

Contents

Gradient-Guided Dimensionality Reduction for Ideal Observers: Conjugate Gradient Channels as a Signal-Processing Bridge Between Tractability and Optimality in Medical Image Quality Assessment	1
Abstract	1
Introduction	2
Background	3
Synthesis	4
Discussion	8
Limitations	9
Selection Process	10
Weakly-Connected Addendum	10
Conclusion	10
Response to Review	11

Gradient-Guided Dimensionality Reduction for Ideal Observers: Conjugate Gradient Channels as a Signal-Processing Bridge Between Tractability and Optimality in Medical Image Quality Assessment

Saluca Agentic AI Research Team — Saluca LLC. AI-drafted synthesis from recent arXiv preprints; for human review, not peer-reviewed.

Abstract

Task-based image quality assessment (IQ) in medical imaging faces a structural tension: the theoretical optimality of Bayesian and Hotelling ideal observers is well-established, yet their direct application to high-dimensional imaging data is computationally intractable. The gap between principled statistical detection theory and practical system optimisation is not merely an engineering inconvenience — it is a fundamental bottleneck that determines whether objective figures of merit (FOMs) can actually guide hardware and algorithm design. This synthesis examines a candidate resolution: the use of conjugate gradient (CG)-based channel construction to perform efficient, task-relevant dimensionality reduction that preserves the performance-ranking properties of ideal observers while remaining computationally feasible.

Drawing primarily from a single focused contribution in eess.IV [arXiv:2605.29415](https://arxiv.org/abs/2605.29415), we situate the CG channel method within the broader signal-processing landscape of ideal observer theory, dimensionality reduction, and detection-theoretic optimality. The thesis advanced here — offered explicitly as a heuristic structural reading rather than a derivation from the abstract’s stated content — is that the conjugate gradient method, when applied to channel construction, constitutes a principled iterative projection of the high-dimensional detection problem onto a sequence of Krylov subspaces that monotonically improves observer approximation quality. This framing bridges Krylov-subspace numerical linear algebra, Bayesian detection theory, and medical imaging system optimisation in a way that is mechanistically argued rather than merely analogical, though it remains unconfirmed by the abstract itself.

The falsification path is concrete: if CG-constructed channels fail to monotonically improve Hotelling observer SNR approximations as channel dimensionality increases, the Krylov-subspace interpretation collapses. Similarly, if the channels constructed do not outperform standard principal-component or Laguerre-Gauss channel families on matched detection tasks, the efficiency claim is falsified. Limitations include abstract-only reading of the primary source, a single-paper corpus, the absence of empirical benchmark comparisons from the abstract itself, and uncertainty about whether results in the full paper are validated on simulated or real scanner data, and in 2D or 3D image spaces.

Introduction

Medical imaging systems — CT scanners, MRI machines, digital radiography platforms — are designed, optimised, and regulated on the basis of image quality metrics. Yet the dominant practical metrics (SNR, contrast-to-noise ratio, sharpness) are observer-agnostic: they do not directly measure how well a human radiologist or an automated classifier can perform a specific detection task on the resulting images. This mismatch between engineering convenience and clinical relevance has motivated decades of work in task-based image quality assessment [arXiv:2605.29415](#).

The theoretical solution is elegant: replace heuristic metrics with figures of merit derived from optimal statistical observers. The Bayesian Ideal Observer (IO) maximises the area under the ROC curve for a given signal detection task, providing a gold-standard FOM that is entirely determined by the imaging system’s statistical properties and the task specification. The Hotelling Observer (HO), the optimal linear observer, provides a tractable approximation when the full posterior is unavailable or when linear observers are appropriate models for human performance in specific tasks.

The practical problem is severe: both the IO and the HO require operating on the full image vector. For modern medical images — volumetric CT data with millions of voxels, or high-resolution digital radiographs — direct computation of the HO requires inverting or decomposing covariance matrices of dimension equal to the number of image pixels or voxels. For a 512×512 image, this is a 262,144-dimensional covariance matrix. Its storage alone requires tens of gigabytes; its eigendecomposition is computationally prohibitive for routine system optimisation workflows.

The channel mechanism addresses this by projecting image data through a dimensionality-reducing linear transformation (the channel) before applying the observer. If the channel is well-chosen, the channelised Hotelling observer (CHO) closely approximates the full HO, and the computational cost drops to the manageable problem of inverting a small channel-space covariance matrix. The central question — the motivating question of this synthesis — is: *how should channels be constructed to maximally preserve observer performance while minimally expanding computational cost?*

The recent contribution reported in [arXiv:2605.29415](#) proposes a conjugate gradient (CG)-based method for constructing such channels. This paper introduces and situates that contribution. The interpretive thesis — that CG channel construction exploits the Krylov subspace geometry of the detection problem — is advanced here as a heuristic structural reading of the CG method’s known mathematical properties applied to the channel construction setting. It is not derived from the abstract of [arXiv:2605.29415](#), which does not invoke Krylov subspaces explicitly, and it should be treated as a candidate mechanistic explanation pending validation from the full paper’s content.

Background

2.1 Ideal Observers and the Detection-Theoretic Framework

The Bayesian Ideal Observer solves the binary hypothesis testing problem optimally in the sense of maximising the area under the receiver operating characteristic (ROC) curve, integrated over all possible decision thresholds. For a signal detection task — “is this lesion present?” — the IO computes the likelihood ratio of the observed image data under the signal-present and signal-absent hypotheses. The area under the ROC curve (AUC) produced by the IO is the highest achievable by any observer, making it an absolute upper bound on system performance for the specified task.

The Hotelling Observer specialises this to the class of linear decision rules. It computes a scalar test statistic as the inner product of the image vector with an optimal template — the Hotelling template — and thresholds this statistic. The HO template is the product of the inverse covariance matrix (of the image data under the background hypothesis) with the mean difference signal. The resulting SNR, the Hotelling SNR (SNR_H), serves as the FOM for the HO. When image data are multivariate Gaussian, the HO is equivalent to the IO, and SNR_H is a sufficient statistic for system comparison.

Both observers are defined in the full image space. Their practical intractability in high dimensions is not a limitation of the theory but a computational constraint imposed by the curse of dimensionality. The abstract of [arXiv:2605.29415](#) characterises this intractability as occurring “often” — a qualifier preserved here, as the abstract does not claim universal intractability across all imaging configurations.

2.2 Channel Mechanisms and Dimensionality Reduction

A channel is a linear operator \mathbf{T} that maps the high-dimensional image vector $\mathbf{g} \in \mathbb{R}^N$ to a low-dimensional channel output $\mathbf{v} = \mathbf{T} \mathbf{g} \in \mathbb{R}^M$, with $M \leq N$. The channelised Hotelling observer (CHO) then applies the HO in the M -dimensional channel space, requiring inversion of an $M \times M$ covariance matrix rather than an $N \times N$ matrix. The quality of the CHO as an approximation to the full HO depends entirely on how well the channel \mathbf{T} captures the task-relevant structure of the image data.

Classical channel families include difference-of-Gaussians (DOG) channels, Laguerre-Gauss (LG) channels, and principal component analysis (PCA)-derived channels. DOG and LG channels are fixed, task-agnostic filter banks motivated by models of human visual processing. PCA channels are data-adaptive but optimise for variance explained rather than for detection task performance. None of these approaches directly optimises the channel for the observer approximation quality it produces.

The channel construction problem can be stated precisely: find \mathbf{T} such that the CHO SNR approximates the full HO SNR as closely as possible for a given number of channels M . This is a constrained optimisation problem in the space of linear operators, and its solution structure depends on the spectral properties of the image covariance matrix and the mean difference signal [arXiv:2605.29415](#).

2.3 Conjugate Gradient Methods and Krylov Subspaces

The conjugate gradient (CG) method is an iterative algorithm for solving symmetric positive definite linear systems $\mathbf{A} \mathbf{x} = \mathbf{b}$. Its key property, relevant here, is that after k iterations, the CG iterate lies in the Krylov subspace $\mathcal{K}_k(\mathbf{A}, \mathbf{b}) = \text{span}\{\mathbf{b}, \mathbf{A} \mathbf{b}, \mathbf{A}^2 \mathbf{b}, \dots, \mathbf{A}^{k-1} \mathbf{b}\}$. Krylov subspaces have a natural optimality property: the CG iterate minimises the \mathbf{A} -norm of the error over all

vectors in \mathcal{K} . Equivalently, the CG method extracts the most task-relevant directions of \mathbf{A} as determined by \mathbf{b} , in a greedy, monotonically improving sequence.

When the linear system being solved is the HO template equation — $\mathbf{K} \mathbf{b} = \Delta$, where \mathbf{K} is the background covariance matrix and Δ is the mean difference signal — the Krylov subspace $\mathcal{K}(\mathbf{K}, \Delta)$ captures the directions most relevant to the detection task. The first basis vector is Δ itself (the matched filter direction); subsequent vectors are $\mathbf{K} \Delta$, $\mathbf{K}^2 \Delta$, etc., which encode how the covariance structure of the background “bends” the optimal template away from the naive matched filter. This is the mechanistic reason to expect CG-derived channels to be efficient for detection tasks: they are, by construction, aligned with the Krylov subspace that the HO template inhabits. The argument is mechanistic — it follows from the algebraic definition of the Krylov subspace and the structure of the HO template equation — rather than analogical.

Important caveat on scope of this interpretation. The abstract of [arXiv:2605.29415](#) does not use the term “Krylov subspace” and does not confirm that the channel vectors are the CG iterates or their span. The Krylov-subspace framing advanced in this synthesis is a heuristic reading of the CG method’s known mathematical properties applied to the channel construction setting. It is consistent with the abstract’s claims but is not derived from them. Whether the authors’ implementation exploits Krylov structure explicitly or uses CG as a general-purpose optimiser is not determinable from the abstract alone. Additionally, the abstract does not specify whether the method is evaluated on 2D or 3D image data, or on simulated versus real scanner acquisitions; these distinctions may affect the generalisability of any results reported in the full paper.

Synthesis

3.1 Claim: CG Channel Construction Exploits Task-Relevant Spectral Structure

Claim. The conjugate gradient method, when applied to the HO template equation, generates a sequence of channel vectors that span increasingly complete approximations to the HO template’s spectral support. Each CG iteration adds one dimension to the channel space, and the added dimension is chosen to maximally reduce the residual in the HO template computation. This makes CG channels plausibly better-suited than fixed channel families (DOG, LG) for detection tasks where the background covariance is not well-modelled by the fixed family’s spectral assumptions — though whether this theoretical advantage translates to empirical superiority is not established from the abstract.

Evidence. The abstract of [arXiv:2605.29415](#) states that the CG-based method constructs “efficient channels for approximating the IO and HO performance,” implying that the resulting channels achieve high approximation quality relative to their dimensionality. The framing of the contribution as a method for dimensionality reduction that “facilitates the computation of ideal observers” is consistent with the claim that channel efficiency — performance per channel dimension — is the operative figure of merit.

Caveat. The abstract does not report specific numerical comparisons between CG channels and competing channel families (DOG, LG, PCA). Whether CG channels outperform these alternatives on standard detection tasks, and by how much, is not established from the abstract alone. The claim of structural advantage is a theoretical prediction derived from the Krylov-subspace analysis above; it requires empirical validation from the full paper’s results. The abstract also does not

specify the imaging modality, signal type, background model, or image dimensionality (2D vs. 3D) used in evaluation.

Falsification path. If experimental results in the full paper show that CG-constructed channels do not outperform LG or PCA channels in approximating HO SNR for matched detection tasks (e.g., signal-known-exactly detection in a lumpy background), the structural advantage claim is falsified. Specifically: if the CHO SNR achieved with M CG channels is not monotonically non-decreasing in M , or does not converge faster (in M) than PCA channels to the full HO SNR on at least one reported task, the Krylov-subspace efficiency argument fails.

3.2 Claim: Channel Dimensionality and Approximation Quality Are Linked Through the Krylov Convergence Rate

Claim. The number of CG channels M required to achieve a given approximation quality (e.g., CHO SNR within % of full HO SNR) is plausibly determined by the effective rank of the detection problem — specifically, by the number of significant eigencomponents of the background covariance matrix \mathbf{K} that project onto the mean difference signal $\mathbf{\Delta}$. When this effective rank is low (as in many medical imaging tasks where the signal occupies a small spatial region relative to the image), a small number of CG channels may suffice for high-fidelity approximation. This is a theoretical prediction of the Krylov-subspace reading, not a finding reported in the abstract.

Evidence. The abstract of [arXiv:2605.29415](https://arxiv.org/abs/2605.29415) identifies computational intractability as the central problem motivating channel-based dimensionality reduction, and presents the CG method as a solution. The implicit claim is that the resulting channels are “efficient” — that is, they achieve good approximation with small M . This is consistent with the Krylov convergence argument: if the detection problem has low effective rank, CG converges rapidly, and few channels are needed.

Caveat. The abstract does not specify the imaging modalities, signal types, or background models used to evaluate the method. Whether the low-effective-rank assumption holds broadly across medical imaging tasks — or only for specific task configurations (e.g., detecting a known signal in a stationary Gaussian background) — is not established from the abstract. The abstract reports the method applies to both IO and HO approximation; the mechanisms by which CG channels approximate the IO (which is nonlinear) may differ from the HO case (which is linear), and this distinction is not resolved from the abstract alone.

Falsification path. For tasks where the background covariance has high effective rank relative to the signal (e.g., detection in highly structured, non-stationary backgrounds such as breast tissue in digital mammography), the Krylov convergence rate will be slow, and many CG channels will be required. If empirical results show that CG channel count for acceptable approximation scales with image dimension N rather than with the signal’s spectral support, the low-effective-rank efficiency claim is falsified.

3.3 Claim: The CG Framework Provides a Candidate Unified Approximation Pathway for Both IO and HO

Claim. While the IO and HO are distinct observers (nonlinear vs. linear), the CG framework may provide a unified computational pathway for approximating both. For the HO, CG directly

solves the template equation. For the IO, CG-constructed channels reduce the dimensionality of the image space before nonlinear IO computation (e.g., via Markov chain Monte Carlo or Parzen window estimation), making the nonlinear computation feasible in the reduced channel space. This claim is offered as a candidate reading of the abstract’s statement; the mechanism for IO approximation is not described in the abstract, and the conditions under which it holds are not established.

Evidence. The abstract of [arXiv:2605.29415](#) explicitly states that the CG-based method constructs channels for “approximating the IO and HO performance,” indicating that the authors claim applicability to both observer types. The framing of channels as a “framework for dimensionality reduction that can facilitate the computation of ideal observers” (plural) is consistent with the unified pathway interpretation.

Structural condition required. For CG channels to serve as a valid approximation pathway for the IO, they must preserve sufficient information about the image likelihood ratio — not merely about the HO template. This is a stronger condition than HO approximation: the channel must act as an approximate sufficient statistic for the detection task, not merely capture the linear task-relevant subspace. The abstract does not establish whether this condition is met. The argument for why CG channels might satisfy it is as follows: if the IO likelihood ratio is well-approximated by a function of the HO template (which holds when image data are approximately Gaussian), then channels that accurately span the HO template subspace will also approximately preserve the IO likelihood ratio. This argument depends on the Gaussianity assumption and breaks down for strongly non-Gaussian backgrounds. The abstract does not specify the background model, so whether this condition applies to the reported results is unknown.

Caveat. The mechanism by which CG channels facilitate IO approximation is not described in the abstract. The stronger claim — that CG channels are sufficient statistics for the IO — is speculative and not supported by the abstract. The unified pathway claim should be treated as a hypothesis to be tested against the full paper’s results.

Falsification path. If the IO AUC computed in the CG channel space is significantly lower than the full-space IO AUC for tasks where the HO approximation is accurate (i.e., CHO SNR > HO SNR), this would indicate that the CG channels capture the linear task-relevant structure but miss nonlinear structure needed for IO approximation. This would falsify the unified pathway claim and would suggest that separate channel construction strategies are needed for IO and HO approximation.

3.4 Claim: The CG Channel Method Addresses a Structural Bottleneck, Not Merely a Computational Convenience

Claim. The intractability of ideal observers in high-dimensional image spaces is not merely a practical inconvenience to be engineered around — it is a structural bottleneck that determines whether task-based IQ assessment can be integrated into iterative system optimisation workflows. The CG channel method, by providing an efficient and principled dimensionality reduction, may enable a feedback loop between system parameter changes and observer performance estimates that was previously computationally prohibitive.

Evidence. The abstract of [arXiv:2605.29415](#) frames the problem as follows: ideal observer application to high-dimensional image data is “often computationally intractable” and presents channels as providing “an effective framework for dimensionality reduction that can facilitate the computa-

tion of ideal observers.” The language of “facilitate” and “effective framework” suggests that the authors view the contribution as enabling a class of applications (system design and optimisation) that was previously blocked by computational cost.

Caveat. The abstract does not describe integration of the CG channel method into an iterative system optimisation loop. Whether the method is fast enough for real-time or near-real-time system optimisation (as opposed to offline post-hoc evaluation) is not established. The computational cost of constructing CG channels — which requires matrix-vector products with the covariance matrix \mathbf{K} , itself potentially expensive to compute or store — is not quantified in the abstract. The abstract’s use of “often” in “often computationally intractable” also signals that the bottleneck is not universal; for lower-dimensional imaging configurations, existing non-CG methods may already be tractable.

Falsification path. If the computational cost of CG channel construction (including covariance matrix estimation from training data) scales unfavorably with image dimension N — e.g., if it requires $O(N^2)$ storage or $O(N^2 M)$ computation — then the method does not resolve the structural bottleneck for the highest-dimensional imaging tasks (e.g., 3D CT). A concrete test: measure wall-clock time for CG channel construction and CHO computation as a function of image dimension N for fixed M , and compare to the cost of full HO computation. If the ratio does not decrease substantially with N , the bottleneck claim is weakened.

3.5 Claim: Efficient Channel Construction May Be Among the Rate-Limiting Steps for Objective IQ Metrics in Clinical Deployment

Claim. A plausible implication of the CG channel contribution is that efficient channel construction is among the rate-limiting steps preventing wider adoption of task-based IQ metrics in clinical imaging system design and regulatory evaluation. Progress on channel construction methods may directly expand the applicability of the full ideal observer framework. This claim is offered as a candidate hypothesis, not as a finding established by the abstract.

Evidence. The abstract of [arXiv:2605.29415](#) identifies computational intractability as a primary barrier: ideal observer application to high-dimensional image data is “often computationally intractable.” The proposed CG method is presented as addressing this specific barrier. The framing implies that ideal observer theory is mature and that the channel construction problem is an active frontier.

Caveat. This claim extrapolates beyond what the abstract establishes. Other barriers to clinical adoption of task-based IQ metrics — including task specification (what detection task is clinically relevant?), background model selection (what statistical model for anatomical variability is appropriate?), and regulatory acceptance of observer-based metrics — are not addressed in the abstract and may be equally or more limiting in practice. The abstract-only reading does not allow assessment of these competing bottlenecks. The claim that channel construction is *the* rate-limiting step is not supportable from the abstract; the hedged form — that it is *among* the rate-limiting steps — is the most the abstract warrants.

Falsification path. A systematic review of published task-based IQ studies that were abandoned or incomplete would falsify the channel-construction bottleneck hypothesis if it showed that most failures were attributable to task specification or background model inadequacy rather than computational cost. A more targeted test: for imaging configurations where existing non-CG channel

methods (LG, DOG) are already computationally tractable and yield HO approximation quality within accepted tolerances, the marginal value of CG channels is reduced, and the rate-limiting-step claim is weakened for those configurations. Specifically, if a survey of the eess.IV/medical physics literature finds that fewer than half of reported ideal observer studies cite computational cost as the primary limitation, the claim that channel construction is a dominant bottleneck requires revision.

Discussion

What This Synthesis Implies

The CG channel construction method reported in [arXiv:2605.29415](#) represents a structurally motivated approach to a long-standing computational bottleneck in task-based medical image quality assessment. The Krylov-subspace interpretation advanced here — that CG channels span the subspace most relevant to the detection task, as determined by the interaction of the background covariance and the mean difference signal — provides a mechanistic candidate explanation for why the method should be efficient, in the sense of achieving high observer approximation quality with few channel dimensions. The mechanistic argument is grounded in the algebraic properties of the CG method and the structure of the HO template equation, not in surface-level vocabulary similarity between the two domains.

If the Krylov-subspace interpretation is correct, it has implications beyond the specific CG implementation. It suggests that any iterative solver that generates Krylov subspace approximations (MINRES, GMRES, Lanczos iteration) could be used for channel construction, with the choice of solver determined by the spectral properties of the covariance matrix. For symmetric positive definite covariance matrices (the standard case in medical imaging), CG is optimal among Krylov methods. For ill-conditioned or indefinite cases, other Krylov methods may be preferable.

The unified IO/HO pathway claim, if validated, would significantly simplify the computational infrastructure needed for task-based IQ assessment: a single channel construction step would enable both linear and nonlinear observer computation in the reduced channel space. This would make it practical to compare HO and IO performance on the same imaging system and task, providing insight into whether nonlinear observer effects are significant for a given system configuration.

What This Synthesis Does NOT Imply

This synthesis does not claim that CG channels are optimal in any formal sense beyond the Krylov subspace approximation to the HO template. The IO approximation quality of CG channels is not established from the abstract, and the claim that CG channels are sufficient statistics for the IO is speculative.

This synthesis does not claim that the CG channel method resolves all barriers to clinical adoption of task-based IQ metrics. Task specification, background modelling, and regulatory acceptance remain open problems that are orthogonal to channel construction efficiency.

The Krylov-subspace interpretation advanced in the Background section is a heuristic reading of the CG method’s properties applied to the channel construction problem. It is not derived from the abstract of [arXiv:2605.29415](#), which does not use the term “Krylov subspace.” The interpretation is consistent with the abstract’s claims but is not confirmed by them. Whether

the authors’ implementation exploits Krylov structure explicitly or uses CG as a general-purpose optimiser is not determinable from the abstract alone.

This synthesis should not be read as endorsing any specific regulatory or clinical deployment pathway for CG-based IQ metrics. The gap between a published preprint demonstrating a computational method and a validated, deployed clinical tool is substantial and involves considerations entirely outside the scope of this synthesis.

Limitations

Abstract-only reading. The primary methodological limitation of this synthesis is that it is based entirely on the abstract of [arXiv:2605.29415](#). The full paper may contain results, comparisons, and caveats that significantly modify the claims made here. All numerical claims, comparisons, and mechanistic interpretations are either directly stated in the abstract or are theoretical predictions that require validation from the full paper. This limitation is acknowledged throughout the Synthesis section with explicit hedges.

Single-paper corpus. This synthesis is built on a single arXiv preprint. The corpus breadth is minimal, and the synthesis is correspondingly narrow. The broader context of ideal observer theory, channel methods, and medical imaging system optimisation is drawn from general signal-processing and detection-theory knowledge, not from additional corpus papers. The absence of comparative papers means that claims about CG channels outperforming competing methods are theoretical predictions, not corpus-supported findings.

Krylov-subspace interpretation as heuristic. The mechanistic bridge between CG channel construction and Krylov subspace theory is a heuristic reading, not a derivation from the abstract. The abstract does not invoke Krylov subspaces, and the interpretation depends on assumptions (that the CG method is applied directly to the HO template equation, that the channel vectors are the CG iterates or their span) that are not confirmed by the abstract.

IO approximation mechanism unestablished. The claim that CG channels provide a unified pathway for IO and HO approximation is supported by the abstract’s statement that the method applies to both, but the mechanism for IO approximation is not described in the abstract. The stronger claim — that CG channels are sufficient statistics for the IO — is speculative and not supported by the abstract.

Preprint status. [arXiv:2605.29415](#) is a preprint as of the corpus date (2026-05-28) and has not been peer-reviewed. All findings should be treated as preliminary until peer review is complete.

Dimensionality and data type unspecified. The abstract does not specify whether the method is evaluated on 2D or 3D image data, or on simulated versus real scanner acquisitions. These distinctions are material: results obtained in 2D simulated backgrounds may not generalise to 3D real-scanner data, where covariance estimation is substantially harder and background statistics are less well-characterised.

No empirical benchmarks from abstract. The abstract does not report specific numerical results (e.g., CHO SNR as a fraction of HO SNR, number of channels required, wall-clock times). All quantitative claims in this synthesis are theoretical predictions or structural arguments, not reported measurements.

“Often” qualifier preserved. The abstract’s characterisation of ideal observer computation as “often computationally intractable” — not universally so — is preserved throughout this synthesis. Claims about the bottleneck are scoped accordingly.

Selection Process

The corpus provided for this synthesis consisted of a single paper from the eess.IV category (image and video processing), specifically [arXiv:2605.29415](#), covering a 30-day recency window. No filtering was required or possible: the candidate pool size was one paper. No papers were dropped.

The kept paper is directly relevant to the synthesis topic: it addresses computational methods for ideal observer approximation in medical image quality assessment, which is a core problem in the eess.IV/eess.SP intersection. The paper’s focus on channel construction as a dimensionality reduction mechanism for ideal observers is a well-defined technical contribution with clear connections to signal detection theory, numerical linear algebra, and medical imaging system optimisation.

The narrow corpus is acknowledged as a significant limitation. The synthesis compensates by grounding the theoretical context (ideal observer theory, Krylov subspaces, channel mechanisms) in well-established signal-processing and numerical linear algebra concepts, while being explicit that these contextual claims are not corpus-supported. All claims that go beyond the single abstract are clearly hedged as theoretical predictions or heuristic interpretations.

Weakly-Connected Addendum

No weakly-connected sources are present in this synthesis, as the corpus contains only one paper. All theoretical context (Krylov subspaces, HO template equations, channel families) is drawn from background knowledge and is clearly distinguished from corpus-supported claims. There are no analogical bridges to other domains in this synthesis; all connections are within the detection theory and numerical linear algebra domains that directly underpin the CG channel construction method.

Conclusion

The conjugate gradient channel construction method reported in [arXiv:2605.29415](#) addresses a structural bottleneck — one that the abstract characterises as occurring “often” rather than universally — in task-based medical image quality assessment: the computational intractability of ideal observers in high-dimensional image spaces. This synthesis has argued, through a heuristic Krylov-subspace reading of the CG method’s properties, that CG channels are structurally well-motivated for detection tasks — they span the subspace most relevant to the Hotelling observer template, as determined by the interaction of the background covariance matrix and the mean difference signal. The mechanistic argument is grounded in the algebraic structure of the CG method and the HO template equation, not in surface vocabulary similarity. The synthesis has identified concrete falsification paths for each major claim, acknowledged the abstract-only reading as a primary methodological limitation, restored the abstract’s “often” qualifier throughout, and noted that the broader barriers to clinical adoption of task-based IQ metrics extend well beyond

channel construction efficiency. The contribution, if validated by the full paper’s results — including specification of the imaging dimensionality and data type used in evaluation — represents a candidate step toward making principled, observer-based image quality metrics computationally accessible for routine medical imaging system design and optimisation workflows.

Response to Review

Heuristic framing in the Introduction. The v1 Introduction presented the Krylov-subspace thesis as though it followed from the problem’s structure, without naming it as a heuristic reading. The v2 Introduction now explicitly states that the Krylov-subspace framing is “a heuristic structural reading of the CG method’s known mathematical properties applied to the channel construction setting” and is “not derived from the abstract.” This matches the hedging already present in the Abstract and Background but absent from the Introduction.

Abstract overshoot in §3.5. The v1 claim that efficient channel construction is “*the* rate-limiting step” extrapolated well beyond what the abstract establishes. The v2 downgrades this to “among the rate-limiting steps,” adds explicit acknowledgement that task specification, background modelling, and regulatory acceptance may be equally limiting, and sharpens the falsification path from a vague “survey of the literature” to a specific empirical criterion (fewer than half of reported ideal observer studies citing computational cost as the primary limitation).

Asserted analogy in §3.3. The v1 unified IO/HO pathway claim asserted that CG channels work for IO approximation without explaining the structural condition required. The v2 adds an explicit “Structural condition required” paragraph naming the approximate sufficient statistic condition, explaining why it might hold (approximate Gaussianity of image data), and specifying when it breaks down (strongly non-Gaussian backgrounds). The claim is also reframed as a “candidate” reading throughout.

Dropped source caveats. The v1 softened the abstract’s “often computationally intractable” to unqualified “is intractable” in several places, and did not flag the absence of 2D/3D or simulated/real-scanner specification. The v2 restores the “often” qualifier throughout the body and Conclusion, adds a dedicated Limitations bullet on dimensionality and data type, and flags this in the Abstract’s limitations list.

“Structurally superior” language in §3.1. The v1 described CG channels as “structurally superior” to fixed channel families, which is stronger than the abstract supports. The v2 replaces this with “plausibly better-suited” and adds an explicit hedge that empirical validation is required.

Falsification path in §3.5. The v1 falsification path for the rate-limiting-step claim referenced “a survey of the medical imaging system optimisation literature” without specifying what outcome would constitute falsification. The v2 adds a concrete criterion: if fewer than half of reported ideal observer studies in the eess.IV/medical physics literature cite computational cost as the primary limitation, the claim requires revision.