

Prospective frame certification and deterministic spectral mask design for nonuniform Fourier-domain optical coherence tomography

Fujio Takenari

`fujio-takenari-ni@alumni.osaka-u.ac.jp`

Independent Researcher <https://orcid.org/0009-0005-8549-0022>

Research Article

Keywords: 42C15, 94A20, 65T50

Posted Date: October 20th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11423787/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: The authors declare no competing interests.

Prospective frame certification and deterministic spectral mask design for nonuniform Fourier-domain optical coherence tomography

Double Anonymous

Abstract

Fourier-domain optical coherence tomography (FD-OCT) reconstructs axial reflectivity from spectral Fourier measurements, but practical acquisition grids are generally nonuniform in wavenumber and may also be deliberately thinned. Frame-based reconstruction from nonuniform OCT samples is established; the unresolved acquisition-side question is whether a proposed spectral mask can be certified as stably invertible before a reconstruction algorithm is chosen. We formulate a finite weighted Fourier model in which the lower frame bound serves as a prospective stability certificate. A known diagonal unit-modulus dispersion phase leaves the frame bounds unchanged, whereas phase-model mismatch obeys an explicit condition-number-weighted error bound. We then introduce frame-certified gap filling (FCGF), a deterministic mask heuristic that combines an unweighted column-pivoted QR core with midpoint filling of the largest admissible wavenumber gaps, followed by frame-bound certification. In a controlled benchmark with 192 wavelength-uniform candidates and a 48-dimensional reconstruction space, FCGF at critical sampling ($M=N=48$) achieved median normalized reconstruction error 0.0179, compared with 0.0203 for nearest k-uniform sampling and 0.0506 for detector-index thinning. Its advantage diminished under oversampling. Across ordinary random masks, the sample-count-normalized lower frame bound strongly anticorrelated with reconstruction error (pooled Spearman $\rho = -0.994$), whereas a coverage-controlled stratified-random baseline largely removed catastrophic failures. The study identifies a narrow but practical role for Fourier-frame certification in acquisition design; it does not replace established nonuniform reconstruction or sparsity-based methods.

¹ Independent Researcher; Alumni, Department of Biophysical Engineering, Graduate School of Medicine, Osaka University, Japan

² Independent Researcher, Hyogo, Japan

Corresponding author: fujio-takenari-ni@alumni.osaka-u.ac.jp

Mathematics Subject Classification: 42C15; 94A20; 65T50

1 Introduction

Let $x \in \mathbb{C}^N$ denote a finite complex axial-reflectivity profile and let its retained spectral measurements be samples of a Fourier map. In Fourier-domain optical coherence tomography (FD-OCT), the physical coordinate is optical wavenumber $k = 2\pi/\lambda$. Grating spectrometers are approximately uniform in wavelength rather than wavenumber, while swept sources may follow nonlinear tuning trajectories. Before k -linearization or resampling, the data are therefore collected at non-equispaced Fourier nodes. Deliberate spectral thinning introduces a second, design-dependent source of nonuniformity. The transition from time-domain to spectral-domain and swept-source OCT exploited the sensitivity and speed advantages of Fourier-domain detection [6,7,11]; nonuniform resampling and reconstruction have consequently become standard numerical concerns [5,10,15], and undersampled OCT has also been studied with sparsity-based recovery [16,17].

Three questions should be separated. **Fast evaluation** asks how to compute forward and adjoint trigonometric sums efficiently; NUFFT algorithms address this cost but do not improve the conditioning of an inverse problem whose sampling geometry is poor [5,8,10,15]. **Prior-driven recovery** asks whether missing data can be compensated by structural assumptions such as sparsity, usually through nonlinear optimization [12,13,16,17]. **Sampling stability** asks whether a prescribed set of nodes stably samples a specified reconstruction space before any image is reconstructed. The present paper addresses only the third question, in a finite-dimensional linear setting.

Stable recovery from nonuniform Fourier measurements has been developed through generalized sampling and weighted exponential frames [1-3]. Frame theory has also been applied directly to OCT: Nagib et al. reconstructed A-scans from redundant, nonuniformly spaced spectral samples using a scaled nonuniform discrete Fourier transform [14]. The novelty claimed here is therefore not Fourier-frame reconstruction itself, NUFFT/NUDFT, compressed sensing, least-squares stability, or Gershgorin theory. The acquisition-side problem is narrower: given a finite candidate grid and a target reconstruction space, can the geometry of a

proposed spectral mask be certified before reconstruction, and can that certificate guide a deterministic thinning rule?

We make four contributions. First, we use the lower frame bound of the weighted sampling matrix as a prospective, signal-independent certificate for a proposed mask. Second, we separate sampling geometry from dispersion calibration: any known diagonal unit-modulus phase leaves the frame bounds invariant, while a mismatched phase model obeys an explicit reconstruction-error bound. Third, we state the standard noise and Gershgorin estimates needed to interpret the certificate. Fourth, we propose **frame-certified gap filling (FCGF)**, which selects a critical core by unweighted column-pivoted QR, fills the largest admissible nonperiodic gaps in normalized wavenumber, and then audits the completed mask through its frame bounds. The method is a transparent heuristic, not a global E-optimal design.

All experiments are controlled and synthetic. They test the mathematical relation between node geometry, frame stability and reconstruction error. They do not constitute experimental or clinical OCT validation.

2 Finite Fourier model for nonuniform FD-OCT

2.1 Measurement model and modeling assumptions

The model assumes that preprocessing has isolated a background-subtracted complex interferometric term, as in a complex-conjugate-resolved or full-range representation. Raw real-valued interferograms, DC terms, autocorrelation terms and mirror-image suppression are outside the model.

Let the candidate wavelengths be λ_c , $c = 1, \dots, C$, with wavenumbers $k_c = 2\pi/\lambda_c$. We map them to the normalized Fourier coordinate

$$\tau_c = \frac{k_c - k_{\min}}{k_{\max} - k_{\min}} \frac{C - 1}{C} \in \left[0, \frac{C - 1}{C}\right],$$

where the factor $(C - 1)/C$ avoids duplicating the periodic endpoint. For the numerical model, the centered depth-grid indices are

$$v_\ell = \ell - \frac{N - 1}{2}, \quad \ell = 0, \dots, N - 1.$$

For a selected index set $\Omega \subset \{1, \dots, C\}$ of size M , the unweighted Fourier matrix is

$$F_{j\ell} = \exp(-2\pi i \tau_j v_\ell), \quad j \in \Omega.$$

Positive density-compensation weights μ_j define the weighted sampling matrix

$$\Phi = W^{1/2} F, \quad W = \text{diag}(\mu_j).$$

The unweighted measurements are modeled as

$$z = D_\phi F x + \xi,$$

where $D_\phi = \text{diag}(e^{i\phi_j})$ is a dispersion phase and ξ is complex measurement noise. Multiplying by $W^{1/2}$ gives the weighted model

$$y = D_\phi \Phi x + \eta, \quad \eta = W^{1/2} \xi,$$

because the diagonal matrices D_ϕ and $W^{1/2}$ commute. After compensation with a phase model $\tilde{\phi}$, the residual phase is

$$E = D_{\tilde{\phi}}^* D_\phi.$$

Exact phase knowledge gives $E = I$.

2.2 Frame bounds, normalization and numerical rank

Let $s_{\min}(\Phi)$ and $s_{\max}(\Phi)$ be the smallest and largest singular values. We compute the frame bounds from the singular values,

$$A_\Omega = s_{\min}(\Phi)^2, \quad B_\Omega = s_{\max}(\Phi)^2,$$

and define

$$\kappa_\Omega = \frac{s_{\max}(\Phi)}{s_{\min}(\Phi)} = \sqrt{\frac{B_\Omega}{A_\Omega}}.$$

Stable unregularized reconstruction requires $A_\Omega > 0$. We call $M = N$ **critical sampling** and $M > N$ **oversampling**.

The numerical study uses half-neighbour Voronoi cell widths, with endpoint cells extended by half the nearest-neighbour spacing. The weights are rescaled to mean one, so $\sum_j \mu_j = M$. Because this normalization makes the absolute frame bounds scale with sample count, cross-budget comparisons also use A_Ω/M .

The least-squares estimate is formed with a truncated singular-value pseudoinverse using $\text{rcond} = 10^{-12}$. The effective numerical rank is the number of singular values exceeding $10^{-12}s_{\max}$. Frame bounds are nevertheless reported from the untruncated SVD. A separate machine-resolution flag uses the threshold $\max(M, N) \epsilon_{\text{mach}} s_{\max}$.

Reconstruction error is measured in the same coefficient space as the propositions below:

$$\text{NRMSE} = \frac{\|\hat{x} - x\|_2}{\|x\|_2}.$$

3 Stability statements

Sampling geometry and dispersion phase act on different sides of the forward model. A unit-modulus diagonal phase rotates measurement coordinates but does not alter the Gram operator, whereas compensation with an incorrect phase leaves a residual multiplicative error.

Proposition 1 (known-phase invariance). Let D_ϕ be diagonal with unit-modulus entries. Then

$$(D_\phi \Phi)^* (D_\phi \Phi) = \Phi^* \Phi.$$

Consequently, A_Ω , B_Ω and κ_Ω are unchanged by any known diagonal unit-modulus spectral phase.

Proof. Since $D_\phi^* D_\phi = I$,

$$(D_\phi \Phi)^* (D_\phi \Phi) = \Phi^* D_\phi^* D_\phi \Phi = \Phi^* \Phi. \quad \square$$

Proposition 2 (phase-model mismatch). Assume that Φ has full column rank and that $y = D_\phi \Phi x + \eta$ is compensated with D_ϕ^* . Let $E = D_\phi^* D_\phi$. The least-squares reconstruction

$$\hat{x} = \Phi^\dagger D_\phi^* y$$

satisfies

$$\|\hat{x} - x\|_2 \leq \kappa_\Omega \|E - I\|_2 \|x\|_2 + \frac{\|\eta\|_2}{\sqrt{A_\Omega}}.$$

Because E is diagonal and unitary,

$$\|E - I\|_2 = 2 \max_j \left| \sin \frac{\phi_j - \tilde{\phi}_j}{2} \right|.$$

Proof. We have

$$\hat{x} - x = \Phi^\dagger (E - I) \Phi x + \Phi^\dagger D_\phi^* \eta.$$

Using $\|\Phi^\dagger\|_2 = A_\Omega^{-1/2}$, $\|\Phi\|_2 = B_\Omega^{1/2}$ and $\|D_\phi^* \eta\|_2 = \|\eta\|_2$ gives the bound. $\square$

Corollary 3 (noise stability). If $y = \Phi x + \eta$ and Φ has full column rank, then

$$\|\hat{x} - x\|_2 \leq \frac{\|\eta\|_2}{\sqrt{A_\Omega}}.$$

This is the standard pseudoinverse estimate [9]. It motivates the use of A_Ω as a pre-reconstruction noise-amplification certificate.

Remark 1 (fixed-weight monotonicity). If measurements are added while all existing positive weights remain fixed, the unnormalized Gram matrix gains positive semidefinite rank-one terms, and its smallest eigenvalue cannot decrease [9]. This observation is not used to compare the cross-budget experiments below, because the Voronoi weights are recomputed and renormalized for every mask and every value of M .

Proposition 4 (Gershgorin certificate). Let $G = \Phi^* \Phi$, let $D = \text{diag}(G_{nn})$, and define $H = D^{-1/2} G D^{-1/2}$. If

$$\rho = \max_n \sum_{m \neq n} |H_{nm}| < 1,$$

then

$$A_\Omega \geq d_{\min}(1 - \rho), \quad d_{\min} = \min_n G_{nn}.$$

This standard Gershgorin estimate [9] is deliberately conservative. It certifies positivity when $\rho < 1$ but is inconclusive when $\rho \geq 1$; it is not used as a quantitative predictor for the ill-conditioned random masks.

4 Deterministic spectral mask design

4.1 Frame-certified gap filling

A retention percentage does not determine stability. At fixed M , different node sets can have different local clustering, boundary coverage, maximum gaps and Gram spectra. FCGF is designed to produce a candidate mask and then certify it; it does not claim global optimality.

Let $\Gamma_C = \{\tau_1, \dots, \tau_C\}$ be the candidate set. The algorithm proceeds as follows.


1. Form the **unweighted** $C \times N$ candidate Fourier matrix F_C .
2. Apply Businger-Golub column-pivoted QR to F_C^T and retain the first N pivoted columns, which correspond to N candidate measurement rows [4].
3. If $M > N$, partition $[\tau_{\min}, \tau_{\max}]$ by the selected nodes, including the two boundary intervals. Among intervals containing at least one unused candidate, select the widest; insert the unused candidate nearest its midpoint. Repeat until M nodes are selected. No wrap-around gap is used because the acquired spectrum occupies a finite interval rather than a circular domain.
4. Compute the Voronoi weights and certify the completed mask through A_Ω , A_Ω/M , κ_Ω and effective rank.

The pivoted QR stage is used as an algebraic independence heuristic, not as a claim of strong rank-revealing QR. Stronger RRQR and direct E-optimal selection are natural comparators for future work.

Lemma 5 (one-step gap reduction). Let Δ_1 and Δ_2 be the largest and second-largest fillable nonperiodic gaps of the current selected set. Let (a, b) be the largest gap, $m = (a + b)/2$, and let $c \in (\Gamma_C \setminus \Omega) \cap [a, b]$ be the candidate chosen by FCGF. Define $\varepsilon_C = |c - m|$. After inserting c ,

$$\Delta_{\text{new}} \leq \max \left\{ \Delta_2, \frac{\Delta_1}{2} + \varepsilon_C \right\}.$$

Proof. The inserted point splits (a, b) into two intervals of lengths $c - a$ and $b - c$. Since $|c - m| \leq \varepsilon_C$, each is at most $\Delta_1/2 + \varepsilon_C$. Every other fillable gap is at most Δ_2 . $\square$



Sampling locations at critical sampling ($M = N = 48$) for the three deterministic masks. The nonperiodic maximum gap includes the boundary intervals of the normalized spectral band

5 Controlled numerical benchmark

5.1 Setup

We used $C = 192$ detector samples uniformly spaced in wavelength from 800 to 900 nm. After conversion to wavenumber, the nodes were mapped by the formula in Section 2.1 to $[0, (C - 1)/C]$. The reconstruction space had $N = 48$ centered complex depth coefficients. The baseline reflectivity contained six localized complex Gaussian scatterers and a weak diffuse background.

Five thinning strategies were evaluated at $M = 48, 60, 72$ and 96: uniform detector-index thinning; nearest-candidate k -uniform sampling; FCGF; fully random sampling without replacement; and a coverage-controlled stratified-random baseline that partitions the ordered candidate grid into M strata and samples one candidate per stratum. Deterministic masks were tested over 30 independent noise realizations. For each random baseline, 80 independent geometries were generated at each M , and each geometry was tested with three noise realizations; its median NRMSE was used in geometry-level summaries.

Measurements included a known quadratic-plus-cubic dispersion phase and independent circular complex Gaussian noise. SNR was defined as $20\log_{10}(\text{signal RMS}/\text{noise RMS})$ and fixed at 35 dB; the real and imaginary noise components had equal variance. Reconstruction used Voronoi-weighted least squares with an SVD pseudoinverse threshold 10^{-12} . The software versions, random seeds, QR pivots, singular values and mask-level diagnostics are included in Online Resource 1.

5.2 Critical sampling

At $M = N = 48$, FCGF gave median NRMSE 0.0179, $A_\Omega/M = 0.6562$ and $\kappa_\Omega = 1.44$. Nearest k -uniform sampling gave median NRMSE 0.0203 and $\kappa_\Omega = 7.25$; detector-index thinning gave 0.0506 and 37.25, respectively. Fully random masks were frequently near singular, with median NRMSE 7104, whereas stratified-random sampling reduced the median to 0.0520. The contrast indicates that the most extreme random-mask failures arose from spectral coverage and conditioning rather than from randomness alone.

Two of the 80 fully random masks at $M = 48$ had effective rank below 48 under the 10^{-12} pseudoinverse threshold; one also had $s_{\min}$ below the machine-resolution criterion. All 80 stratified-random masks retained effective rank 48. The large random-mask errors therefore remain visible even after truncated-SVD stabilization.

5.3 Oversampling

As M increased, the deterministic masks converged in reconstruction error. At $M = 60$, FCGF was slightly less accurate than detector-index and k -uniform sampling. At $M = 72$, the three deterministic masks were nearly indistinguishable, and at $M = 96$ FCGF remained competitive but was not the lowest-error design. FCGF is therefore useful chiefly in the critical or low-redundancy regime; it does not dominate across all budgets.

5.4 Frame certificate and reconstruction failure

Across 320 fully random masks, $\log_{10}(A_\Omega/M)$ and $\log_{10}(\text{NRMSE})$ had pooled Spearman correlation $\rho = -0.994$ ($p < 10^{-300}$). The association remained strong within each budget, from $\rho = -0.994$ at $M = 48$ to $\rho = -0.830$ at $M = 96$ (Table 4). Excluding the single mask below the machine-resolution criterion left the pooled estimate essentially unchanged ($\rho = -0.994$, $n = 319$). The nonperiodic maximum gap also correlated positively with error, although less uniformly.

For stratified-random masks, the frame-bound association weakened as redundancy increased and was not detected at $M = 96$ ($\rho = 0.098$, $p = 0.387$). Once both coverage and redundancy are adequate, many masks occupy an already stable regime and the certificate becomes less discriminative.

5.5 Robustness across reflectivity profiles

We generated 50 additional complex reflectivity profiles with four to eight randomly located Gaussian scatterers and random amplitudes, widths and phases. Each deterministic mask at $M = N = 48$ was evaluated with three noise realizations per signal. The median per-signal NRMSE was 0.0178 for FCGF, 0.0214 for k -uniform sampling and 0.0585 for detector-index thinning (Table 3). The ranking of the medians therefore did not depend on the single baseline profile.


5.6 Known phase and phase-model mismatch

The mismatch study used the FCGF mask at $M = 72$ and 30 noise realizations per condition. A mismatch level q means that the quadratic phase coefficient 7.0 used in the compensation model was replaced by $7.0(1 + q)$; the cubic coefficient 1.5 was unchanged. The weighted Gram-matrix difference before and after applying the known phase was 5.85×10^{-14} in spectral norm, consistent with Proposition 1 to floating-point precision.


As q increased from 0 to 0.10, median NRMSE increased monotonically from 0.0151 to 0.1290 (Table 2). The phase-only term $\kappa_\Omega \|E - I\|_2$ and the complete right-hand side of Proposition 2 remained conservative. At 10% mismatch, the observed median NRMSE was 0.1290, the phase-only term was 0.2733, and the median full bound was 0.2949.

Table 1. Numerical stability and reconstruction summary. For deterministic masks, NRMSE is the median over 30 noise realizations and the geometry metrics are fixed. For random masks, each entry is a marginal median across 80 independently generated geometries; the values in one row need not come from the same geometry. The fully random baseline at $M = 48$ contained two geometries below the pseudoinverse effective-rank threshold; all other groups retained median effective rank 48.


M	Pattern	Median NRMSE	$A\Omega/M$	$\kappa\Omega$	Max gap
48	uniform-index	0.05058	0.001547	37.25	0.02816
48	k-uniform	0.02026	0.04453	7.25	0.02535
48	FCGF	0.01788	0.6562	1.44	0.02475
48	random	7104	4.389e-14	9.47e6	0.07298
48	stratified-random	0.05203	0.001341	39.11	0.03868
60	uniform-index	0.01642	0.6285	1.581	0.02319
60	k-uniform	0.01655	0.6647	1.614	0.02200
60	FCGF	0.01763	0.6170	1.498	0.02217
60	random	0.6352	8.940e-6	519.1	0.06180
60	stratified-random	0.01831	0.1288	3.861	0.02926
72	uniform-index	0.01504	0.6366	1.588	0.01754
72	k-uniform	0.01496	0.6709	1.458	0.01750
72	FCGF	0.01475	0.6778	1.382	0.02080
72	random	0.07516	0.0007894	53.68	0.05422
72	stratified-random	0.01635	0.2769	2.434	0.02704
96	uniform-index	0.01262	0.8717	1.203	0.01557
96	k-uniform	0.01268	0.8495	1.243	0.01510
96	FCGF	0.01285	0.5213	1.452	0.01453
96	random	0.01804	0.06198	5.303	0.03889
96	stratified-random	0.01305	0.6688	1.415	0.01745



Median NRMSE versus retained sample count. Random and stratified-random curves summarize 80 independent geometries per budget




Sample-count-normalized lower frame bound versus retained sample count. The normalization removes the trivial scaling induced by mean-one quadrature weights




Relationship between the sample-count-normalized lower frame bound and NRMSE for the two random baselines. The left-pointing marker denotes the single geometry below the machine-resolution criterion

Table 2. Dispersion-phase mismatch for the FCGF mask at $M = 72$ (30 noise realizations per condition). The mismatch is the relative perturbation of the quadratic phase coefficient; the cubic coefficient is fixed. The median noise-only term was 0.0217 for every condition.

Mismatch	Median NRMSE	$\ E-I\ _2$	Phase-only term	Full bound
0%	0.0151	0.0000	0.0000	0.0217
1%	0.0197	0.0198	0.0274	0.0491
2%	0.0295	0.0396	0.0547	0.0764
5%	0.0658	0.0990	0.1368	0.1585
10%	0.1290	0.1977	0.2733	0.2949



Observed reconstruction error and the phase-only and full bounds from Proposition 2 under deliberate mismatch of the quadratic dispersion coefficient



Representative FCGF reconstruction at critical sampling ($M = N = 48$) and 35 dB measurement SNR

Table 3. Multi-signal robustness at critical sampling ($M = N = 48$). Entries summarize per-signal median NRMSE across three noise realizations for 50 synthetic reflectivity profiles.

Pattern	Median	Mean	SD	Min	Max
uniform-index	0.0585	0.0586	0.0192	0.0240	0.0961
k-uniform	0.0214	0.0217	0.0025	0.0160	0.0295
FCGF	0.0178	0.0179	0.0010	0.0156	0.0198

Table 4. Within-budget rank correlations for random mask geometries. Each entry gives Spearman ρ followed by its two-sided p value in parentheses. The frame-bound analysis uses $\log_{10}(A_\Omega/M)$ and $\log_{10}(\text{NRMSE})$; the gap analysis uses the nonperiodic maximum gap and $\log_{10}(\text{NRMSE})$.

Baseline	M	Frame bound vs error	Max gap vs error
random	48	-0.994 (8.5e-77)	0.442 (4.0e-5)
random	60	-0.985 (9.1e-61)	0.675 (6.7e-12)
random	72	-0.968 (1.8e-48)	0.821 (1.2e-20)
random	96	-0.830 (1.9e-21)	0.746 (2.0e-15)
stratified-random	48	-0.667 (1.4e-11)	0.060 (0.595)
stratified-random	60	-0.393 (3.2e-4)	0.239 (0.0329)
stratified-random	72	-0.388 (3.8e-4)	0.177 (0.117)
stratified-random	96	0.098 (0.387)	-0.118 (0.297)

6 Discussion

The principal distinction is between reconstructing data already acquired at nonuniform Fourier nodes and certifying a proposed acquisition geometry before reconstruction. Existing NUFFT and NUDFT methods accelerate transform evaluation, compressed sensing introduces signal priors, and prior frame-based OCT work reconstructs from given nonuniform samples [5,8,10,12-17]. Here the lower frame bound is used prospectively: it asks whether a mask stably samples the full target space before a solver, regularizer or test signal is chosen.

The numerical results support two qualified conclusions. First, the certificate is most informative near the stability boundary. At $M = N = 48$, small geometric differences produced large changes in A_Ω/M , condition number and error. FCGF combined an algebraically independent CPQR core with gap control and performed best among the tested critical-sampling masks. Its advantage largely disappeared with oversampling, where several designs were already well conditioned. Second, catastrophic fully random failures were not evidence that randomness is intrinsically unsuitable. Stratification restored coverage and removed most failures. The comparison therefore favors geometry-aware design, not a blanket deterministic-versus-random conclusion.

The phase analysis isolates a separate issue. Exact knowledge of a unit-modulus dispersion phase cannot repair a poorly conditioned mask, because it does not change the frame bounds. Conversely, a well-conditioned mask cannot remove calibration error: a residual phase enters through the condition-number-weighted term in Proposition 2. This separation is useful whenever acquisition geometry and phase calibration are tuned by different parts of an imaging system.

FCGF should be read as a reproducible design heuristic. CPQR, Voronoi weighting, gap filling and frame certification are individually standard ingredients. The contribution is their acquisition-side combination for nonuniform FD-OCT and the explicit use of the lower frame bound to accept, reject or compare candidate masks before reconstruction. The controlled benchmark indicates when this combination is useful; it does not establish a universal optimal design.

7 Limitations and next experiments

The study uses a 48-dimensional finite model, a single 800-900 nm candidate band, synthetic complex reflectivity and circular Gaussian noise. Real OCT also contains source relative-intensity noise, detector response, apodization, speckle, depth-dependent sensitivity and calibration drift. The model assumes that a complex interferometric term has already been isolated and therefore does not address raw real-valued interferogram processing.

The mask is selected on a mathematical candidate grid rather than under a specific detector readout or swept-source clocking constraint. The CPQR core is not a strong-RRQR or E-optimal solution, and the gap rule is local. The next mathematical comparisons should include strong RRQR, density-based designs, direct

maximization of A_Ω and hardware-constrained combinatorial selection. The next optical test should replay raw experimental SD-OCT or SS-OCT spectra under prospective thinning and ask whether A_Ω predicts point-spread width, sidelobe energy, sensitivity falloff and layer-boundary localization. Only such experiments could support a biomedical-engineering claim.

8 Conclusion

Nonuniform FD-OCT reconstruction is established, but the geometry of a proposed spectral mask can also be examined before reconstruction. In the finite weighted Fourier model studied here, the lower frame bound provides a signal-independent stability certificate. A known unit-modulus dispersion phase leaves that certificate unchanged, while phase-model error obeys a condition-number-weighted bound. FCGF combines a CPQR core, nonperiodic gap filling and post-selection frame certification. It was most useful near critical sampling and became comparable with simpler designs under oversampling. The resulting claim is intentionally narrow: spectral thinning should be specified not only by sample percentage and reconstruction method, but also by the stability of the induced sampling operator on the target space.

Data and code availability

All results were generated from synthetic data. The complete simulation and figure-generation code, random seeds, QR pivots, software-version metadata, mask-level diagnostics, singular values, aggregate tables and representative reconstructions are supplied with the submission as **Online Resource 1**. No patient or human-subject data were used.

Funding

No external funding was received for this work.

Competing interests

The authors declare no competing interests relevant to this mathematical and computational study.

Author contributions

T.F. conceived the study, developed the mathematical framework, designed and executed the computational experiments, and drafted the manuscript. H.F. contributed to interpretation and critical revision. Both authors reviewed and approved the manuscript.

References

- [1] Adcock, B., Gataric, M., Hansen, A.C.: On stable reconstructions from nonuniform Fourier measurements. *SIAM J. Imaging Sci.* **7**(3), 1690-1723 (2014). <https://doi.org/10.1137/130943431>
- [2] Adcock, B., Gataric, M., Hansen, A.C.: Weighted frames of exponentials and stable recovery of multidimensional functions from nonuniform Fourier samples. *Appl. Comput. Harmon. Anal.* **42**(3), 508-535 (2017). <https://doi.org/10.1016/j.acha.2015.09.006>
- [3] Berger, P., Gröchenig, K., Matz, G.: Sampling and reconstruction in distinct subspaces using oblique projections. *J. Fourier Anal. Appl.* **25**, 1080-1112 (2019). <https://doi.org/10.1007/s00041-018-9620-8>
- [4] Businger, P., Golub, G.H.: Linear least squares solutions by Householder transformations. *Numer. Math.* **7**, 269-276 (1965). <https://doi.org/10.1007/BF01436084>
- [5] Chan, K.K.H., Tang, S.: High-speed spectral domain optical coherence tomography using non-uniform fast Fourier transform. *Biomed. Opt. Express* **1**(5), 1309-1319 (2010). <https://doi.org/10.1364/BOE.1.001309>
- [6] Choma, M.A., Sarunic, M.V., Yang, C., Izatt, J.A.: Sensitivity advantage of swept source and Fourier domain optical coherence tomography. *Opt. Express* **11**, 2183-2189 (2003). <https://doi.org/10.1364/OE.11.002183>
- [7] de Boer, J.F., Cense, B., Park, B.H., Pierce, M.C., Tearney, G.J., Bouma, B.E.: Improved signal-to-noise ratio in spectral-domain compared with time-domain optical coherence tomography. *Opt. Lett.* **28**, 2067-2069 (2003). <https://doi.org/10.1364/OL.28.002067>

- [8] Fessler, J.A., Sutton, B.P.: Nonuniform fast Fourier transforms using min-max interpolation. *IEEE Trans. Signal Process.* **51**(2), 560-574 (2003). <https://doi.org/10.1109/TSP.2002.807005>
- [9] Horn, R.A., Johnson, C.R.: *Matrix Analysis*, 2nd edn. Cambridge University Press, Cambridge (2012). <https://doi.org/10.1017/CBO9781139020411>
- [10] Ke, J., Lam, E.Y.: Image reconstruction from nonuniformly spaced samples in spectral-domain optical coherence tomography. *Biomed. Opt. Express* **3**(4), 741-752 (2012). <https://doi.org/10.1364/BOE.3.000741>
- [11] Leitgeb, R., Hitzenberger, C.K., Fercher, A.F.: Performance of Fourier domain vs. time domain optical coherence tomography. *Opt. Express* **11**, 889-894 (2003). <https://doi.org/10.1364/OE.11.000889>
- [12] Liu, X., Kang, J.U.: Compressive SD-OCT: the application of compressed sensing in spectral domain optical coherence tomography. *Opt. Express* **18**(21), 22010-22019 (2010). <https://doi.org/10.1364/OE.18.022010>
- [13] Liu, X., Kang, J.U.: Sparse OCT: optimizing compressed sensing in spectral domain optical coherence tomography. *Proc. SPIE* **7904**, 79041C (2011). <https://doi.org/10.1117/12.874058>
- [14] Nagib, K., Mezgebo, B., Fernando, N., Kordi, B., Sherif, S.S.: Generalized image reconstruction in optical coherence tomography using redundant and non-uniformly-spaced samples. *Sensors* **21**(21), 7057 (2021). <https://doi.org/10.3390/s21217057>
- [15] Wang, K., Ding, Z., Wu, T., Wang, C., Meng, J., Chen, M., Xu, L.: Development of a non-uniform discrete Fourier transform based high speed spectral domain optical coherence tomography system. *Opt. Express* **17**(14), 12121-12131 (2009). <https://doi.org/10.1364/OE.17.012121>
- [16] Xu, D., Huang, Y., Kang, J.U.: Compressive sensing with dispersion compensation on non-linear wavenumber sampled spectral domain optical coherence tomography. *Biomed. Opt. Express* **4**(9), 1519-1532 (2013). <https://doi.org/10.1364/BOE.4.001519>
- [17] Xu, D., Huang, Y., Kang, J.U.: GPU-accelerated non-uniform fast Fourier transform-based compressive sensing spectral domain optical coherence tomography. *Opt. Express* **22**(12), 14871-14884 (2014). <https://doi.org/10.1364/OE.22.014871>