Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Controlled Observation Fibres and Robust Causal Prefix Stopping in Sequential MRI

Takashi Matsuhisa

takashimatsuhisa.mri.bsbh@gmail.com

MRI BSBH Foundation for Sci. Res.

Research Article

Keywords: adaptive MRI, observational equivalence, robust e-process, selective inference, prefix stopping

Posted Date: October 26th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10724034/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Controlled Observation Fibres and Robust Causal Prefix Stopping in Sequential MRI

Takashi Matsuhisa[1,2*]

[1*]MRI BSBH Foundation for Scientific Research, Mito, Japan.
[2]IAMR, KarRC RAS, Petrozavodsk, Russia.

Corresponding author(s). E-mail(s): takashimatsuhisa.mri.bsbh@gmail.com, takashi.matsuhisa@krc.karelia.ru;

Abstract

Sequential MRI permits adaptive sampling and early stopping, but neither resolves observational equivalence nor likelihood misspecification. We represent acquisition histories by controlled observation fibres and restrict stopping to causal prefixes. Under a declared one-step model-error envelope, predictable normalisation of likelihood-ratio factors yields time-uniform outer confidence fibres and bounds the probability of a wrong, non-rejected report at any stopping time. A reproducible multicoil DCE-MRI phantom separates identifiable, singular, misspecified, and colliding regimes. Stored-image replay evaluates a frozen classifier and split-conformal label sets at predeclared gates; it concerns recorded prefixes rather than unrecorded scanner actions, raw-k-space reconstruction, or clinical performance. The framework separates identification, robust uncertainty, rejection, and acquisition cost.

Keywords: adaptive MRI, observational equivalence, robust e-process, selective inference, prefix stopping

1 Introduction

Sequential sensing couples three questions that are often treated separately: which observation to acquire, when to stop, and which task value to report. Controlled experiments can improve discrimination because the next observation law may depend on the selected action and the preceding history [1–3]. They cannot, however, separate two states that induce the same controlled path law. A high terminal score does not remove this obstruction, and a stopping rule based on a misspecified likelihood need not retain its nominal error level.

Recent MRI methods learn sequential sampling policies for reconstruction or pathology prediction from undersampled k-space data [4–7]. Du et al. also evaluate a multi-objective reinforcement-learning sampler on two sequential knee-pathology decisions [8]. Those studies address reconstruction or diagnostic performance; we ask which tasks survive equality of controlled path laws and derive a time-uniform joint error guarantee for selective stopping. E-values provide the probability tool [9, 10], whereas the observation fibre determines which unresolved states must remain together.

Indistinguishability quotients are classical in deterministic observability, and equality of observation-process laws is central to stochastic filtering [11, 12]. Here the fibre is indexed by a declared family of causal finite-horizon policies

and is used as the set retained by a stopping rule. The task is identified precisely when it is constant on each fibre. Stopped relative entropy quantifies the information accumulated away from exact collisions and yields a sample lower bound near singular ones. A predictable conditional envelope converts nominal likelihood-ratio scores into robust e-processes. Its outer fibre inversion reports a label only when the retained set is homogeneous. The guarantee bounds a wrong, non-rejected report jointly; it is not conditional on reporting.

The numerical study uses a reduced multicoil DCE-MRI model with a Tofts-type recursion, the Parker input, an SPGR response, analytic coils, and Cartesian complex k-space [13, 14]. Its linearised response defines an exact finite-hypothesis Gaussian experiment; the nonlinear response is a misspecification stress test.

A separate interface replays recorded prefixes through a frozen classifier and split-conformal label sets [15]. It concerns the recorded protocol, not unrecorded actions or raw-k-space reconstruction. All supplied arrays and labels are synthetic and establish no clinical performance.

Organisation. Sections 2–5 develop the law-level and time-uniform theory. Section 6 isolates what retrospective prefix replay can establish. Section 7 reports the synthetic experiment; the appendices contain the proofs and the full forward model.

2 Sequential MRI as a controlled experiment

Let the latent state space Ω, action space $\mathcal{U}$, and observation space $\mathcal{Y}$ be standard Borel. A state $\omega \in \Omega$ indexes the task label, structural and nuisance parameters, and the law of any latent process; it does not contain a realised innovation path unless the experiment is expressly conditioned on that path. Put $H_0 = \varnothing$ and let $\mathcal{F}_0$ be the trivial sigma-field. For $n \geq 1$, set

$$H_n = (U_1, Y_1, \dots, U_n, Y_n),$$
$$\mathcal{F}_n = \sigma(H_n).$$

At stage n, the scanner chooses $U_n \in \mathcal{U}$ from H_{n-1} and observes $Y_n \in \mathcal{Y}$. A causal policy is a sequence of kernels $\pi_n(\mathrm{d}u \mid H_{n-1})$. Conditional on ω, the observation kernel is

$$Q_{\omega,n}(\mathrm{d}y \mid H_{n-1}, U_n). \tag{1}$$

The policy is common to all states: a state-dependent policy would reveal the unknown state by assumption.

Proposition 2.1 (Path law) *Suppose that the policy and observation kernels are jointly measurable in their conditioning variables. For every ω, policy π, and finite horizon N, there is a unique probability law $P^{\pi}_{\omega,N}$ on $(\mathcal{U} \times \mathcal{Y})^N$ satisfying*

$$P^{\pi}_{\omega,N}(\mathrm{d}h_N) = \prod_{n=1}^{N} \pi_n(\mathrm{d}u_n \mid h_{n-1}) \times Q_{\omega,n}(\mathrm{d}y_n \mid h_{n-1}, u_n). \tag{2}$$

Its restriction to the first n stages is $P^{\pi}_{\omega,n}$.

Set $P^{\pi}_{\omega,0} = \delta_{\varnothing}$.

The assertion is the finite-horizon Ionescu–Tulcea construction [16]; Appendix A gives the kernel-by-kernel proof. Equality and distance between controlled experiments below always refer to the laws in (2), not merely to noiseless images.

The scientific task is a measurable map $g : \Omega \to \mathcal{D}$, with $\mathcal{D}$ finite. Let τ be an $\{0, \dots, N\}$-valued stopping time and D_τ an $\mathcal{F}_\tau$-measurable action in $\mathcal{A} = \mathcal{D} \cup \{\perp\}$, where $\perp$ denotes rejection. Fix $c \geq 0$ and a measurable loss $L : \mathcal{D} \times \mathcal{A} \to [0, \infty]$; the acquisition cost is $c\tau$ and the terminal loss is $L(g(\omega), D_\tau)$. No history-based decision can distinguish states with equal path laws.

2.1 Declared and actual observation laws

Let $\mu_n(\mathrm{d}y \mid h, u)$ be a sigma-finite kernel, measurable in (h, u) and fixed before Y_n is observed. The declared kernel has density $q_{\omega,n} = \mathrm{d}Q_{\omega,n}/\mathrm{d}\mu_n$. Write $Q^{\star}_{\omega,n}$ for the actual conditional kernel and $p^{\star}_{\omega,n} = \mathrm{d}Q^{\star}_{\omega,n}/\mathrm{d}\mu_n$ whenever the derivative exists. Alternating $Q^{\star}_{\omega,n}$ with π defines the actual path law $P^{\star,\pi}_{\omega,N}$. The algorithm evaluates q; validity instead rests on the predictable conditional envelope of Section 5. When $p^{\star} = q$, the ordinary likelihood-ratio process is valid. An omitted systematic shift may otherwise force repeated nominal factors across the boundary, so the envelope must hold before Y_n is observed.

2.2 Finite-hypothesis DCE-MRI control

The numerical experiment specialises ω to four hypotheses formed by two synthetic enhancement patterns and two nuisance locations. A causal mask policy drives a scalar DCE state-space model with multicoil complex Gaussian observations. Section 7 gives the state and observation equations; Appendix E fixes every implemented array and parameter. Its linear response is the declared Gaussian experiment and its nonlinear SPGR response a misspecification stress test.

2.3 Chronological stored-image prefixes

The limited retrospective experiment begins with a recorded series $Z_{1:N}$ in a fixed chronological order. Its filtration is $\mathcal{G}_n = \sigma(Z_1, \dots, Z_n)$. At predeclared gates, a classifier fitted on the training split and calibrated on a disjoint calibration split produces a label set from features measurable with respect to $\mathcal{G}_n$. The rule stops at the first singleton set and otherwise rejects at the final gate. Training subjects, calibration subjects, and test subjects are kept disjoint.

Replay concerns only the recorded protocol and requires documented prefix-only reconstruction; Section 6 states the audit conditions.

3 Controlled observation fibres

We use the finite-horizon standard Borel experiment of Section 2. Its common policy kernels and state-specific observation kernels determine the history laws $P^{\pi}_{\omega,n}$. The task map $g : \Omega \to \mathcal{D}$ is measurable and $\mathcal{D}$ is finite.

Definition 3.1 (Controlled observational equivalence) For a declared family $\mathfrak{P}$ of causal policies and $0 \le n \le N$, set

$$\omega \sim_{\mathfrak{P},n} \omega' \quad \Longleftrightarrow \quad P^{\pi}_{\omega,n} = P^{\pi}_{\omega',n} \quad \text{for every } \pi \in \mathfrak{P}.$$

The equivalence class

$$\mathfrak{F}_{\mathfrak{P},n}(\omega) = \{\omega' : \omega' \sim_{\mathfrak{P},n} \omega\}$$

is the controlled observation fibre through ω.

The horizon-zero convention in Section 2 gives $\mathfrak{F}_{\mathfrak{P},0}(\omega) = \Omega$.

The definition is a policy-indexed noisy-history quotient. It is related to, but distinct from, comparison of experiments, deterministic-input observability, and observation-law observability in nonlinear filtering [11, 12, 17].

Proposition 3.2 (Order of controlled fibres) *The relation in Definition 3.1 is an equivalence relation. If* $\mathfrak{P}_1 \subseteq \mathfrak{P}_2$, *then*

$$\mathfrak{F}_{\mathfrak{P}_2,n}(\omega) \subseteq \mathfrak{F}_{\mathfrak{P}_1,n}(\omega).$$

If the n*-stage experiment is the marginal of the* $(n+1)$*-stage experiment, then*

$$\mathfrak{F}_{\mathfrak{P},n+1}(\omega) \subseteq \mathfrak{F}_{\mathfrak{P},n}(\omega).$$

Definition 3.3 (Task identification) The task g is identified by $(\mathfrak{P}, n)$ if it is constant on every controlled observation fibre:

$$\omega \sim_{\mathfrak{P},n} \omega' \quad \Longrightarrow \quad g(\omega) = g(\omega').$$

Write

$$\mathcal{E}_{\mathfrak{P},n}(\omega) = (P^{\pi}_{\omega,n})_{\pi \in \mathfrak{P}}$$

for the complete experiment-family map.

Theorem 3.4 (Task descent) *The task* g *is identified by* $(\mathfrak{P}, n)$ *if and only if there is a unique set map*

$$\bar{g} : \mathcal{E}_{\mathfrak{P},n}(\Omega) \longrightarrow \mathcal{D} \quad \text{such that} \quad g = \bar{g} \circ \mathcal{E}_{\mathfrak{P},n}.$$

Endow the image $\mathcal{E}_{\mathfrak{P},n}(\Omega)$ *with a standard-Borel sigma-field. If* $\mathcal{E}_{\mathfrak{P},n}$ *is a Borel quotient map onto this image, then measurable* g *has measurable descent* $\bar{g}$.

Theorem 3.5 (Policy-uniform no-go) *Let* $\omega_0, \omega_1 \in \Omega$. *Suppose that, for every* $1 \le i \le n$, *every* $h \in (\mathcal{U} \times \mathcal{Y})^{i-1}$, *and every* $u \in \mathcal{U}$,

$$Q_{\omega_0,i}(\cdot \mid h, u) = Q_{\omega_1,i}(\cdot \mid h, u). \tag{3}$$

Then $P^{\pi}_{\omega_0,n} = P^{\pi}_{\omega_1,n}$ *for every causal policy* π. *If* $g(\omega_0) \neq g(\omega_1)$, *no decision based on the declared history identifies* g *on a model containing both states.*

3.1 Local observation Fitting diagnostic

Let $H : (\mathbb{R}^d, \theta_0) \to (\mathbb{R}^q, H(\theta_0))$ be a real analytic observation germ and let $A = \mathbb{R}\{x_1 - \theta_{0,1}, \dots, x_d - \theta_{0,d}\}$ be the local ring of convergent real power-series germs at θ_0. For an ideal $I \subseteq A$, write $V_{\mathbb{R}}(I)$ for its real zero germ at θ_0. For a matrix B over A, write $I_k(B)$ for the ideal generated by its $k \times k$

minors, with $I_0(B) = A$ and $I_k(B) = (0)$ for $k > \min\{d, q\}$ when B is $q \times d$. We use the standard presentation convention for Fitting ideals [18, 19].

Definition 3.6 (Observation Fitting module) For the Jacobian J_H, define

$$\mathcal{M}_H = \operatorname{coker}(J_H : A^d \to A^q),$$
$$\operatorname{Fitt}_j(\mathcal{M}_H) = I_{q-j}(J_H), \qquad 0 \le j < q. \tag{4}$$

For $j \ge q$, set $\operatorname{Fitt}_j(\mathcal{M}_H) = A$. The real rank-drop germ $\{\theta : \operatorname{rank} J_H(\theta) \le r\}$ is $V_{\mathbb{R}}(I_{r+1}(J_H))$.

For a fixed endpoint index u, let s_u be a positive integer and $L_u : \mathbb{R}^d \to \mathbb{R}^{s_u}$ an additional analytic endpoint. Write $D(H, L_u)$ for the derivative of the stacked map (H, L_u) and set

$$\rho(u \mid H; \theta_0) = \operatorname{rank} D(H, L_u)(\theta_0) - \operatorname{rank} DH(\theta_0). \tag{5}$$

Theorem 3.7 (Constant-rank local fibre reduction) *At every θ_0,*

$$\rho(u \mid H; \theta_0) = \operatorname{rank}\Big(DL_u(\theta_0)|_{\ker DH(\theta_0)}\Big). \tag{6}$$

If DH and $D(H, L_u)$ have constant ranks r and r' near θ_0, their local fibres through θ_0 are embedded submanifolds of dimensions $d-r$ and $d-r'$. The endpoint therefore reduces local fibre dimension by exactly $\rho = r' - r$.

For $H(x, y) = x$ and $L(x, y) = y$, the increment is one and the local line fibre becomes a point. The rank identity is pointwise; the dimension conclusion is not. At zero, $H(x) = x^2$ and $L(x) = x$ give $\rho = 1$, although $H^{-1}(0)$ and $(H, L)^{-1}(0, 0)$ are both locally $\{0\}$. For $H(x, y) = xy$, the rank-drop locus is the origin, whereas the zero fibre is the singular union of the coordinate axes. The Fitting flag therefore locates rank loss but neither resolves fibre branches and multiplicities nor determines a global controlled fibre or statistical information; a derivative at zero, for example, does not record the separation of $N(0, \sigma^2)$ from $N(\delta^2, \sigma^2)$.

Corollary 3.8 (Boundary of retrospective evidence) *Let π^0 be the deterministic open-loop policy reproducing the action sequence recorded in a retrospective data set. For every family $\mathfrak{P}$ containing π^0,*

$$\mathfrak{F}_{\mathfrak{P},n}(\omega) \subseteq \mathfrak{F}_{\{\pi^0\},n}(\omega).$$

Thus a replay of one schedule can establish properties of that schedule, but cannot establish equality or separation under unrecorded counterfactual controls.

Example 3.9 (A separating and a colliding experiment) If $Y \mid (U = u, \theta) \sim N(u\theta, 1)$ and $u = 1$ is admissible, distinct values of θ have distinct one-step laws. By contrast, if every admissible control gives $N(a_u\theta^2, 1)$, then θ and $-\theta$ lie in the same controlled fibre at every horizon. The sign task is therefore unidentified, however adaptive the policy.

Fibres depend on full laws: mean equality may be defeated by covariance, and local rank by remote symmetry. An unidentified nuisance parameter is harmless when g is constant on the resulting fibre.

4 Stopped information and singular separation

The next result concerns probability kernels, not only their means. For two states ω_0, ω_1, write

$$k_i(h, u) = \mathrm{KL}(Q_{\omega_0,i}(\,\cdot \mid h, u) \| Q_{\omega_1,i}(\,\cdot \mid h, u)).$$

Let $\tau \le N$ be a stopping time. For a state ω, let $P_\omega^{\pi,\tau}$ denote the law of the stopped transcript $(\tau, U_1, Y_1, \ldots, U_\tau, Y_\tau)$, and put $P_j^{\pi,\tau} = P_{\omega_j}^{\pi,\tau}$. This transcript takes values in the finite disjoint union

$$\begin{aligned} \mathsf{H}_N^\tau &= \bigsqcup_{t=0}^{N} \mathsf{H}_t, \\ \mathsf{H}_t &= \{t\} \times (\mathcal{U} \times \mathcal{Y})^t. \end{aligned} \tag{7}$$

Theorem 4.1 (Stopped conditional-information identity) *Suppose that, for each i, there is a jointly measurable nonnegative function $R_i(y \mid h, u)$ such that, for almost every conditioning pair under the law of (H_{i-1}, U_i) induced by $P_{\omega_0,N}^\pi$,*

$$Q_{\omega_0,i}(\mathrm{d}y \mid h, u) = R_i(y \mid h, u) Q_{\omega_1,i}(\mathrm{d}y \mid h, u) \tag{8}$$

and suppose that the policy kernel is common to the two states and that

$$\mathbb{E}_{\omega_0}^\pi \sum_{i=1}^{\tau} k_i(H_{i-1}, U_i) < \infty. \tag{9}$$

Then

$$\mathrm{KL}(P_0^{\pi,\tau} \| P_1^{\pi,\tau}) = \mathbb{E}_{\omega_0}^\pi \sum_{i=1}^{\tau} k_i(H_{i-1}, U_i). \tag{10}$$

The identity holds for deterministic or randomised causal policies.

The density hypothesis holds under a common measurable dominating kernel and covers the Gaussian model below; pointwise choices of Radon–Nikodym derivatives do not suffice. The common policy cancels from the likelihood ratio. Bounded N and (9) reduce the proof to a finite sum; no limiting argument is used.

Corollary 4.2 (Predictable Gaussian acquisitions) *Let $\mathcal{Y} = \mathbb{R}^d$ for a fixed positive integer d. Suppose that, conditionally on $(H_{i-1}, U_i) = (h, u)$, state j gives $N(m_{j,i}(h,u), \Sigma_{j,i}(h,u))$ on $\mathbb{R}^d$, where the displayed means and positive-definite covariances are measurable in (h,u). Suppressing these arguments, the summand in (10) is*

$$k_i(h,u) = \frac{1}{2}\{\operatorname{tr}(\Sigma_{1,i}^{-1}\Sigma_{0,i}) - d + (m_{1,i} - m_{0,i})^{\mathsf{T}}\Sigma_{1,i}^{-1}(m_{1,i} - m_{0,i}) + \log(\det \Sigma_{1,i}/\det \Sigma_{0,i})\}. \tag{11}$$

If $\Sigma_{0,i} = \Sigma_{1,i} = \Sigma_i$, this reduces to one half of the squared Σ_i^{-1}-norm of the mean difference.

Predictive covariance may depend on the structural hypothesis, so a Fitting analysis of the mean map alone does not determine information [18, 19]. For example, $N(0, 1+\theta)$ has zero mean Jacobian but divergence of order δ^2 between $\theta = 0$ and δ. Mean equality gives a collision only when the remaining kernel parameters also agree.

Definition 4.3 (Uniform controlled information order) For this local family the finite horizon may depend on δ; denote it by N_δ. For $(\omega_0, \omega_\delta)_{0<\delta\le\delta_0}$ with opposite task values, put

$$k_{i,\delta}(h,u) = \mathrm{KL}\big(Q_{\omega_0,i}(\cdot \mid h,u) \| Q_{\omega_\delta,i}(\cdot \mid h,u)\big).$$

The pair has controlled information order at least the positive integer m if a constant $0 < C < \infty$, independent of δ, satisfies

$$\sup_{1\le i\le N_\delta} \sup_{h,u} k_{i,\delta}(h,u) \le C\delta^{2m}. \tag{12}$$

The second supremum is over admissible histories and controls. A bound only for the control ultimately selected is not policy-uniform. Because finite relative entropy implies absolute continuity and $\tau \le N$, (12) also supplies the density and integrability hypotheses needed to apply Theorem 4.1 in the lower-bound proof below.

For $p, q \in (0,1)$, put

$$\mathrm{kl}(p,q) = p\log(p/q) + (1-p)\log\{(1-p)/(1-q)\}.$$

The reduction of a complete transcript to this binary event is the classical two-point method [20].

Theorem 4.4 (Singular sequential lower bound) *Assume (12). For each δ, let a common causal policy π_δ, a stopping time $\tau_\delta \le N_\delta$, and a forced binary decision $\widehat{d}_\delta$, measurable with respect to the stopped transcript in $\mathsf{H}_{N_\delta}^{\tau}$, satisfy, for some $0 < \alpha < 1/2$,*

$$P_{\omega_0}^{\pi_\delta,\tau_\delta}\{\widehat{d}_\delta \ne g(\omega_0)\} \le \alpha,$$
$$P_{\omega_\delta}^{\pi_\delta,\tau_\delta}\{\widehat{d}_\delta \ne g(\omega_\delta)\} \le \alpha.$$

Then

$$\mathbb{E}_{\omega_0}^{\pi_\delta}\tau_\delta \ge \frac{\mathrm{kl}(1-\alpha,\alpha)}{C\delta^{2m}}. \tag{13}$$

Because $\tau_\delta \le N_\delta$, the assumptions cannot hold at a value of δ for which the right-hand side of (13) exceeds N_δ.

Example 4.5 (Sharp order and exact collision) For independent $N(\theta^m, \sigma^2)$ observations and the pair $(0, \delta)$, a fixed-sample likelihood-ratio test with $n = n_\delta \asymp \delta^{-2m}$ attains the lower-bound order. If every admissible kernel depends on θ only through θ^2, the pair $(\delta, -\delta)$ has zero one-step divergence and remains a controlled collision at every prefix.

Theorem 4.4 is a two-state lower bound for a forced decision. It gives neither a composite minimax theorem nor a lower bound for a procedure allowed to reject. An upper-order statement requires an explicit test and an attainable control; it does not follow from a local rank flag.

5 Robust prefix-fibre stopping

Nominal likelihood ratios need not remain martingales when the denominator model is misspecified. We therefore impose a conditional moment envelope, rather than deducing optional-stopping validity from an average goodness-of-fit measure. Let Θ be standard Borel and let the measurable map $\vartheta : \Omega \to \Theta$ retain the component of the structural state used by a nominal model. The remaining component of ω may include a nuisance kernel or model discrepancy, and g remains defined on the full state.

Write $Q^\star_{\omega,n}(\mathrm{d}y \mid h,u)$ for the actual conditional kernel associated with $p^\star_{\omega,n}$ in Section 2.1, and $P^{\star,\pi}_{\omega,N}$ for the path law obtained by alternating it with the policy kernels. For every Borel $A \subseteq \mathcal{Y}$, assume that $(\omega,h,u) \mapsto Q^\star_{\omega,n}(A \mid h,u)$ is measurable.

For every $\theta \in \Theta$, stage n, history h, and control u, let $\mathcal{K}_{\theta,n}(h,u)$ be a nonempty set of probability measures on $\mathcal{Y}$. We require

$$Q^\star_{\omega,n}(\cdot \mid h,u) \in \mathcal{K}_{\vartheta(\omega),n}(h,u) \tag{14}$$

for every $\omega \in \Omega$, $h \in (\mathcal{U}\times\mathcal{Y})^{n-1}$, and $u \in \mathcal{U}$. Choose a nonnegative score $f_{\theta,n}(y \mid h,u)$, jointly measurable in its displayed variables, and a finite, strictly positive bound $\kappa_{\theta,n}(h,u)$, jointly measurable in (θ,h,u) and fixed after U_n is selected but before Y_n is observed, such that

$$\sup_{P\in\mathcal{K}_{\theta,n}(h,u)} \int f_{\theta,n}(y \mid h,u)P(\mathrm{d}y) \le \kappa_{\theta,n}(h,u). \tag{15}$$

The score and its bound may depend on the preceding history and current control, but not on the current observation. Define

$$E_n(\theta) = \prod_{i=1}^{n} \frac{f_{\theta,i}(Y_i \mid H_{i-1},U_i)}{\kappa_{\theta,i}(H_{i-1},U_i)}, \qquad E_0(\theta) = 1. \tag{16}$$

In the exact nominal case, let $Q^0_{\theta,n}$ have density $q_{\theta,n}$ with respect to μ_n. Let $r_n(\cdot \mid h,u)$ be a conditional probability density with respect to μ_n, selected measurably before Y_n, and suppose that its measure is absolutely continuous with respect to $Q^0_{\theta,n}(\cdot \mid h,u)$. Set

$$f_{\theta,n} = \begin{cases} r_n/q_{\theta,n}, & q_{\theta,n} > 0, \\ 0, & q_{\theta,n} = 0, \end{cases} \qquad \kappa_{\theta,n} = 1,$$

where the displayed densities retain their arguments $(y \mid h,u)$. Then $\int f_{\theta,n}\,\mathrm{d}Q^0_{\theta,n} = 1$. The process then reduces to the ordinary likelihood-ratio martingale. Under misspecification, $\kappa = 1$ is invalid unless the same conditional moment identity has been established for the actual kernel.

Lemma 5.1 (Conditional-envelope e-process) *Under (14)–(15), $(E_n(\vartheta(\omega)))_{n=0}^N$ is a nonnegative supermartingale under $P^{\star,\pi}_{\omega,N}$ for every state ω and every causal policy π. Consequently,*

$$P^{\star,\pi}_{\omega,N}\left\{\max_{0\le n\le N} E_n(\vartheta(\omega)) \ge \alpha^{-1}\right\} \le \alpha \tag{17}$$

for every $0<\alpha<1$.

Set

$$C_n = \{\omega : E_n(\vartheta(\omega)) < \alpha^{-1}\},$$
$$S_n = \mathrm{sat}_{\mathfrak{P},n}(C_n) = \bigcup_{\omega\in C_n} \mathfrak{F}_{\mathfrak{P},n}(\omega). \tag{18}$$

Exact saturation may be difficult in a continuous model. Let $\widehat{F}_n$ be a random outer approximation. We assume directly that, for every n, the events $\{S_n \subseteq \widehat{F}_n\}$, $\{\mathfrak{F}_{\mathfrak{P},n}(\omega) \subseteq \widehat{F}_n\}$ for fixed ω, and $\{g(\widehat{F}_n) = \{d\}\}$ for $d \in \mathcal{D}$ are $\mathcal{F}_n$-measurable. These assumptions make the construction a causal prefix rule and the stopped decision below measurable; a measurable graph alone does not ensure them for arbitrary random subsets of a standard Borel space. Suppose, for some $0 \le \beta < 1-\alpha$, that

$$\inf_{\omega,\pi} P^{\star,\pi}_{\omega,N}\{S_n \subseteq \widehat{F}_n \text{ for every } 0\le n\le N\} \ge 1-\beta. \tag{19}$$

The infimum is over the declared states and causal policies. A deterministic certified outer approximation has $\beta = 0$; an inner approximation does not satisfy (19).

For a set $F \subseteq \Omega$, write $g(F) = \{g(\omega) : \omega \in F\}$. At a stopping time $\tau \le N$, define

$$\delta(\widehat{F}_\tau) = \begin{cases} d, & g(\widehat{F}_\tau) = \{d\}, \\ \perp, & \text{otherwise.} \end{cases} \tag{20}$$

Rejection expresses the usual error–coverage trade-off [21].

Theorem 5.2 (Robust prefix-fibre stopping) *Suppose the preceding measurability conditions and (14), (15), and (19) hold. Then, for every state ω_0, every causal policy π, and every stopping time $\tau \le N$,*

$$P^{\star,\pi}_{\omega_0,N}\{\mathfrak{F}_{\mathfrak{P},n}(\omega_0) \subseteq \widehat{F}_n \text{ for all } n \le N\} \ge 1-\alpha-\beta, \tag{21}$$

and

$$P^{\star,\pi}_{\omega_0,N}\{\delta(\widehat{F}_\tau) \ne g(\omega_0),\ \delta(\widehat{F}_\tau) \ne \perp\} \le \alpha+\beta. \tag{22}$$

We use $\mathrm{TV}(P,Q) = \sup_{A\in\mathcal{B}(\mathcal{Y})} |P(A)-Q(A)|$.

Corollary 5.3 (A conditional total-variation envelope) *Let $Q^0_{\theta,n}(\cdot \mid h,u)$ be nominal and suppose that every $P \in \mathcal{K}_{\theta,n}(h,u)$ satisfies*

$$\mathrm{TV}\Big(P, Q^0_{\theta,n}(\cdot \mid h,u)\Big) \le \varepsilon_{\theta,n}(h,u).$$

If $0 \le f_{\theta,n}(\cdot \mid h,u) \le M_{\theta,n}(h,u) < \infty$ and

$$a_{\theta,n}(h,u) = \int f_{\theta,n}(y \mid h,u) Q^0_{\theta,n}(\mathrm{d}y \mid h,u),$$

then

$$\kappa_{\theta,n}(h,u) = a_{\theta,n}(h,u) + M_{\theta,n}(h,u)\varepsilon_{\theta,n}(h,u) \quad (23)$$

is valid in (15), provided it is positive. A conditional bound $\mathrm{KL}\{P \| Q^0_{\theta,n}(\cdot \mid h,u)\} \le \eta_{\theta,n}(h,u)$ permits $\varepsilon_{\theta,n}(h,u) = \sqrt{\eta_{\theta,n}(h,u)/2}$.

For $0 \le \varepsilon \le 1$ and $M < \infty$, the contamination class $P = (1-\varepsilon)Q^0 + \varepsilon R$ permits the sharper choice $\kappa = (1-\varepsilon)a + \varepsilon M$ whenever this quantity is positive. Thus a bounded, clipped score pays for the declared contamination in each e-factor.

Proposition 5.4 (Failure without clipping) *There are nominal laws q_δ, true laws p_δ, and numerator laws r_δ on a two-point space such that*

$$\begin{aligned} &\mathrm{TV}(p_\delta, q_\delta) \longrightarrow 0, \\ &\mathbb{E}_{q_\delta}(r_\delta/q_\delta) = 1, \\ &\mathbb{E}_{p_\delta}(r_\delta/q_\delta) \longrightarrow \infty. \end{aligned}$$

Hence total-variation proximity alone does not robustify an unbounded likelihood ratio.

The bound is joint, not conditional on reporting; conditional error may equal one when reporting is rare. A broad kernel class may force rejection. Acquisition saving additionally requires prefix-measurability of τ, preprocessing, and all scores. The fibres belong to the declared experiment, whereas coverage is evaluated under the actual law; robustness of the equivalence relation requires adjoining the discrepancy kernel to the state.

6 Retrospective causal-prefix replay

Let W and each Z_t take values in standard Borel spaces. Let $Z_{1:T}$ be a stored chronological image series acquired under one fixed protocol, and let $1 \le t_1 < \cdots < t_G \le T$ be gates fixed before the test subject is examined. Write W for all objects fixed independently of that subject: the training and calibration records, the gates, any fitted maps, and any external random seed. The available information at frame t is

$$\mathcal{G}_t = \sigma(W, Z_1, \ldots, Z_t).$$

At gate t_j, a prefix rule returns a label set $\Gamma_j \subseteq \mathcal{D}$ that is $\mathcal{G}_{t_j}$-measurable. Put

$$J = \inf\{j : |\Gamma_j| = 1\}, \qquad \inf \varnothing = \infty.$$

If $J \le G$, set $\tau = t_J$ and let $\widehat{D}$ be the unique element of Γ_J; otherwise set $\tau = T$ and $\widehat{D} = \perp$.

The word *causal* here refers to measurability and to invariance of the pre-stop data-generating kernels. It does not assert that a stored vendor image was reconstructed without future frames. Without further monotonicity assumptions, future leakage can make the first singleton occur earlier or later; the sign of the stopping-time bias is not identified.

Proposition 6.1 (Causal-prefix replay identity) *Consider a full acquisition regime and a prospective stop-only regime. Assume that (i) W has the same law in both regimes and, conditionally on W, both use the same initial law and the same conditional kernel for Z_t given $Z_{1:t-1}$ while acquisition remains active; (ii) stopping changes no earlier observation; and (iii) the feature maps, label sets, stopping time, and terminal report are adapted to $(\mathcal{G}_t)_{t \le T}$. Then the law of the stopped prefix, stopping time, and report obtained by replaying the full series equals their law in the prospective stop-only regime.*

The identity concerns termination of the recorded protocol. Retrospective undersampling can introduce further reconstruction and distributional artefacts [22]. In particular, the identity does not cover a policy that would choose a different phase-encoding line, pulse sequence, or contrast-injection schedule after seeing a prefix. Corollary 3.8 gives the corresponding fibre inclusion.

Finite gates permit a distribution-free label-set construction. Fix integers $m, G \ge 1$ and $0 < \alpha < 1$. Let $D_* \in \mathcal{D}$ be the task label of a future subject, R_{ij} the nonconformity score of calibration subject

i at gate j, and $R_{*j}(d)$ the candidate score under label $d \in \mathcal{D}$. Define

$$p_{j,d} = \frac{1 + \sum_{i=1}^{m} \mathbf{1}\{R_{ij} \geq R_{*j}(d)\}}{m+1},$$
$$\Gamma_j = \{d \in \mathcal{D} : p_{j,d} > \alpha/G\}. \tag{24}$$

The use of $\geq$ makes ties conservative.

Proposition 6.2 (Finite-gate conformal rejection) *Let $\mathcal{T}$ be generated by the training data and external randomisation. Conditionally on $\mathcal{T}$, suppose that the m calibration subjects and one future subject are exchangeable at the subject level, and that the gates and gate-specific score functions are $\mathcal{T}$-measurable. Thus they are not selected from the calibration or future-subject records. Then*

$$\mathbb{P}\{D_* \in \Gamma_j \text{ for every } j\} \geq 1 - \alpha. \tag{25}$$

If the rule reports the unique element of the first singleton Γ_j and otherwise reports $\perp$, then

$$\mathbb{P}\{\widehat{D} \neq D_*,\ \widehat{D} \neq \perp\} \leq \alpha. \tag{26}$$

No independence across gates is required.

This is marginal finite-sample coverage for one exchangeable future subject, as in conformal classification [15]. It is not class-conditional or site-conditional coverage. Reusing patients across splits, selecting gates after viewing calibration outcomes, or a temporal or site shift invalidates the stated argument; extensions beyond exchangeability require additional structure [23].

The distributed program accepts NPZ arrays or conventional single-frame MR DICOM with a separate subject-level split file. At each gate it uses only the stored magnitude prefix, fits a diagonal-Gaussian classifier on the training split, and calibrates (24) on the disjoint calibration split. The DICOM loader enforces subject, series, geometry, and ordinal-time consistency; the temporal tag, start, and phase count are command-line inputs. It neither infers clock time nor certifies prefix-only upstream processing. Online Resource 1 gives the complete input contract.

The synthetic self-test verifies that perturbing future frames leaves an earlier prefix feature unchanged. Reported visited-gate coverage is descriptive and differs from (25). No patient series was analysed; prospective use requires latency measurement, unchanged pre-stop acquisition, external validation, and safety analysis.

7 Reduced DCE-MRI phantom experiment

Four structural hypotheses $h \in \{0,1,2,3\}$ combine two synthetic spatio-kinetic phenotypes with two nuisance locations. Their task labels are $\ell_h = \lfloor h/2 \rfloor$; all arrays are synthetic and estimate no clinical performance. For hypothesis h, the scalar concentration follows

$$\begin{aligned} c_{h,0} &\sim N(0, \sigma_0^2), \\ c_{h,t} &= a_h c_{h,t-1} + b_h \bar{C}_{p,t-1} + w_t, \\ w_t &\sim N(0, \sigma_w^2). \end{aligned} \tag{27}$$

where $a_h = \exp(-k_{ep,h}\Delta t)$, $b_h = v_{e,h}(1 - a_h)$, $k_{ep,h} = K_h^{\text{trans}}/v_{e,h}$, and $\bar{C}_{p,j} = 0$ for $j < 4$ while $\bar{C}_{p,j} = C_p\{(j-4)\Delta t\}$ for $j \geq 4$ [13]. The input C_p is the population function of Parker et al. [14]. The separation of evolution and complex k-space observation equations has a precedent in dynamic radial MRI [24]; our state is a scalar concentration, not a reconstructed image. The initial draw and every transition in (27) match the filtering likelihood.

For Cartesian mask a, let S_a retain its six phase-encode rows. For coil $c \in \{0,1,2,3\}$, the observation is

$$\begin{aligned} Y_{t,a,c} &= S_a \mathcal{F}\left[s_c\{B_c + A(c_{h,t})\Phi_h\}\right] + \eta_{t,a,c}, \\ \eta_{t,a,c} &\sim \mathcal{CN}(0, \sigma^2 I). \end{aligned} \tag{28}$$

Under each hypothesis, $c_{h,0}$, the innovations $(w_t)_{t\geq 1}$, and the receiver noises are mutually independent; the latter two sequences are also independent across time, and receiver noise is independent across coils. Here B_c is the complex baseline defined in Appendix E. The exact experiment uses $A(c) = \beta c$, with $\beta = S'(0)$ for the SPGR response. The stress experiment generates $A(c) = S(c) - S(0)$ but retains the linear filter. Only the first experiment has the declared linear–Gaussian likelihood and its e-process guarantee. Appendix E specifies every array, parameter, and density.

There are 36 frames at five-second spacing, four coils, 7 candidate masks, and 6 phase-encode lines per selected mask. The fixed policy cycles through masks. The variance-only policy minimises the mean next-state variance under fixed uniform hypothesis weights; in this model

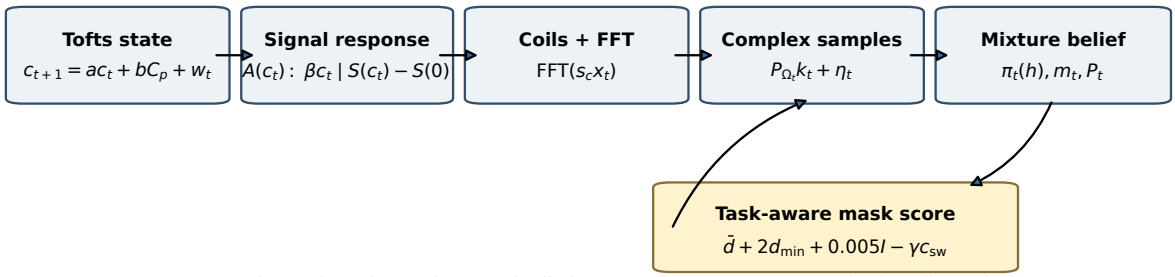

Fig. 1 Implemented sequential pipeline. The control selects one of seven six-line Cartesian masks from the preceding history. The displayed task score is a separation surrogate, not a Bellman solution or a one-step expected terminal Bayes loss. State noise and proper complex receiver noise are distinct

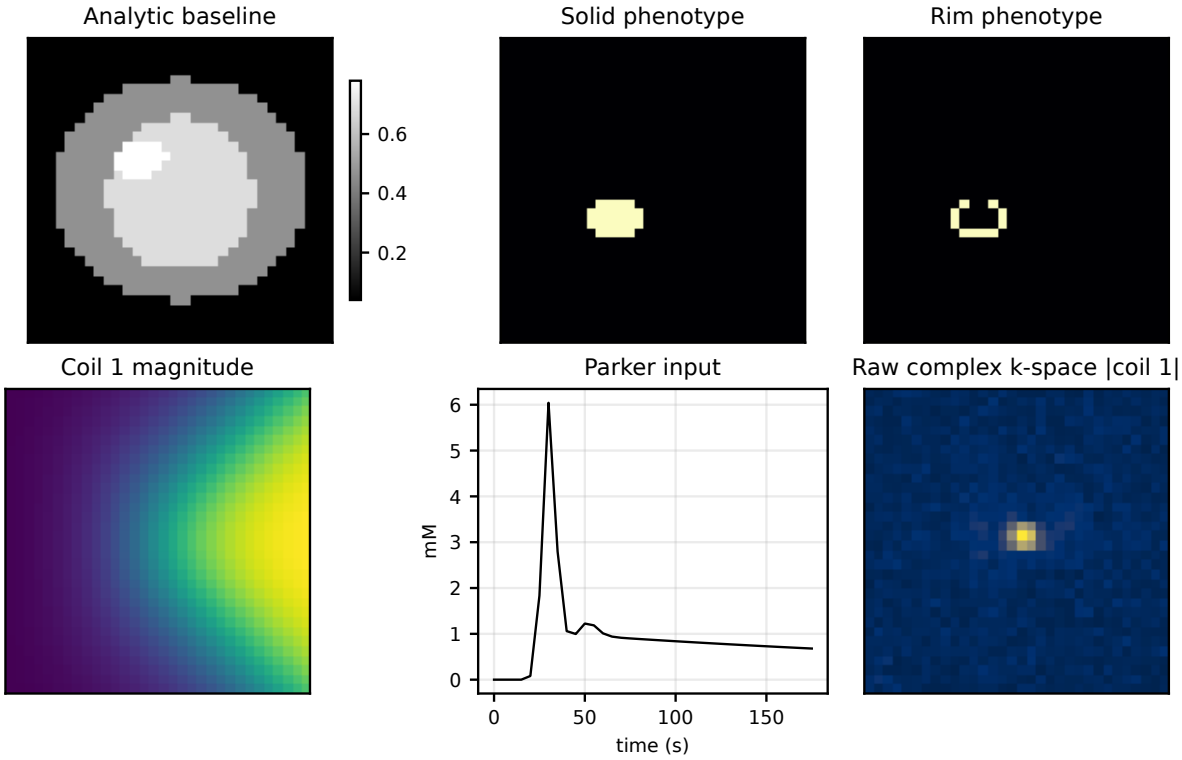

Fig. 2 Analytic 32×32 digital experiment: baseline magnitude, solid and equal-DC rim bases, one coil sensitivity, the Parker input, and one coil of the generated complex k-space. Displayed magnitudes do not replace the complex observations in (28)

it is observation independent. The task-aware policy maximises the causal separation score in Appendix E. The stopping guarantee does not depend on optimality of this heuristic. For the controlled-fibre relation, the declared family contains all causal policies on these seven masks and therefore every deterministic one-stage mask policy. The three implemented policies are members of this larger family.

For each hypothesis, the filter supplies the exact one-step predictive density $q_{h,t}$ in the linear experiment. With the predictable Bayes-mixture

density r_t, the implementation updates

$$E_t(h) = \prod_{i=1}^{t} \frac{r_i(Y_i)}{q_{h,i}(Y_i)},$$
$$C_t = \mathrm{sat}\{h : E_t(h) < 1/0.10\}. \qquad (29)$$

The saturation uses a conservative outer moment-signature relation described in Appendix E. All policies stop at the first prefix after at least eight frames have been acquired for which the set of phenotype labels represented in C_t is a singleton; otherwise they reject at the full horizon.

7.1 Exact linear experiment

The same truth and receiver-noise arrays were reused across policies in each of 1000 runs. Each of the four hypotheses occurs exactly 250 times in every policy–scenario cell. Table 1 reports Monte Carlo means; Appendix E defines the metrics and pointwise intervals. Hypothesis-specific summaries and paired contrasts are distributed in machine-readable form. For the task-aware heuristic, hypothesis-specific decision coverage ranged from 99.2–99.6%, and selective risk ranged from 0.0–3.2%.

For the task-aware policy, the true finite hypothesis remained in the confidence set through every acquired prefix in 96.1% of runs (Wilson interval 94.7–97.1%). Its empirical wrong-and-non-rejected rate was 1.7% (Wilson interval 1.1–2.7%) at nominal $\alpha = 0.10$, and its descriptive 95% mixture-state interval had mean within-run acquired-frame coverage 95.0%. The first two quantities audit the exact sequential construction; the third is not a theorem-backed confidence sequence. These finite Monte Carlo frequencies neither prove nor replace the probability bounds in Section 5. The paired mean frame difference, task-aware heuristic minus fixed, was -1.36; the 95% interval was [-1.76, -0.97]. This contrast assesses the implemented heuristic, not an optimal policy.

7.2 Misspecification and collision controls

These synthetic controls isolate misspecification and exact collision; they do not estimate scanner or patient error.

Under nonlinear SPGR generation, task-aware decision coverage was 98.4%, with log loss 0.123 and complex dynamic NRMSE 0.480; variance-only decision coverage was 96.8%. These calculations deliberately use the misspecified linear filter.

The exact no-go pair has prior $(1/2, 1/2)$ and duplicates every controlled state-space and observation kernel under opposite labels. Its posterior remains 0.5 for both labels at every prefix, and the saturated rule rejected in 100.0% of runs, illustrating Theorem 3.5.

Figure 5 also reports a scalar model-error control with 100000 paths of length 40. An uncorrected likelihood ratio crossed the nominal threshold in 81.8% of paths. Division by the exact conditional factor $\kappa = 1.822$ reduced this to 5.5%. The δ^2 and δ^4 curves in the same figure are prescribed algebraic sequences used to check information order. They are not rates fitted to phantom observations.

The finite MRI model admits the analogous calculation. At each realised history and selected mask, its actual conditional law under h is

$$P^{\star}_{h,t} = (1-\varepsilon)Q_{h,t} + \varepsilon Q_{\bar{h},t},$$

where $\bar{h}$ has the opposite phenotype at the same nuisance location. For $f_{h,t} = \min\{r_t/q_{h,t}, M\}$, one has $\mathbb{E}_{Q_{h,t}} f_{h,t} \le 1$ and $\mathbb{E}_{Q_{\bar{h},t}} f_{h,t} \le M$. Hence the conditional envelope in Lemma 5.1 holds analytically with $\kappa = (1-\varepsilon) + \varepsilon M$. We used $\varepsilon = 0.10$, $M = 2.0$, and $\kappa = 1.100$ in 1000 balanced paired runs per policy. For the task-aware heuristic, the uncorrected clipped process rejected the true hypothesis at some prefix in 1.1% of runs. The κ-corrected process did so in 0.0% (Wilson interval 0.0–0.4%), while its empirical wrong-and-non-rejected rate was 0.0%. Its decision coverage was 35.7%.

8 Discussion

Controlled fibres, confidence fibres, and selective decisions answer distinct questions. The first is a quotient of path laws; the second retains states under a declared error model; the third adds loss, rejection, and acquisition cost. Local rank does not exclude remote collision, and no score identifies a task that varies within an exact fibre.

Table 1 Exact linear–Gaussian experiment. The labels are synthetic phenotypes. NRMSE is computed for the complex dynamic component through the stopping frame.

Policy	Decision coverage	Selective risk	Mean frames	Acquired	Log loss	NRMSE
Fixed	98.2%	2.0%	19.6	54.6%	0.102	0.462
Variance only	97.6%	1.6%	20.3	56.4%	0.085	0.430
Task-aware heuristic	99.3%	1.7%	18.3	50.8%	0.085	0.457

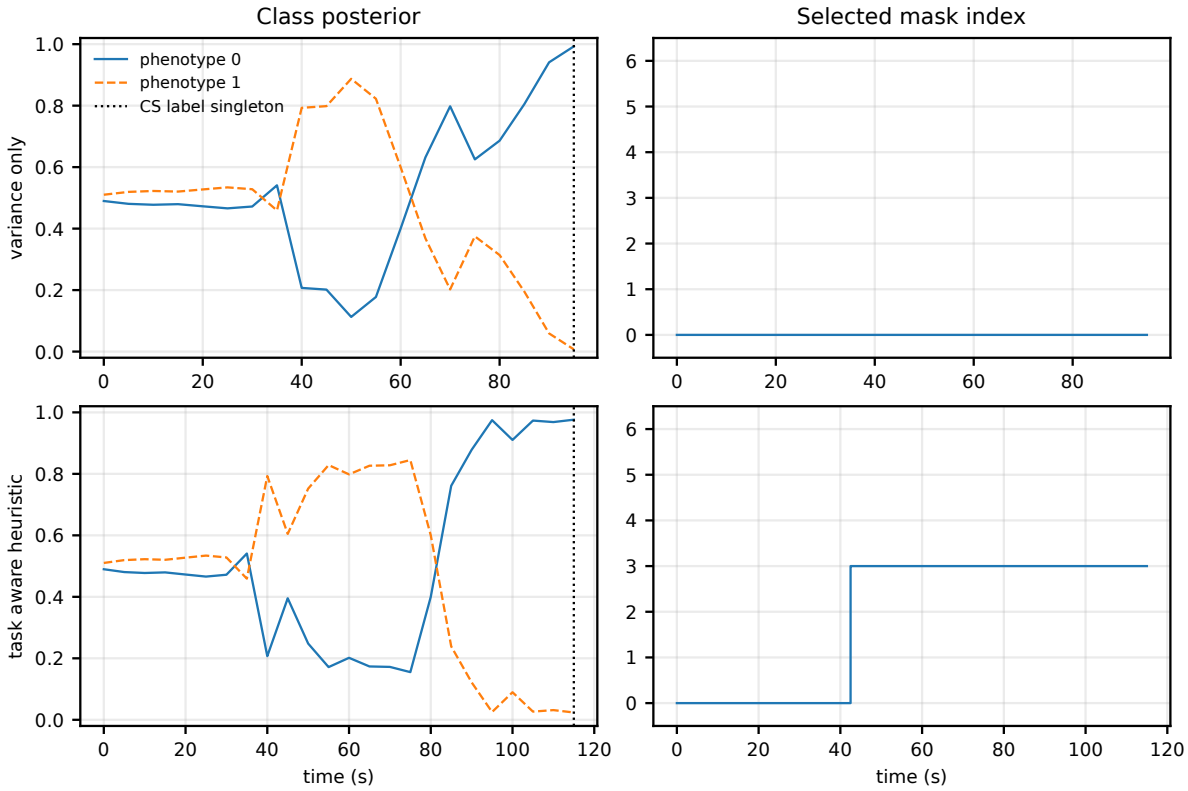

Fig. 3 One fixed-seed linear run. Class posteriors and selected mask indices are shown for the variance-only and task-aware policies. The dotted line marks the first label-singleton confidence set when it occurs. This is a path illustration, not a population comparison

The stopping guarantee follows from a one-step conditional envelope, as in anytime-valid e-value methods [9, 10]; average fit does not imply it. The parameter β records failure of outer-fibre containment. If an independently estimated envelope fails with probability β_{env}, a union bound gives $\alpha + \beta + \beta_{\mathrm{env}}$. Broad envelopes preserve validity at the cost of rejection. Higher-dimensional latent-state Gaussian and nonlinear extensions require new predictive calculations and, in the latter case, a new envelope.

The linear phantom checks the declared Gaussian experiment; the nonlinear and duplicated-law controls expose misspecification and collision. Stored-image replay instead requires chronological features, fixed splits and gates, and subject-level exchangeability [15, 23]; it does not recover unrecorded actions or remove reconstruction leakage [22]. External study therefore requires provenance-preserving raw data or a documented causal image pipeline, fixed endpoints, and prospective validation. Relevant infrastructure includes ISMRMRD and fastMRI Breast

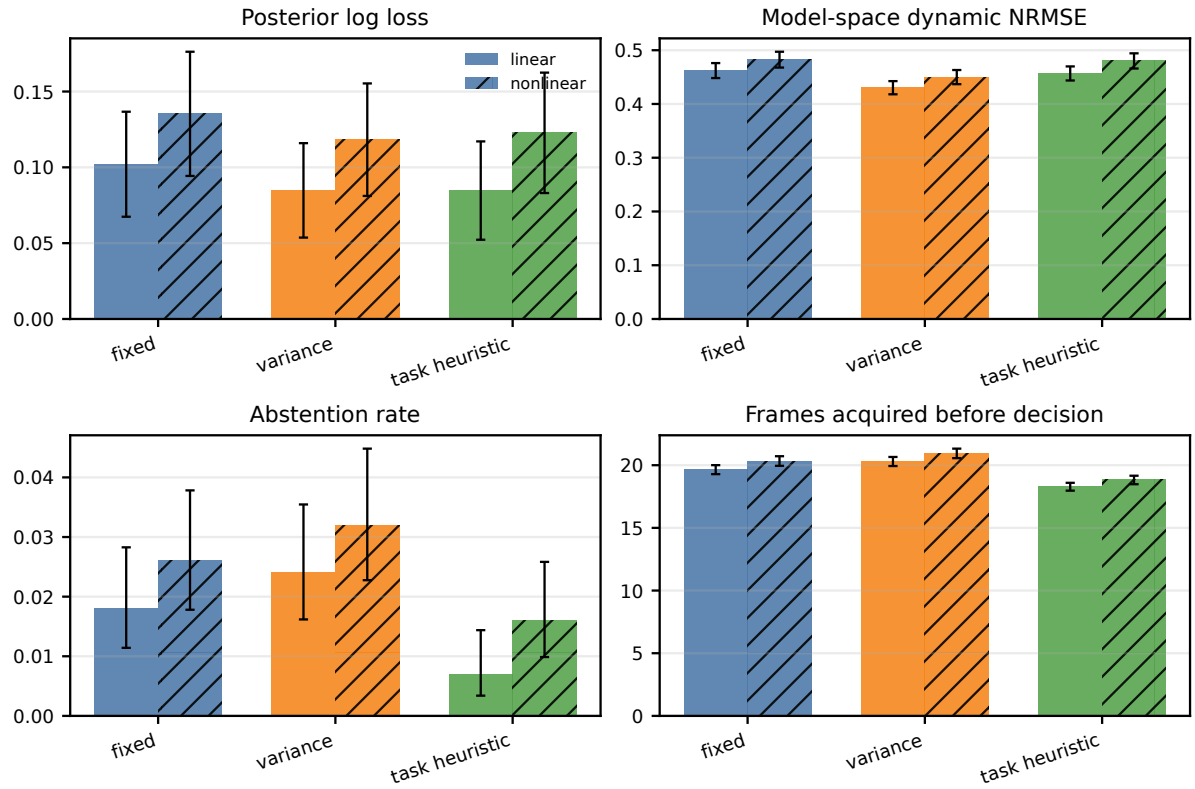

Fig. 4 Paired Monte Carlo summaries for the exact linear experiment and the nonlinear SPGR stress experiment. Error bars are two-sided 95% normal Monte Carlo intervals for continuous means and Wilson score intervals for abstention. The hatched nonlinear bars carry no exact e-process guarantee

[25, 26]; comparisons must share the action budget and reconstruction target [5, 7, 24, 27].

9 Conclusion

Controlled fibres locate the identification limit of sequential MRI. Predictably normalised evidence protects stopping under an explicit error envelope, while rejection records unresolved task heterogeneity. The phantom and replay illustrate these distinctions without conflating simulation, retrospective evidence, and prospective acquisition.

Acknowledgements

OpenAI Codex assisted with language editing, LaTeX, code, and package assembly; the author reviewed all outputs.

Declarations

Authorship. Takashi Matsuhisa is the sole author and accepts responsibility for the submitted work.

Ethics and clinical scope. The computations use synthetic arrays only. No human, animal, or patient data are included. The replay software contains no clinical data; no diagnostic claim is made.

Data Availability. The source code, fixed random seed, numerical summaries, figures, and a retrospective-prefix self-test are supplied as Online Resource 1. Clinical data, raw scanner data, and a prospective acquisition implementation are not supplied.

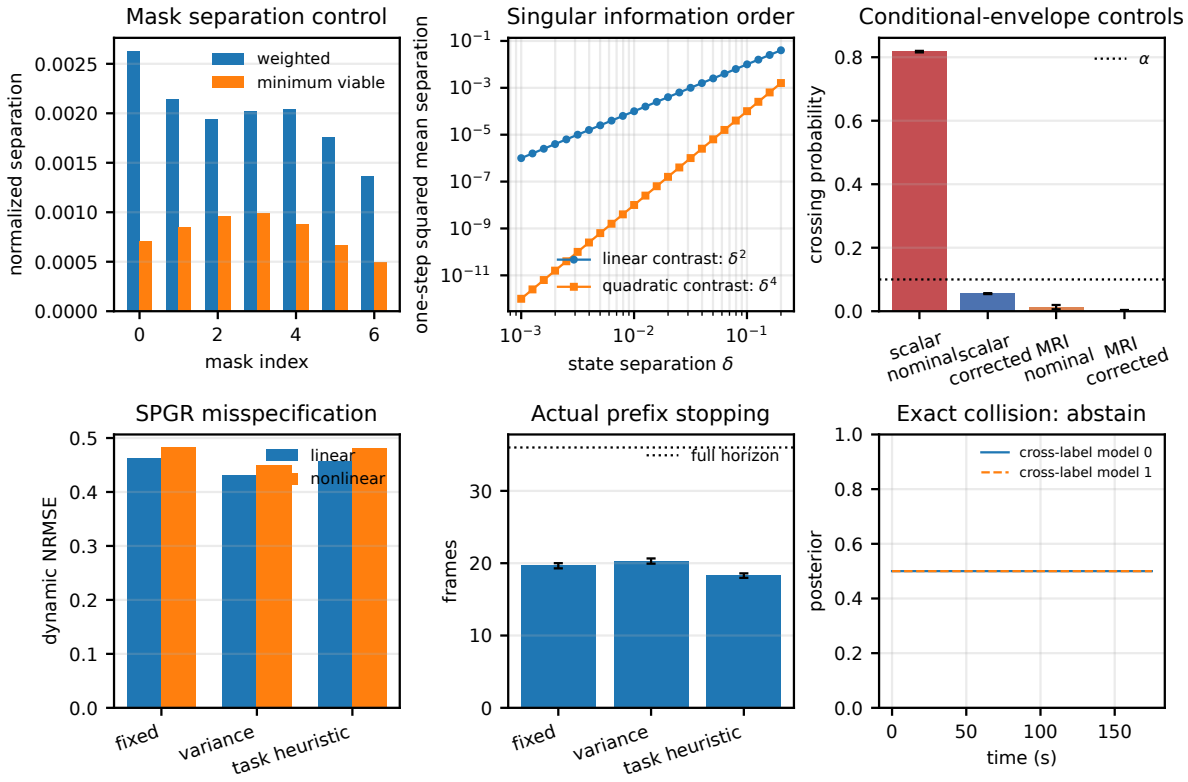

Fig. 5 Boundary checks. From left to right and top to bottom: mask-wise cross-label separation; prescribed linear δ^2 and quadratic δ^4 squared-separation orders; scalar and MRI-kernel nominal and κ-corrected crossing rates with Wilson intervals; model-space complex NRMSE under linear and nonlinear SPGR generation; policy-wise mean prefix lengths with 95% normal Monte Carlo intervals in the linear experiment; and the exact cross-label collision. The second panel is algebraic. The scalar pair in the third panel is separate from the MRI forward model

A Path laws and fibre proofs

Proof of Proposition 2.1 For a bounded measurable function f of the first pair, set

$$\int f \, \mathrm{d}P^{\pi}_{\omega,1} = \int \pi_1(\mathrm{d}u) \int Q_{\omega,1}(\mathrm{d}y \mid \varnothing, u) f(u, y).$$

Suppose that $P^{\pi}_{\omega,n-1}$ has been constructed. For bounded measurable f on histories of length n, define

$$\int f \, \mathrm{d}P^{\pi}_{\omega,n} = \int P^{\pi}_{\omega,n-1}(\mathrm{d}h) \int \pi_n(\mathrm{d}u \mid h) \times \int Q_{\omega,n}(\mathrm{d}y \mid h, u) f(h, u, y).$$

Kernel measurability makes each integral measurable in the preceding variables. The construction gives a probability measure and the prescribed conditional kernels. Any other such measure agrees on measurable rectangles at stage one and, by the displayed recursion, at each later stage. Induction proves existence and uniqueness. This is the finite Ionescu–Tulcea construction. □

Proof of Proposition 3.2 Equality of probability measures is reflexive, symmetric, and transitive; hence so is equality under every policy in $\mathfrak{P}$. Equality for every policy in $\mathfrak{P}_2$ implies equality for every policy in its subset $\mathfrak{P}_1$, which proves the first inclusion. Equality of the $(n+1)$-stage laws implies equality of their n-stage marginals, which proves the second. □

Proof of Theorem 3.4 Assume first that g is constant on controlled fibres. For $z \in \mathcal{E}_{\mathfrak{P},n}(\Omega)$ choose a representative ω and put $\bar{g}(z) = g(\omega)$. If ω' is another representative, then $\omega \sim_{\mathfrak{P},n} \omega'$, so the definition is independent of the choice. It gives $g = \bar{g} \circ \mathcal{E}_{\mathfrak{P},n}$. Surjectivity onto the displayed image makes $\bar{g}$ unique.

Conversely, a factorisation through $\mathcal{E}_{\mathfrak{P},n}$ gives equal task values whenever the experiment-family images agree. This is precisely Definition 3.3.

For the measurable assertion, let $d \in \mathcal{D}$. The equality

$$\mathcal{E}^{-1}_{\mathfrak{P},n}(\bar{g}^{-1}\{d\}) = g^{-1}\{d\}$$

has a Borel right-hand side. The quotient property makes $\bar{g}^{-1}\{d\}$ Borel in the image. Since $\mathcal{D}$ is finite and discrete, $\bar{g}$ is measurable. □

A.1 Fitting rank locus

At a point θ, a matrix has rank at most r exactly when all its $(r+1)\times(r+1)$ minors vanish. Applied to $J_H(\theta)$, this proves the zero-germ statement in Definition 3.6. The formula for $\mathrm{Fitt}_j(\mathcal{M}_H)$ is the defining determinantal formula for the displayed presentation $A^d \xrightarrow{J_H} A^q \to \mathcal{M}_H \to 0$.

Proof of Theorem 3.7 Put $A_0 = DH(\theta_0)$, $B_0 = DL_u(\theta_0)$, and $K = \ker A_0$. The kernel of the stacked map (A_0, B_0) is $K \cap \ker B_0$. Rank–nullity gives

$$\begin{aligned}\operatorname{rank}(A_0, B_0) - \operatorname{rank} A_0 &= \dim K - \dim(K \cap \ker B_0) \\ &= \operatorname{rank}(B_0|_K),\end{aligned}$$

which is (6).

If a C^1 map $F : \mathbb{R}^d \to \mathbb{R}^s$ has constant rank k near a point, the constant-rank theorem supplies local source and target coordinates in which

$$F(x_1, \dots, x_d) = (x_1, \dots, x_k, 0, \dots, 0).$$

Its fibre through the point is therefore an embedded submanifold of dimension $d - k$. Apply this normal form to H and then to (H, L_u). Their fibre dimensions are $d-r$ and $d-r'$, and their difference is $r'-r = \rho$. □

Proof of Theorem 3.5 At stage one the states use the same policy kernel and, by (3), the same observation kernel. Their joint laws of (U_1, Y_1) therefore agree. Suppose the history laws agree through stage $i-1$. Conditional on any such history, the common policy selects U_i by the same kernel, after which (3) selects Y_i by the same kernel. The recursion in Proposition 2.1 gives equality through stage i. Induction reaches n. Opposite task values in one fibre contradict Definition 3.3. □

Proof of Corollary 3.8 Apply Proposition 3.2 with $\mathfrak{P}_1 = \{\pi^0\}$ and $\mathfrak{P}_2 = \mathfrak{P}$. A single recorded law contains no observations from policies in $\mathfrak{P} \setminus \{\pi^0\}$, so the final evidential statement follows. □

B Proofs for stopped information

Proof of Theorem 4.1 Use the jointly measurable conditional Radon–Nikodym derivatives stipulated in Theorem 4.1:

$$R_i(y \mid h, u) = \frac{\mathrm{d}Q_{\omega_0,i}(\cdot \mid h, u)}{\mathrm{d}Q_{\omega_1,i}(\cdot \mid h, u)}(y).$$

On the event $\{\tau = t\}$, the recursive construction in Proposition 2.1 factors the transcript law into alternating policy and observation kernels. The policy factors are identical under the two states. The event $\{\tau = t\}$ is the same measurable subset of the history space under both laws. Conditional absolute continuity therefore propagates through each finite kernel product. On the tth component of the disjoint union in (7), it gives

$$\frac{\mathrm{d}P_0^{\pi,\tau}}{\mathrm{d}P_1^{\pi,\tau}} = \prod_{i=1}^{\tau} R_i(Y_i \mid H_{i-1}, U_i) \quad P_0^{\pi,\tau}\text{-almost surely.} \tag{30}$$

There are only $N+1$ components, so no limit is taken. For a realised conditioning pair, put $R = R_i(\cdot \mid h, u)$ and $Q = Q_{\omega_1,i}(\cdot \mid h, u)$. Then

$$\int_{\{R<1\}} -R \log R \,\mathrm{d}Q \le e^{-1}.$$

Thus the negative part of each stopped log factor is integrable, while (9) controls its positive part. Put $\ell_i = \log R_i(Y_i \mid H_{i-1}, U_i)$. Taking the logarithm and expectation under state ω_0 gives

$$\mathrm{KL}(P_0^{\pi,\tau} \| P_1^{\pi,\tau}) = \sum_{i=1}^{N} \mathbb{E}_{\omega_0}^{\pi}[\mathbf{1}_{\{\tau \ge i\}} \ell_i].$$

The event $\{\tau \ge i\}$ is measurable with respect to H_{i-1}, and U_i is selected before Y_i. Conditional expectation given (H_{i-1}, U_i) therefore replaces the logarithm by $k_i(H_{i-1}, U_i)$. The tower property and finite summation now yield (10). □

Proof of Corollary 4.2 For $X \sim N(m_0, \Sigma_0)$, put $D = m_1 - m_0$. Direct subtraction of the two Gaussian log-densities and evaluation under X gives

$$\begin{aligned}\mathbb{E}_0 \log \frac{p_0(X)}{p_1(X)} = \frac{1}{2}\{ &\log(\det \Sigma_1 / \det \Sigma_0) - d \\ &+ \operatorname{tr}(\Sigma_1^{-1}\Sigma_0) + D^{\mathsf{T}}\Sigma_1^{-1}D\}.\end{aligned}$$

This proves (11); the common-covariance reduction is immediate. The formula is applied conditionally at each (i, h, u). □

Lemma B.1 (Binary processing of relative entropy) *If $P \ll Q$ and A is measurable, then*

$$\mathrm{KL}(P\|Q) \ge \mathrm{kl}(P(A), Q(A)), \tag{31}$$

with the usual lower-semicontinuous extension at zero and one. If $P(A) \ge 1-\alpha$ and $Q(A) \le \alpha < 1/2$, the right-hand side is at least $\mathrm{kl}(1-\alpha, \alpha)$.

Proof Put $p = P(A)$ and $q = Q(A)$. When $0 < q < 1$, conditional Jensen inequalities on A and A^c give

$$\int_A \frac{\mathrm{d}P}{\mathrm{d}Q} \log \frac{\mathrm{d}P}{\mathrm{d}Q}\,\mathrm{d}Q \geq p \log(p/q),$$
$$\int_{A^c} \frac{\mathrm{d}P}{\mathrm{d}Q} \log \frac{\mathrm{d}P}{\mathrm{d}Q}\,\mathrm{d}Q \geq (1-p) \log\{(1-p)/(1-q)\}.$$

Their sum proves (31). The boundary cases follow by lower semicontinuity and absolute continuity. On the region $p \geq 1-\alpha > \alpha \geq q$, differentiation shows that $\mathrm{kl}(p, q)$ increases with p and decreases with q; its minimum is therefore attained at $(1-\alpha, \alpha)$. □

Proof of Theorem 4.4 Let $A = \{\widehat{d}_\delta = g(\omega_0)\}$. The error assumptions give

$$P_{\omega_0}^{\pi_\delta, \tau_\delta}(A) \geq 1-\alpha, \qquad P_{\omega_\delta}^{\pi_\delta, \tau_\delta}(A) \leq \alpha.$$

Define

$$d_{\delta,i}(h,u) = \mathrm{KL}\big(Q_{\omega_0,i}(\cdot \mid h,u) \big\| Q_{\omega_\delta,i}(\cdot \mid h,u)\big)\,.$$

Lemma B.1 and Theorem 4.1 then give

$$\begin{aligned}\mathrm{kl}(1-\alpha, \alpha) &\leq \mathbb{E}_{\omega_0}^{\pi_\delta} \sum_{i=1}^{\tau_\delta} d_{\delta,i}(H_{i-1}, U_i) \\ &\leq C\delta^{2m} \mathbb{E}_{\omega_0}^{\pi_\delta} \tau_\delta.\end{aligned}$$

Rearrangement proves (13). □

In the first part of Example 4.5, let Φ_{N} denote the standard normal distribution function. The midpoint test based on the mean of n observations has both errors

$$\Phi_{\mathrm{N}}\left(-\frac{\sqrt{n}\,\delta^m}{2\sigma}\right).$$

Thus a fixed target error is attained with n of order δ^{-2m}. For the variance example in Section 4, write K_δ for the divergence from $N(0,1)$ to $N(0, 1+\delta)$. Then

$$\begin{aligned}K_\delta &= \frac{1}{2}\{(1+\delta)^{-1} - 1 + \log(1+\delta)\} \\ &= \frac{1}{4}\delta^2 + O(\delta^3).\end{aligned}$$

C Proofs for robust prefix stopping

Proof of Lemma 5.1 Fix a state ω and put $\theta = \vartheta(\omega)$. Conditional on $(H_{n-1}, U_n) = (h, u)$, membership (14) and the envelope (15) give

$$\mathbb{E}_\omega^{\star,\pi}\left[\frac{f_{\theta,n}(Y_n \mid h,u)}{\kappa_{\theta,n}(h,u)} \middle| H_{n-1} = h, U_n = u\right] \leq 1.$$

The current control is drawn from the common policy after H_{n-1} and before Y_n. Conditioning once more on H_{n-1} therefore yields

$$\mathbb{E}_\omega^{\star,\pi}[E_n(\theta) \mid H_{n-1}] \leq E_{n-1}(\theta).$$

Nonnegativity and $E_0 = 1$ prove the supermartingale assertion.

For $a > 1$, let $\sigma = \inf\{n \leq N : E_n(\theta) \geq a\}$, with $\sigma = N$ if no crossing occurs. Bounded optional sampling for the nonnegative supermartingale gives $\mathbb{E}E_\sigma \leq 1$. On the crossing event, $E_\sigma \geq a$; hence

$$P\{\max_{n \leq N} E_n(\theta) \geq a\} \leq a^{-1}.$$

Taking $a = \alpha^{-1}$ proves (17). □

Proof of Theorem 5.2 Let

$$A = \{E_n(\vartheta(\omega_0)) < \alpha^{-1} \text{ for every } 0 \leq n \leq N\}$$

and let B be the simultaneous outer-containment event in (19). Lemma 5.1 gives $P_{\omega_0}^{\star,\pi}(A) \geq 1-\alpha$, while (19) gives $P_{\omega_0}^{\star,\pi}(B) \geq 1-\beta$.

On A, the true state ω_0 belongs to C_n for every n. The definition of saturation then gives

$$\mathfrak{F}\mathfrak{P},_n(\omega_0) \subseteq S_n \quad \text{for every } n.$$

On B, $S_n \subseteq \widehat{F}_n$ simultaneously. Thus the inclusions in (21) hold on $A \cap B$, whose probability is at least $1-\alpha-\beta$.

In particular, on $A \cap B$ one has $\omega_0 \in \widehat{F}_\tau$. If the rule in (20) reports $d \neq \perp$, every state in $\widehat{F}_\tau$ has task value d; hence $d = g(\omega_0)$. A wrong, non-rejected report is therefore contained in $(A \cap B)^c$, proving (22). □

Proof of Corollary 5.3 For any probability measures P, Q and any measurable $0 \leq f \leq M$, the layer-cake representation gives

$$\begin{aligned}\int f\,\mathrm{d}P - \int f\,\mathrm{d}Q &= \int_0^M \{P(f > t) - Q(f > t)\}\,\mathrm{d}t \\ &\leq M\,\mathrm{TV}(P, Q).\end{aligned}$$

With $Q = Q_{\theta,n}^0$ this proves $\int f_{\theta,n}\,\mathrm{d}P \leq a_{\theta,n} + M_{\theta,n}\varepsilon_{\theta,n}$, which is (15) with (23).

For the KL implication, Lemma B.1 gives, for every event A,

$$\mathrm{KL}(P\|Q) \geq \mathrm{kl}(P(A), Q(A)).$$

The scalar inequality $\mathrm{kl}(p, q) \geq 2(p-q)^2$ follows because, for fixed q, the second derivative in p of $\mathrm{kl}(p,q) - 2(p-q)^2$ is $1/\{p(1-p)\} - 4 \geq 0$, and the function and its first derivative vanish at $p = q$. Taking the supremum over A yields $\mathrm{TV}(P, Q) \leq \sqrt{\mathrm{KL}(P\|Q)/2}$. □

For the contamination calculation following Corollary 5.3,

$$\int f \,\mathrm{d}\{(1-\varepsilon)Q^0 + \varepsilon R\} = (1-\varepsilon)a + \varepsilon \int f \,\mathrm{d}R \\ \leq (1-\varepsilon)a + \varepsilon M.$$

Proof of Proposition 5.4 On $\{0,1\}$ and for $0 < \delta < 1$, set

$$q_\delta(1) = \delta^3, \qquad p_\delta(1) = \delta^2, \qquad r_\delta(1) = 1.$$

All complementary probabilities are determined by normalisation. Since q_δ has full support, $r_\delta \ll q_\delta$. Direct calculation gives

$$\mathrm{TV}(p_\delta, q_\delta) = \delta^2 - \delta^3 \longrightarrow 0,$$

whereas

$$\mathbb{E}_{q_\delta} \frac{r_\delta}{q_\delta} = \delta^3 \delta^{-3} = 1,$$
$$\mathbb{E}_{p_\delta} \frac{r_\delta}{q_\delta} = \delta^2 \delta^{-3} = \delta^{-1} \longrightarrow \infty.$$

Thus no calibration factor tending to one can be justified from the displayed total-variation convergence for this unbounded likelihood ratio. □

D Proofs and assumptions for causal-prefix replay

Proof of Proposition 6.1 Absorb any exogenous randomisation used by the stopping rule into W. In the full regime, let $K_t(\mathrm{d}z_t \mid z_{1:t-1}, W)$ be the regular conditional kernel of frame t. By assumption, the same kernel is used in the stop-only regime on the event that acquisition is active at t. Induction therefore gives the same law for $Z_{1:t}$ under the two regimes on every active prefix:

$$\mu(\mathrm{d}z_{1:t} \mid W) = K_1(\mathrm{d}z_1 \mid W) \prod_{s=2}^{t} K_s(\mathrm{d}z_s \mid z_{1:s-1}, W). \tag{32}$$

Let $A_{j,d}$ be the set of prefixes for which no earlier gate has a singleton, the rule stops at t_j, and it reports d. Adaptedness makes $A_{j,d}$ measurable in $\sigma(W, Z_{1:t_j})$. For every measurable set B of t_j-prefixes, both regimes consequently assign

$$\int \mathbf{1}\{z_{1:t_j} \in B \cap A_{j,d}\} \mu(\mathrm{d}z_{1:t_j} \mid W) \tag{33}$$

to the event $\{Z_{1:t_j} \in B, \tau = t_j, \widehat{D} = d\}$, conditionally on W. The analogous expression with the event that no gate is a singleton covers the terminal rejection at T. Summing over the finitely many gates and then integrating over W proves equality of the laws on the disjoint union of the prefix spaces.

Assumption (ii) is essential: without it the full acquisition could alter an earlier reconstructed frame, so the kernels in (32) would differ. The argument also uses no kernel after stopping and therefore says nothing about an action that is absent from the recorded protocol. □

Proof of Proposition 6.2 Condition on $\mathcal{T}$, which includes the training data and external randomisation. At gate j, apply its fixed score map to the m calibration subjects and to the true label of the future subject. Subject-level exchangeability makes the $m+1$ scores exchangeable. The conservative rank p-value

$$p_{j,D_*} = \frac{1 + \sum_{i=1}^{m} \mathbf{1}\{R_{ij} \geq R_{*j}(D_*)\}}{m+1}$$

is therefore super-uniform: for every $u \in [0,1]$, $\mathbb{P}\{p_{j,D_*} \leq u \mid \mathcal{T}\} \leq u$. One way to see this is to break equal-score ranks uniformly at random. The resulting rank is uniform on $\{1, \ldots, m+1\}$; replacing the random tie break by counting every tie against rejection can only increase the p-value. Thus

$$\mathbb{P}\{D_* \notin \Gamma_j \mid \mathcal{T}\} = \mathbb{P}\{p_{j,D_*} \leq \alpha/G \mid \mathcal{T}\} \leq \alpha/G.$$

The union bound, which requires no independence among the gates, gives (25) conditionally on $\mathcal{T}$ and hence unconditionally.

If the rule makes a wrong non-rejected report at gate j, then Γ_j is a singleton containing a label different from D_*. Hence

$$\{\widehat{D} \neq D_*, \widehat{D} \neq \perp\} \subseteq \bigcup_{j=1}^{G} \{D_* \notin \Gamma_j\}.$$

The preceding union bound proves (26). □

Proposition 6.1 is a kernel identity for the recorded protocol; Proposition 6.2 is marginal across exchangeable subjects. Neither implies the other or identifies an unrecorded acquisition action.

E DCE-MRI forward model and numerical definitions

E.1 Analytic image and coil arrays

For $N = 32$, put

$$x_i = \frac{i - (N-1)/2}{N/2}, \qquad 0 \leq i < N,$$
$$y_j = \frac{j - (N-1)/2}{N/2}, \qquad 0 \leq j < N.$$

For compactness, write

$$I_{a,b;c,d}(x,y) = \mathbf{1}\left\{\left(\frac{x-c}{a}\right)^2 + \left(\frac{y-d}{b}\right)^2 \le 1\right\}. \tag{34}$$

The real baseline magnitude is

$$\begin{aligned} B_0 = 0.04 &+ 0.42 I_{0.84,0.72;0,0} \\ &+ 0.22 I_{0.48,0.50;0,-0.03} \\ &+ 0.10 I_{0.18,0.12;-0.27,0.20}, \end{aligned} \tag{35}$$

with phase

$$\varphi(x,y) = 0.24x - 0.18y + 0.10\sin(\pi x)\cos(\pi y).$$

The complex baseline is $B_c = B_0 e^{\mathrm{i}\varphi}$; each real lesion basis is likewise multiplied by $e^{\mathrm{i}\varphi}$ to define Φ_h.

The two nuisance locations are $(-0.25, -0.18)$ and $(0.25, 0.18)$. At centre (x_0, y_0) set

$$q(x,y) = \left(\frac{x-x_0}{0.19}\right)^2 + \left(\frac{y-y_0}{0.14}\right)^2.$$

The solid basis is $\mathbf{1}\{q \le 1\}$. The rim basis is $\mathbf{1}\{q \le 1\} - \mathbf{1}\{q \le 0.42\}$, rescaled to have the same sum as the solid basis. The equal DC component deliberately makes central k-space comparatively uninformative about the task without making all mask-specific laws coincide.

Hypotheses $h = 0, 1$ use the solid basis at locations $0, 1$; hypotheses $h = 2, 3$ use the rim basis at the same respective locations. Their task labels are $\ell_h = \lfloor h/2 \rfloor$.

For coil $c \in \{0, 1, 2, 3\}$, let $\theta_c = 2\pi c/4$ and

$$\begin{aligned} \rho_c^2(x,y) &= (x - 1.25\cos\theta_c)^2 + (y - 1.25\sin\theta_c)^2, \\ \psi_c(x,y) &= x\cos\theta_c + y\sin\theta_c, \\ \widetilde{s}_c(x,y) &= \exp\{-\rho_c^2(x,y)/1.25 + 0.45\mathrm{i}\psi_c(x,y)\}. \end{aligned} \tag{36}$$

The code normalises these pointwise so that $\sum_c |s_c(x,y)|^2 = 1$. The forward and adjoint transforms are centred, orthonormal two-dimensional FFTs; their numerical inner-product discrepancy is 2.79e-15 [28, 29].

Each mask contains the two central phase-encode rows and two symmetric row pairs. In zero-based centred offsets, the seven pair choices are $(2,3), (4,5), \ldots, (14,15)$; mask a contains offsets $\{-1, 0, -k_1, k_1, -k_2, k_2\}$. All readout points and all four coils are retained, giving $4 \times 6 \times 32 = 768$ complex observations per frame. Let S_a denote the linear operator that retains these six rows in every coil.

E.2 Kinetics and SPGR response

Following the standard DCE notation in [13], the phenotype-zero parameters are $K^{\mathrm{trans}} = 0.100\,\mathrm{min}^{-1}$ and $v_e = 0.300$; phenotype one uses $0.130\,\mathrm{min}^{-1}$ and 0.280. With $k_{ep} = K^{\mathrm{trans}}/v_e$ and $\Delta t = 1/12$ min, holding C_p constant within one interval gives the recursion in (27). For minutes t after injection, the Parker population input [14] is

$$\begin{aligned} C_p(t) = &\frac{A_1}{s_1\sqrt{2\pi}} e^{-(t-T_1)^2/(2s_1^2)} \\ &+ \frac{A_2}{s_2\sqrt{2\pi}} e^{-(t-T_2)^2/(2s_2^2)} \\ &+ \frac{\alpha_P e^{-\beta_P t}}{1 + e^{-\gamma(t-\tau_P)}}, \end{aligned}$$

for $t \ge 0$, and $C_p(t) = 0$ for $t < 0$, where

$$\begin{aligned} (A_1, A_2) &= (0.809, 0.330), \\ (T_1, T_2) &= (0.17046, 0.365), \\ (s_1, s_2) &= (0.0563, 0.132), \\ (\alpha_P, \beta_P) &= (1.050, 0.1685), \\ (\gamma, \tau_P) &= (38.078, 0.483). \end{aligned}$$

Injection begins after four baseline frames. The simulation uses $\sigma_0 = 0.030$ mM and $\sigma_w = 0.0015$ mM, with independent innovations at every $t > 0$. The initial state, transition innovations, and receiver noises are mutually independent; innovations and receiver noises are independent across time, and receiver noises are independent across coils.

For concentration c, define

$$\begin{aligned} E(c) &= \exp[-\mathrm{TR}(T_{10}^{-1} + r_1 c)], \\ S(c) &= e^{-\mathrm{TE}R_2^*} \sin\alpha \, \frac{1 - E(c)}{1 - \cos\alpha E(c)}. \end{aligned} \tag{37}$$

This is the spoiled gradient-echo response used for the tangent and stress layers [30, 31]. The

constants are

$$\begin{aligned} \mathrm{TR} &= 5\ \mathrm{ms}, & \mathrm{TE} &= 2\ \mathrm{ms}, \\ \alpha &= 15^\circ, & T_{10} &= 1.20\ \mathrm{s}, \\ r_1 &= 4.5\ \mathrm{mM}^{-1}\mathrm{s}^{-1}, & R_2^* &= 18\ \mathrm{s}^{-1}. \end{aligned}$$

Writing $E_0 = e^{-\mathrm{TR}/T_{10}}$, the derivative used in the exact layer is

$$\beta = S'(0) = r_1 e^{-\mathrm{TE}R_2^*} \sin\alpha\, \mathrm{TR} E_0 \frac{1-\cos\alpha}{(1-\cos\alpha E_0)^2}. \tag{38}$$

The analytic derivative agrees with a centred finite difference to 4.25×10^{-11} in the distributed test.

Let Φ_h be the complex lesion basis and let s_c be the coil map. The generated full k-space is

$$\begin{aligned} K_{h,t,c} &= \mathcal{F}[s_c\{B_c + A(c_{h,t})\Phi_h\}] + \eta_{t,c}, \\ \eta_{t,c} &\sim \mathcal{CN}(0, \sigma^2 I). \end{aligned} \tag{39}$$

where $\sigma = 0.060$ and the real and imaginary components each have variance $\sigma^2/2$. The policy observes only the selected rows. In the exact layer, $A(c) = \beta c$; in the nonlinear stress layer, $A(c) = S(c) - S(0)$. The Gaussian concentration state is untruncated and may be negative. This choice makes the exact likelihood auditable but prevents the construction from being interpreted as a physiological simulator.

E.3 Exact predictive density and confidence set

Condition on a hypothesis h, a selected mask a, and the preceding history. After baseline subtraction, write

$$\begin{aligned} Z &= v_{a,h} C + \eta, \\ C &\sim N(m, V), \\ \eta &\sim \mathcal{CN}(0, \sigma^2 I_d). \end{aligned}$$

where $v_{a,h} = \beta S_a \mathcal{F}(s_c \Phi_h)_{c=0}^3$ and $d = 768$. For $v \in \mathbb{C}^d$, define $\widetilde{v} = (\mathrm{Re}\, v^\mathsf{T}, \mathrm{Im}\, v^\mathsf{T})^\mathsf{T}$. Because C is real, the real-stacked covariance is $\sigma^2 I_{2d}/2 + V\widetilde{v}\widetilde{v}^\mathsf{T}$, not a proper-complex rank-one covariance. If $e = z - vm$, $H = \|v\|^2$, and $\rho = \mathrm{Re}(v^* e)$, direct application of the determinant and Sherman–Morrison formulae gives

$$\begin{aligned} \log q_h(z) = & -d\log\pi - d\log\sigma^2 \\ & - \tfrac{1}{2}\log(1 + 2VH/\sigma^2) \\ & - \|e\|^2/\sigma^2 \\ & + \frac{2V\rho^2}{\sigma^2(\sigma^2 + 2VH)}. \end{aligned} \tag{40}$$

The scalar update is

$$\begin{aligned} V^+ &= (V^{-1} + 2H/\sigma^2)^{-1}, \\ m^+ &= V^+\{m/V + 2\,\mathrm{Re}(v^* z)/\sigma^2\}. \end{aligned} \tag{41}$$

The state transition is then applied before the next acquisition. An independent real-multivariate-normal calculation verifies (40) to machine precision.

Numbering acquired frames from one in the likelihood recursion, set $p_0(h) = 1/4$. Let $p_{t-1}(h)$ be the posterior mass before frame t and set $r_t = \sum_h p_{t-1}(h) q_{h,t}$. After observing Y_t, put

$$p_t(h) = \frac{p_{t-1}(h) q_{h,t}(Y_t)}{r_t(Y_t)}.$$

The selected mask is measurable with respect to the preceding history, so under hypothesis h

$$\mathbb{E}_h\left[\frac{r_t(Y_t)}{q_{h,t}(Y_t)}\middle|\mathcal{F}_{t-1}\right] = 1.$$

Thus (29) is a test martingale in the exact layer even under adaptive mask choice.

For the deterministic one-stage policy selecting mask a, the initial conditions $m = 0$ and $V = \sigma_0^2 > 0$ give covariance trace

$$d\sigma^2 + \sigma_0^2 \|v_{a,h}\|^2.$$

Equality of controlled path laws implies equality of this one-stage marginal for every a, hence equality of the signature $(\|v_{a,h}\|^2)_a$, but not conversely. The code groups exact rational sums of squares of the stored IEEE-754 components, without a tolerance. These classes may merge distinct laws but cannot split an equal pair. They are singletons for the four hypotheses of the main experiment; the no-go pair duplicates every transition and observation vector under opposite labels. A supplied counterexample has equal initial signatures but distinct transitions: its second-frame

covariance traces differ by 3.529e-06. The signature is therefore a certified outer relation, not a controlled-fibre algorithm.

E.4 Acquisition scores and reported metrics

After $t \in \{0, \dots, 35\}$ frames and before frame $t+1$, let p_h, m_h, V_h be the current posterior mass, mean, and variance. Put $G_{a,hg} = \mathrm{Re}(v_{a,h}^* v_{a,g})$ and $\mathcal{P}_\times = \{(h,g) : 0 \le h < g \le 3,\ \ell_h \ne \ell_g\}$. For $(h,g) \in \mathcal{P}_\times$, define

$$d_a(h,g) = \frac{m_h^2 G_{a,hh} + m_g^2 G_{a,gg} - 2m_h m_g G_{a,hg}}{d\sigma^2 + V_h G_{a,hh} + V_g G_{a,gg}}. \tag{42}$$

Let $w_{hg} = p_h p_g$, $W = \sum_{\mathcal{P}_\times} w_{hg}$, and $\mathcal{V} = \{(h,g) \in \mathcal{P}_\times : p_h, p_g \ge 0.025\}$. Define

$$\bar{d}_a = \frac{\sum_{\mathcal{P}_\times} w_{hg} d_a(h,g)}{\max\{W, 10^{-15}\}},$$
$$I_a = \sum_h p_h \log(1 + 2V_h G_{a,hh}/\sigma^2),$$

and let $d_a^{\min} = \min_{(h,g)\in\mathcal{V}} d_a(h,g)$ when $\mathcal{V} \ne \varnothing$, and zero otherwise. The task-aware policy maximises $\bar{d}_a + 2d_a^{\min} + 0.005 I_a - 0.0002\Delta_t(a)$, where

$$\Delta_t(a) = \begin{cases} 0, & t = 0, \\ |a - a_{t-1}|/7, & t > 0. \end{cases}$$

It is a causal heuristic, not a Bellman solution. The variance-only policy minimises

$$\frac{1}{4}\sum_h (V_h^{-1} + 2G_{a,hh}/\sigma^2)^{-1}$$

with fixed uniform weights. Ties are resolved by the smallest mask index. The fixed policy uses $a_t = t \bmod 7$, beginning with mask zero. The mask-separation panel in Figure 5 evaluates these terms at the peak of the deterministic kinetic trajectory, with $V_h = \sigma_w^2$ and $p_h = 1/4$ for all four hypotheses.

Let $D \in \{0, 1, \perp\}$ be the terminal report. Decision coverage is $\mathbf{1}\{D \ne \perp\}$; selective risk is the mean of $\mathbf{1}\{D \ne \ell_h\}$ over non-rejected runs. Anytime coverage records whether the true finite hypothesis belongs to every confidence set up to stopping, and log loss uses the final phenotype posterior. In scenario s, let $j \in \{0, \dots, \tau - 1\}$ index the acquired frames from zero and let $p_{g,j+1}, m_{g,j+1}, V_{g,j+1}$ be the filtered quantities after frame j. Define

$$X_j^{\mathrm{dyn}} = A_s(c_{h,j})\Phi_h,$$
$$\widehat{X}_j^{\mathrm{dyn}} = \beta \sum_g p_{g,j+1} m_{g,j+1} \Phi_g,$$
$$\widehat{c}_j = \sum_g p_{g,j+1} m_{g,j+1},$$
$$\widehat{M}_{2,j} = \sum_g p_{g,j+1}(V_{g,j+1} + m_{g,j+1}^2),$$
$$s_j^2 = \max\{\widehat{M}_{2,j} - \widehat{c}_j^2, 0\},$$

where $A_s(c) = \beta c$ in the linear scenario and $A_s(c) = S(c) - S(0)$ in the nonlinear scenario. Thus complex dynamic NRMSE is

$$\left\{ \frac{\sum_{j<\tau} \|\widehat{X}_j^{\mathrm{dyn}} - X_j^{\mathrm{dyn}}\|_2^2}{\max\{\sum_{j<\tau} \|X_j^{\mathrm{dyn}}\|_2^2, 10^{-15}\}} \right\}^{1/2}.$$

The displayed state interval is $[\widehat{c}_j - 1.96 s_j, \widehat{c}_j + 1.96 s_j]$; it is descriptive. The estimator $\widehat{X}_j^{\mathrm{dyn}}$ retains the linear response βc in both scenarios.

The scalar robustness control uses declared denominator $N(0,1)$, numerator $N(1,1)$, and actual null $N(0.6, 1)$. Its nominal factor is $\exp(Y - 1/2)$ and has actual conditional expectation $\kappa = \exp(0.6)$. Division by this exact factor gives a mean-one martingale. This control is intentionally separate from the MRI forward model.

The MRI robustness control is defined directly at the predictive-kernel level. Let $Q_{h,t}(\cdot \mid H_{t-1}, a_t)$ denote (40) with the recursively updated mean and variance. For each h, let $\bar{h}$ be the unique hypothesis with the same nuisance location and opposite phenotype, and put

$$P_{h,t}^\star = (1 - \varepsilon) Q_{h,t} + \varepsilon Q_{\bar{h},t}. \tag{43}$$

This is a prescribed conditional law, not a hidden-state approximation. With $f_{h,t} = \min\{r_t/q_{h,t}, M\}$,

$$\mathbb{E}_{P_{h,t}^\star} f_{h,t} = (1-\varepsilon)\mathbb{E}_{Q_{h,t}} f_{h,t} + \varepsilon \mathbb{E}_{Q_{\bar{h},t}} f_{h,t}$$
$$\le (1 - \varepsilon) + \varepsilon M =: \kappa.$$

The first term is at most one because clipping cannot increase $\mathbb{E}_{Q_{h,t}}(r_t/q_{h,t}) = 1$; the second is at most M pointwise. Consequently $\prod_{i\le t} f_{h,i}/\kappa$ is the robust test supermartingale of Lemma 5.1. The implementation uses $\varepsilon = 0.10$ and $M = 2.0$, samples (43) conditionally at every prefix, and reuses the same contamination uniforms and Gaussian arrays across policies. It continues each path to the full horizon to audit any crossing, while separately recording the first label-singleton stopping prefix. This construction proves the envelope for the stipulated class only; it does not estimate ε from empirical MRI data.

Binary frequencies in the distributed summaries carry two-sided Wilson score intervals. Continuous means and paired continuous policy contrasts use normal Monte Carlo intervals, while paired binary differences use a bounded Hoeffding interval. No familywise coverage claim is made. The four hypotheses are assigned cyclically, giving 250 observations per hypothesis in each main policy–scenario cell.

All distributed numerical outputs are reproduced from the package root, in the following order, by

```
python3 code/run_retrospective_prefix.py
  --self-test
  --output data/retrospective_prefix_
selftest.json
  python3 code/run_prefix_phantom.py
  --runs 1000 --seed 20260814
  --package-root .
  python3 code/run_prefix_phantom.py
  --test-only
```

The main simulation generated 6000 policy–scenario rows in 39.95 seconds in the recorded environment. Runtime is hardware dependent. The manifest records software versions, file hashes, the seed rule, and the exact/stress/clinical scope boundaries.

References

[1] H. Chernoff, Sequential design of experiments, *Ann. Math. Statist.* **30**(3), 755–770 (1959). https://doi.org/10.1214/aoms/1177706205

[2] M. Naghshvar and T. Javidi, Active sequential hypothesis testing, *Ann. Statist.* **41**(6), 2703–2738 (2013). https://doi.org/10.1214/13-AOS1144

[3] S. Nitinawarat, G.K. Atia, and V.V. Veeravalli, Controlled sensing for multihypothesis testing, *IEEE Trans. Automat. Control* **58**(10), 2451–2464 (2013). https://doi.org/10.1109/TAC.2013.2261188

[4] L. Pineda, S. Basu, A. Romero, R. Calandra, and M. Drozdzal, Active MR k-space sampling with reinforcement learning, in *Medical Image Computing and Computer Assisted Intervention—MICCAI 2020*, Lecture Notes in Computer Science, vol. 12262, pp. 23–33 (Springer, Cham, 2020). https://doi.org/10.1007/978-3-030-59713-9_3

[5] T. Bakker, M. Muckley, A. Romero-Soriano, M. Drozdzal, and L. Pineda, On learning adaptive acquisition policies for undersampled multi-coil MRI reconstruction, in *Proceedings of the 5th International Conference on Medical Imaging with Deep Learning*, Proc. Mach. Learn. Res. **172**, 63–85 (2022). https://proceedings.mlr.press/v172/bakker22a.html

[6] C.-Y. Yen, R. Singhal, U. Sharma, R. Ranganath, S. Chopra, and L. Pinto, Adaptive sampling of k-space in magnetic resonance for rapid pathology prediction, in *Proceedings of the 41st International Conference on Machine Learning*, Proc. Mach. Learn. Res. **235**, 57018–57032 (2024). https://proceedings.mlr.press/v235/yen24a.html

[7] Y. Du, R. Dharmakumar, and S.A. Tsaftaris, The MRI scanner as a diagnostic: image-less active sampling, in *Medical Image Computing and Computer Assisted Intervention—MICCAI 2024*, Lecture Notes in Computer Science, vol. 15003, pp. 467–476 (Springer, Cham, 2024). https://doi.org/10.1007/978-3-031-72384-1_44

[8] Y. Du, J. Liu, R. Dharmakumar, and S.A. Tsaftaris, Active sampling for MRI-based sequential decision making, arXiv:2505.04586v2 [eess.IV], preprint under review (2025; revised 13 February 2026). https://arxiv.org/abs/2505.04586v2

[9] S.R. Howard, A. Ramdas, J. McAuliffe, and J. Sekhon, Time-uniform, nonparametric, nonasymptotic confidence sequences, *Ann. Statist.* **49**(2), 1055–1080 (2021). https://doi.org/10.1214/20-AOS1991

[10] P. Grünwald, R. de Heide, and W. Koolen, Safe testing, *J. R. Stat. Soc. Ser. B Stat. Methodol.* **86**(5), 1091–1128 (2024). https://doi.org/10.1093/jrsssb/qkae011

[11] R. Hermann and A.J. Krener, Nonlinear controllability and observability, *IEEE Trans. Automat. Control* **22**(5), 728–740 (1977). https://doi.org/10.1109/TAC.1977.1101601

[12] R. van Handel, Observability and nonlinear filtering, *Probab. Theory Related Fields* **145**(1–2), 35–74 (2009). https://doi.org/10.1007/s00440-008-0161-y

[13] P.S. Tofts, G. Brix, D.L. Buckley, J.L. Evelhoch, E. Henderson, M.V. Knopp, H.B.W. Larsson, T.-Y. Lee, N.A. Mayr, G.J.M. Parker, R.E. Port, J. Taylor, and R.M. Weisskoff, Estimating kinetic parameters from dynamic contrast-enhanced T_1-weighted MRI of a diffusable tracer: standardized quantities and symbols, *J. Magn. Reson. Imaging* **10**(3), 223–232 (1999). https://doi.org/10.1002/(SICI)1522-2586(199909)10:3⟨223::AID-JMRI2⟩3.0.CO;2-S

[14] G.J.M. Parker, C. Roberts, A. Macdonald, G.A. Buonaccorsi, S. Cheung, D.L. Buckley, A. Jackson, Y. Watson, K. Davies, and G.C. Jayson, Experimentally-derived functional form for a population-averaged high-temporal-resolution arterial input function for dynamic contrast-enhanced MRI, *Magn. Reson. Med.* **56**(5), 993–1000 (2006). https://doi.org/10.1002/mrm.21066

[15] Y. Romano, M. Sesia, and E.J. Candès, Classification with valid and adaptive coverage, in *Advances in Neural Information Processing Systems 33*, pp. 3581–3591 (2020). NeurIPS proceedings record

[16] C.T. Ionescu Tulcea, Mesures dans les espaces produits, *Atti Accad. Naz. Lincei Rend. Cl. Sci. Fis. Mat. Nat.*, Ser. VIII, **7**, 208–211 (1949).

[17] D. Blackwell, Equivalent comparisons of experiments, *Ann. Math. Statist.* **24**(2), 265–272 (1953). https://doi.org/10.1214/aoms/1177729032

[18] D. Eisenbud, *Commutative Algebra: With a View Toward Algebraic Geometry*, Graduate Texts in Mathematics, vol. 150 (Springer-Verlag, New York, 1995). https://doi.org/10.1007/978-1-4612-5350-1

[19] The Stacks Project Authors, Fitting ideals, *The Stacks Project*, More on Algebra, Sect. 15.8, Tags 07Z6 and 07Z9 (accessed 14 August 2026). https://stacks.math.columbia.edu/tag/07Z6

[20] L. Le Cam, *Asymptotic Methods in Statistical Decision Theory* (Springer-Verlag, New York, 1986). https://doi.org/10.1007/978-1-4612-4946-7

[21] C.K. Chow, On optimum recognition error and reject tradeoff, *IEEE Trans. Inform. Theory* **16**(1), 41–46 (1970). https://doi.org/10.1109/TIT.1970.1054406

[22] J.R. Rajput, S. Weinmueller, J. Endres, P. Dawood, F. Knoll, A. Maier, and M. Zaiss, Death by retrospective undersampling—caveats and solutions for learning-based MRI reconstructions, in *Medical Image Computing and Computer Assisted Intervention—MICCAI 2024*, Lecture Notes in Computer Science, vol. 15007, pp. 233–241 (Springer, Cham, 2024). https://doi.org/10.1007/978-3-031-72104-5_23

[23] R.F. Barber, E.J. Candès, A. Ramdas, and R.J. Tibshirani, Conformal prediction beyond exchangeability, *Ann. Statist.* **51**(2), 816–845 (2023). https://doi.org/10.1214/23-AOS2276

[24] V.-V. Wettenhovi, V. Kolehmainen, M. Kettunen, O. Gröhn, and M. Vauhkonen, State estimation of time-varying MRI with radial golden angle sampling, *J. Math. Imaging Vis.* **64**(8), 825–844 (2022). https://doi.org/10.1007/s10851-022-01095-x

[25] S.J. Inati, J.D. Naegele, N.R. Zwart, V. Roopchansingh, M.J. Lizak, D.C. Hansen, C.-Y. Liu, D. Atkinson, P. Kellman, S. Kozerke, H. Xue, A.E. Campbell-Washburn, T.S. Sørensen, and M.S. Hansen, ISMRM raw data format: a proposed standard for MRI raw datasets, *Magn. Reson. Med.* **77**(1), 411–421 (2017). https://doi.org/10.1002/mrm.26089

[26] E. Solomon, P.M. Johnson, Z. Tan, R. Tibrewala, Y.W. Lui, F. Knoll, L. Moy, S.G. Kim, and L. Heacock, fastMRI Breast: a publicly available radial k-space dataset of breast dynamic contrast-enhanced MRI, *Radiol. Artif. Intell.* **7**(1), e240345 (2025). https://doi.org/10.1148/ryai.240345

[27] E. Levine and B. Hargreaves, On-the-fly adaptive k-space sampling for linear MRI reconstruction using moment-based spectral analysis, *IEEE Trans. Med. Imaging* **37**(2), 557–567 (2018). https://doi.org/10.1109/TMI.2017.2766131

[28] P.B. Roemer, W.A. Edelstein, C.E. Hayes, S.P. Souza, and O.M. Mueller, The NMR phased array, *Magn. Reson. Med.* **16**(2), 192–225 (1990). https://doi.org/10.1002/mrm.1910160203

[29] K.P. Pruessmann, M. Weiger, M.B. Scheidegger, and P. Boesiger, SENSE: sensitivity encoding for fast MRI, *Magn. Reson. Med.* **42**(5), 952–962 (1999). https://doi.org/10.1002/(SICI)1522-2594(199911)42:5⟨952::AID-MRM16⟩3.0.CO;2-S

[30] A. Haase, J. Frahm, D. Matthaei, W. Hänicke, and K.-D. Merboldt, FLASH imaging: rapid NMR imaging using low flip-angle pulses, *J. Magn. Reson.* **67**(2), 258–266 (1986). https://doi.org/10.1016/0022-2364(86)90433-6

[31] M.C. Schabel and D.L. Parker, Uncertainty and bias in contrast concentration measurements using spoiled gradient echo pulse sequences, *Phys. Med. Biol.* **53**(9), 2345–2373 (2008). https://doi.org/10.1088/0031-9155/53/9/010

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- ESM1.zip