the low-order construction persists under nonlinear upstream-to-downstream couplings. Appendix F.2 studies the gap between low-order certification and higher-order observability-identifiability analysis. Appendix F.3 evaluates the numerical stability of the functional-identifiability diagnostic near degenerate regimes in the reduced physiological state space. Finally, Appendix F.4 examines how prescribed input profiles affect numerical certification and distinguishes these effects from structural ambiguity imposed by the observation pathway.

F.1 Nonlinear-Coupling Scalability

The synthetic experiments in Appendix E.1 increase the number of upstream latent components while using an additive upstream-to-downstream coupling. Here, we add an additional set of synthetic scalability experiments to further test whether the computational advantage persists when the coupling map itself becomes nonlinear. Specifically, we compare the following four couplings:

$$h_{\mathrm{sum}}(x_U) = \sum_{i=1}^{K} X_i, \quad h_{\mathrm{sep}}(x_U) = \sum_{i=1}^{K} \sin(X_i), \quad h_{\mathrm{pair}}(x_U) = \sum_{1 \le i < j \le K} X_i X_j,$$

$$h_{\mathrm{dense}}(x_U) = \sin\left(\sum_{i=1}^{K} \frac{i}{K+1} X_i \right) + \sum_{1 \le i < j \le K} \frac{i+j}{2K+1} X_i X_j.$$

The separable case tests channel-wise nonlinear dependence, the pairwise case introduces explicit interactions among upstream components, and the dense mixed case combines a nonlinear aggregate with weighted pairwise interactions. For each coupling, we compare symbolic construction of the proposed low-order blocks, using derivative orders 0–2, against a higher-order observability-identifiability reference using derivative orders 0–6.

Table 3: Runtime ratio between higher-order construction and the proposed low-order construction across nonlinear couplings. Larger value means larger speedup from using the low-order blocks.

Coupling	$K=1$	$K=2$	$K=3$	$K=4$	$K=5$	$K=6$
Additive sum	$2.8\times$	$5.1\times$	$6.0\times$	$6.7\times$	$7.4\times$	$7.9\times$
Separable nonlinear	$12.2\times$	$18.5\times$	$22.5\times$	$26.4\times$	$29.4\times$	$31.6\times$
Pairwise polynomial	$1.3\times$	$6.4\times$	$9.8\times$	$13.2\times$	$17.4\times$	$16.3\times$
Dense mixed	$13.0\times$	$47.1\times$	$116.3\times$	$212.4\times$	$345.1\times$	$\geq 413.0\times$

The results show that the low-order construction remains substantially faster across all tested coupling maps. For the additive coupling, the speedup grows from $2.8\times$ at $K=1$ to $7.9\times$ at $K=6$. The advantage becomes larger under nonlinear couplings, reaching $31.6\times$ for separable nonlinear coupling and at least $413.0\times$ for the dense mixed coupling at $K=6$, where the higher-order construction reached the 120s timeout while the low-order construction completed in 0.291s.

These results refine the computational claim. Nonlinear couplings are not cost-free, since the low-order blocks still inherit the algebraic complexity of the coupling map. However, the higher-order reference construction repeatedly propagates the same nonlinear expressions through additional Lie derivatives, which leads to much faster expression growth. Thus, the advantage of the proposed criterion is not specific to the additive coupling in Appendix E.1. It comes from restricting symbolic differentiation to the structurally relevant low-order blocks.

F.2 Low-Order Certification and Higher-Order Recovery

The criterion in Theorem 3.4 provides a sufficient low-order certificate for local observability and structural identifiability. Its failure indicates non-certification at derivative orders $0-2$, but does not preclude recovery by the full observability-identifiability matrix. We therefore evaluate a staged certification procedure that first constructs the proposed low-order blocks and invokes higher-order observability-identifiability analysis only for non-certified cases.

We use controlled symbolic cascades that represent the possible outcomes of this workflow. In the low-order certified setting, each upstream component has its own downstream readout path:

$$\dot{D}_i = \sin(X_i) + \frac{1}{5} X_i^2, \quad \dot{X}_i = \theta_i + \frac{1}{7} \sin^2(X_i), \quad \dot{\theta}_i = 0, \quad y_i = D_i.$$

Here, derivative orders 0–2 expose the downstream states, upstream states, and parameters. In the higher-order recovery setting,

$$\dot{D} = X^q, \quad \dot{X} = 1, \quad y = D,$$

evaluated at $X = 0$, with $q \in \{3, 4, 5, 6, 7\}$. In this case, the low-order rows miss the upstream direction, but higher-order derivatives recover it. Finally, in the persistently uncertified setting, the downstream dynamics depend only on an aggregate $\sum_i X_i$, matching the confounding mechanism in Section E.1, so individual upstream directions remain unresolved even after higher-order construction.

For each system, we compare three procedures: constructing only the low-order rows through order 2, constructing the higher-order reference through order 7, and the gated workflow that first constructs the low-order rows and constructs the higher-order reference only if the low-order rank test fails.

Table 4: Runtime of the gated workflow compared with always constructing the higher-order reference. The three outcome columns count whether the low/high order test passes or fails.

Setting	n	LP/HP	LF/HP	LF/HF	Gated time	Higher-order time
Low-order certified	5	5	0	0	0.421s	13.078s
Higher-order recovery	5	0	5	0	0.252s	0.134s
Persistently uncertified	8	0	0	8	0.227s	0.112s
All systems	18	5	5	8	0.900s	13.323s

LP/HP denotes systems where both the low-order and higher-order constructions reach full rank, LF/HP denotes systems where the low-order construction fails but the higher-order construction reaches full rank, and LF/HF denotes systems where both fail. The gated workflow preserves the intended computational advantage because it avoids higher-order construction whenever the low-order certificate succeeds. In the low-order certified systems, the gated workflow stops after order 2 and is $31.0\times$ faster than always constructing the higher-order reference. Across all 18 systems, including the cases that require escalation, the gated workflow remains $14.8\times$ faster overall. Thus, low-order failure is not treated as evidence of non-identifiability. Instead, the low-order criterion acts as a cheap positive certificate, while higher-order analysis remains available for inconclusive cases.

F.3 Stability Near Degenerate Regimes

The functional-identifiability result in Theorem 3.9 is local and relies on the row space of the target Jacobian being stable near the evaluated reduced state. In physiological models, possible rank changes may occur near degenerate regimes such as very low flow, small volume, delayed transit, or weak observed signal. We therefore evaluate whether the row-space diagnostic used in Section 4.2 is sensitive to such regimes or to the numerical SVD threshold used to estimate rank.

We use the same reduced cerebral-perfusion pipeline as in Section 4.2 and evaluate all 44 successfully processed ISLES2017 patients. For each patient and retained minimizer, we construct the same stacked order-$0-2$ matrix $W_1 = D\mathcal{J}(\mathcal{Z})$ and target Jacobian $W_2 = D\psi_{\mathrm{MTT}}(Z)$ used in the main experiment. We then restrict the columns to different ROI subsets in the reduced regional state space:

$$\text{all ROIs}, \quad \text{high MTT top } 20\%, \quad \text{low flow bottom } 20\%,$$
$$\text{low volume bottom } 20\%, \quad \text{low peak signal bottom } 20\%.$$

The low-flow subset directly targets the near-zero-flow regime, while the other subsets probe related near-degenerate physiological or observational regimes. For each subset and each relative SVD threshold, we report two quantities. The first is the MTT row-space residual, which measures whether $D\psi_{\mathrm{MTT}}(Z)$ lies in the row space of $D\mathcal{J}(Z)$. The second is the numerical rank fraction, defined as the numerical rank of the ROI-restricted W_1 divided by the number of retained columns in that restricted matrix. Smaller SVD thresholds discard fewer singular directions, so stability across thresholds indicates that the diagnostic is not driven by a single arbitrary cutoff.

Table 5 and Table 6 show that the functional-identifiability diagnostic is stable across both boundary-like subset choices and SVD thresholds. The MTT row-space residual remains near zero for all tested subsets, including the low-flow regime, and decreases to numerical precision for thresholds at or below 10^{-6}. The numerical rank fraction follows the same pattern: boundary-like subsets do not exhibit degradation relative to the all-ROI calculation and reach full numerical rank fraction from 10^{-6} onward. After the proposed regional reduction, the MTT functional-identifiability diagnostic is not driven by a single threshold choice or by a small set of near-degenerate ROIs.

Table 5: Median MTT row-space residual across 44 patients under different boundary-like ROI subsets and SVD thresholds.

Subset \ Threshold	10^{-4}	10^{-5}	10^{-6}	10^{-7}	10^{-8}	10^{-9}	10^{-10}
All ROIs	2.67×10^{-4}	1.92×10^{-6}	5.08×10^{-8}	2.45×10^{-10}	1.27×10^{-12}	7.55×10^{-16}	7.18×10^{-16}
High MTT top 20%	6.79×10^{-7}	1.76×10^{-9}	8.98×10^{-16}	6.74×10^{-16}	6.40×10^{-16}	6.31×10^{-16}	6.29×10^{-16}
Low flow bottom 20%	3.65×10^{-5}	1.38×10^{-9}	8.04×10^{-16}	6.57×10^{-16}	5.92×10^{-16}	5.82×10^{-16}	5.75×10^{-16}
Low volume bottom 20%	4.64×10^{-4}	4.45×10^{-8}	7.95×10^{-16}	6.57×10^{-16}	6.26×10^{-16}	6.20×10^{-16}	6.15×10^{-16}
Low peak signal bottom 20%	3.52×10^{-5}	9.28×10^{-8}	8.30×10^{-16}	7.71×10^{-16}	6.99×10^{-16}	6.56×10^{-16}	6.16×10^{-16}

Table 6: Median numerical rank fraction across 44 patients under different boundary-like ROI subsets and SVD thresholds.

Subset \ Threshold	10^{-4}	10^{-5}	10^{-6}	10^{-7}	10^{-8}	10^{-9}	10^{-10}
All ROIs	0.8267	0.9448	0.9657	0.9862	0.9875	1.0000	1.0000
High MTT top 20%	0.8500	0.9417	1.0000	1.0000	1.0000	1.0000	1.0000
Low flow bottom 20%	0.8667	0.9500	1.0000	1.0000	1.0000	1.0000	1.0000
Low volume bottom 20%	0.8500	0.9750	1.0000	1.0000	1.0000	1.0000	1.0000
Low peak signal bottom 20%	0.8500	0.9500	1.0000	1.0000	1.0000	1.0000	1.0000

F.4 Input-Dependent Identifiability

In addition to varying coupling structure and derivative order, we examine how the prescribed input affects identifiability in the linear cascade of Section E.1. We compare one and two upstream channels ($K = 1, 2$) under zero input, constant input, the pulse $u(t) = 1.8t^2 \exp(-0.9t)$ used in Section E.1, and a multisine input $u(t) = c[1 + 0.4\sin(t) + 0.4\sin(2.5t)]$. The constant amplitude and multisine scale c are chosen to match the pulse's integrated squared amplitude over $[0, 20]$. Each configuration uses the original parameter setting and 20 additional draws within the original parameter bounds, paired across inputs, with zero initial states. For each trajectory, we evaluate 25 interior time points, giving 525 evaluations per model–input pair. At each point, we construct the output-derivative Jacobian through order six, explicitly including the known-input derivatives, and estimate its rank using a relative SVD threshold of 10^{-8}. These are pointwise matrices, without temporal stacking. The augmented coordinates include the current downstream and upstream states, the decay rates, and the input gains, giving dimensions five and seven for $K = 1$ and $K = 2$, respectively. Because the scalar output precludes the low-order block condition in these configurations, this experiment evaluates the higher-order continuation of the proposed workflow.

Input-dependent numerical certification. For $K = 1$, constant, pulse, and multisine inputs yield full numerical rank at 55.2%, 79.4%, and 100% of the evaluated points (Figure 6, left). Each nonzero input provides at least one full-rank point in all 21 parameter settings. Thus, constant input can already support identification through transient dynamics, while the multisine preserves numerically distinguishable directions over more sampled points. The reported percentages characterize numerical certification across the sampled trajectories under the specified threshold. With zero input and initial states, the trajectories remain at rest and contain no information about input gains.

Input-independent structural ambiguity. For $K = 2$, all nonzero inputs increase the maximum observed rank from two to five, but none achieves the full augmented dimension of seven (Figure 6, right). The remaining deficiency arises because both upstream channels share the same decay rate and input, while the downstream subsystem observes only their sum. Consequently, the summed upstream dynamics depend only on the summed state and summed input gain, so different allocations between channels produce identical outputs under any common input. These two indistinguishable directions explain why the rank remains below seven even with the multisine input. Therefore, richer input variation improves numerical certification of observable directions but cannot overcome the information loss imposed by observing only the aggregate upstream response.

These experiments demonstrate how the analysis accommodates different prescribed inputs through their explicit inclusion in the derivative construction. Changes in numerical certification reflect changes in the information supplied by the input and the resulting trajectory. Specifically, varying the input alters numerical certification within the same model, while the two-channel aggregation

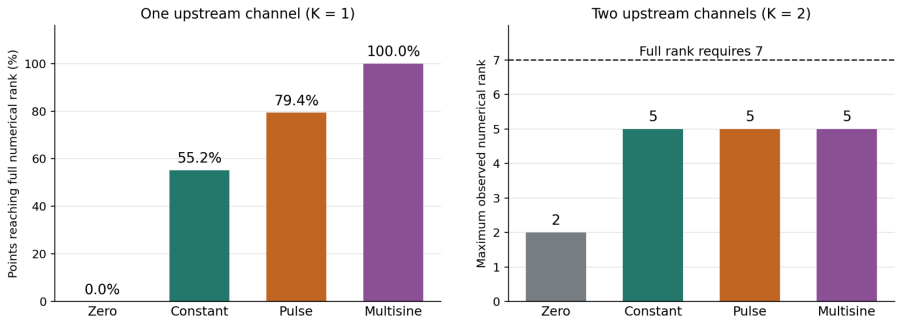


Figure 6: Input-dependent numerical certification in the linear cascade. **Left:** Percentage of the 525 evaluated points attaining full numerical rank for $K = 1$, whose augmented dimension is five. **Right:** Maximum numerical rank across the evaluated points for $K = 2$, whose augmented dimension is seven. All matrices include output derivatives through order six and use a relative SVD threshold of 10^{-8}. Nonzero inputs have matched integrated squared amplitude.

ambiguity persists across inputs, as predicted by the exact aggregate dynamics. The analysis therefore remains applicable across the tested input profiles without requiring identical certification outcomes.