

Fabrication-Tolerant Inverse Design of Chirped Apodized Fiber Bragg Gratings: A Transfer-Matrix and Uncertainty-Aware Computational Study

Connor Noble

c0613727@gmail.com

Research Article

Keywords: fiber Bragg grating, chirped grating, apodization, inverse design, transfer matrix, robust optimization

Posted Date: September 1st, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-10788859/v1>

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

[Read Full License](#)

Additional Declarations: No competing interests reported.

Fabrication-Tolerant Inverse Design of Chirped Apodized Fiber Bragg Gratings: A Transfer-Matrix and Uncertainty-Aware Computational Study

Connor Noble

Independent Interdisciplinary Researcher, United States

Corresponding author: Connor Noble | E-mail: c0613727@gmail.com

Abstract

Chirped fiber Bragg gratings (CFBGs) are compact wavelength-selective devices for dispersion control, filtering, and sensing, yet designs optimized only at nominal parameters can be fragile to inscription and calibration errors. This computational study develops a five-parameter inverse-design framework for a linearly chirped, continuously apodized CFBG and evaluates it with coupled-mode theory and a segmented transfer-matrix solver. The coupling envelope is parameterized by a generalized sinusoidal-cosine window, while grating length, index-modulation amplitude, chirp span, apodization exponent, and shape coefficient are optimized jointly. A nominal objective is compared with an uncertainty-aware objective evaluated over deterministic fabrication corners involving Bragg-wavelength shift, coupling-strength error, and chirp-scale error. High-resolution simulations show that the nominal and robust solutions achieve mean reflectivities of 0.9997 and 0.9991, respectively, over 1548–1552 nm, compared with 0.7275 for the conventional baseline. Under the prescribed corner set, the robust design raises the worst minimum in-band reflectivity from 0.9661 to 0.9820 and reduces worst ripple from 0.1476 dB to 0.0768 dB, at the cost of modestly higher stop-band leakage. A 180-sample Monte Carlo study and discretization convergence analysis further quantify statistical performance and numerical error. The results demonstrate that uncertainty-aware inverse design can improve guaranteed CFBG spectral performance without requiring experimental training data.

Keywords: fiber Bragg grating; chirped grating; apodization; inverse design; transfer matrix; robust optimization

1 Introduction

Fiber Bragg gratings (FBGs) are distributed dielectric reflectors formed by a longitudinal modulation of the effective refractive index in an optical-fiber core. Since the foundational development of photosensitive gratings, they have become important in wavelength selection, dispersion engineering, sensing, filtering, laser cavities, and multiplexed optical networks (Hill and Meltz 1997; Othonos 1997; Erdogan 1997; Kashyap 2010). The basic device is deceptively simple: a submicrometre-scale periodic index perturbation coherently couples a forward-propagating core mode to its backward counterpart when the Bragg phase-matching condition is met. In practice, useful spectral functions are produced by deliberately varying the grating amplitude, period, or phase along the device. This makes the FBG a natural one-dimensional inverse-design problem in which a physically fabricable longitudinal profile must be selected to realize a target complex reflection response.

Chirped FBGs are especially useful when different wavelengths should interact at different axial positions. A longitudinally varying Bragg wavelength introduces a wavelength-dependent reflection depth and therefore a controllable group delay. For optical communication links this mechanism can compensate chromatic dispersion; for pulse shaping it can compress or stretch waveforms; and for interrogation or sensing systems it can create prescribed spectral windows. Apodization is complementary to chirp. By smoothly tapering the coupling strength near the grating ends, apodization suppresses sidelobes generated by abrupt truncation, although aggressive tapering may reduce usable reflectivity or increase device length. The design problem is consequently multi-objective even before fabrication variability is considered.

Recent work in Optical and Quantum Electronics has continued to investigate cascaded, chirped, and apodized gratings for dispersion compensation. Mustafa et al. (2022) examined a new apodization function in a WDM system, Shaalan et al. (2025) combined dispersion-compensation structures with supervised regression, and Abdallah et al. (2026) studied cascaded gratings with a hybrid Nuttall-Bartlett profile. These studies demonstrate continuing interest in spectral shaping and optical-link performance. Much of this literature, however, compares a finite menu of named apodization windows or tunes system parameters for a single nominal model. Such workflows answer the question 'which candidate is best at the assumed design point?' but do not necessarily answer 'which design retains acceptable performance when the fabricated device differs from the nominal prescription?'. The latter question is central to manufacturable photonics.

FBG uncertainty arises from several mechanisms. The written period and the effective index determine the local Bragg wavelength, so small thermal, strain, calibration, or phase-mask errors shift the spectrum. The index-modulation amplitude depends on photosensitivity, exposure dose, writing-beam stability, and material response; these effects perturb the coupling coefficient. Chirp may also deviate systematically from its intended scale. In a high-reflectivity design, these perturbations can interact nonlinearly: increasing coupling may broaden the reflection band and improve its interior while also increasing leakage near a stop-band boundary; a chirp error may move one pass-band edge into a transition region even when the center remains well matched. Robustness therefore cannot be inferred from a single derivative or from nominal reflectivity alone.

The present paper formulates a compact uncertainty-aware inverse-design problem for a single chirped, apodized FBG. Rather than selecting among discrete named windows, the coupling envelope is described by two continuous shape parameters embedded in a five-dimensional design vector. The optical response is evaluated with standard coupled-mode theory and a transfer-matrix discretization. Two optimization strategies are then compared: a nominal objective evaluated at the intended geometry, and a robust objective that augments average performance with worst-corner penalties over perturbations in Bragg wavelength, coupling strength, and chirp scale. This construction is deliberately transparent. It uses no experimental training data and no opaque surrogate model; every spectrum is generated from the governing equations.

Four questions guide the study. First, how much improvement over a conventional smoothly apodized chirped grating can be obtained by joint continuous optimization? Second, does an uncertainty-aware objective improve deterministic worst-case pass-band quality compared with a nominal optimum? Third, how does this protection affect average Monte Carlo performance and stop-band leakage? Fourth, are the reported differences larger than the numerical error introduced by segmentation of the grating? The central contribution is not a claim that one window is universally optimal, but a reproducible framework for quantifying the trade-off between nominal efficiency and guaranteed spectral quality in a low-dimensional, fabrication-aware FBG design space.

The remainder of the paper is organized as follows. Section 2 summarizes the modeling context and motivates the selected parameterization. Section 3 develops the coupled-mode and transfer-matrix formulation. Section 4 defines the nominal and robust inverse-design objectives and uncertainty model. Section 5 describes the computational implementation and verification. Section 6 presents optimized designs, spectra, group-delay behavior, deterministic corner results, Monte Carlo statistics, and convergence tests. Sections 7 and 8 discuss engineering implications and limitations, and Sect. 9 concludes.

2 Modeling context and design motivation

The spectral behavior of fiber gratings is commonly analyzed with coupled-mode theory. Kogelnik's coupled-wave framework for volume gratings established the mathematical structure of resonant mode conversion, and later FBG-specific treatments provided closed-form and numerical descriptions for uniform, apodized, chirped, phase-shifted, and superstructured gratings (Kogelnik 1969; Erdogan 1997). For a uniform lossless grating, analytical expressions show that peak reflectivity increases with the product of coupling coefficient and length. Once the coupling or detuning varies with position, however, a numerical propagation method is generally more convenient. Transfer matrices provide such a method and remain widely used for layered and one-dimensional optical structures (Morozov 2023).

Apodization is conventionally introduced by multiplying the refractive-index modulation by a smooth envelope. Gaussian, raised-cosine, sine, Hamming, Cauchy, hyperbolic-tangent, Blackman, Bartlett, and related windows have all been explored. Naguib et al. (2021) showed that the choice of apodization profile changes both reflectivity and dispersion-related characteristics, while Sayed et al. (2020) used an apodized CFBG in a post-dispersion compensation configuration. More recent OQE work has extended the same design vocabulary to cascaded or hybrid profiles (Mustafa et al. 2022; Abdallah et al. 2026). These results motivate continuous profile optimization because a fixed library of windows samples only a small subset of possible coupling envelopes.

The design variable used here is intentionally simple enough to remain interpretable. Let $x=z/L$ be normalized position. The apodization envelope is $g(x)=\sin(\pi x)^p[1+a \cos(2\pi(x-1/2))]$, clipped at zero if required. The exponent p controls edge tapering: $p=2$ resembles a raised-sine-squared profile, whereas $p<1$ makes the envelope flatter over a larger fraction of the device. The coefficient a redistributes coupling between the center and shoulders without introducing high-frequency axial structure. This two-parameter family does not attempt to approximate every realizable grating. Its purpose is to provide a continuous design space with an obvious fabrication interpretation and a direct connection to common window functions.

The second modeling choice is to optimize spectral quantities at the device level rather than embed the grating in a complete transmitter-channel-receiver simulation. Link-level BER and Q-factor can be valuable application metrics, but they depend on transmitter power, receiver noise, amplifier placement, modulation format, fiber length, and other assumptions. By optimizing reflectivity, stop-band leakage, and group-delay structure directly, the present study isolates what is intrinsic to the grating. This makes comparisons easier to reproduce and permits the resulting design to be reused in multiple system contexts.

The third choice is to distinguish uncertainty protection from expected-value optimization. A design that performs extremely well for most random perturbations may still fail at a physically plausible corner of the tolerance box. Conversely, minimizing only the worst case can yield an unnecessarily conservative design if the corner is extremely unlikely. The robust objective below therefore combines average penalties over a compact set of deterministic corners with explicit penalties on the worst mean and worst minimum pass-band reflectivity. A separate Monte Carlo analysis is then used as an out-of-objective validation rather than as the optimization dataset itself.

3 Mathematical model

3.1 Refractive-index and Bragg-wavelength parameterization

A single-mode, lossless grating of length L is considered. The unperturbed effective index is $n_{\text{eff}}=1.447$, representative of a silica-fiber core near 1550 nm. The longitudinal index perturbation is written as a slowly varying amplitude multiplying a rapidly varying grating carrier. Under the standard slowly varying envelope approximation, only the local modulation amplitude and local Bragg period enter the first-order coupled-mode equations.

The local Bragg wavelength is taken to vary linearly along the grating, centered at $\lambda_{B0}=1550$ nm. If $\Delta\lambda_c$ denotes the total chirp span, then $\lambda_B(z)=\lambda_{B0} + \Delta\lambda_c(z/L-1/2)$. The corresponding local period follows from $\lambda_B(z)=2 n_{\text{eff}} \Lambda(z)$. This parameterization defines the intended grating in terms of an easily interpretable total spectral span rather than a period slope in units of length per length.

The index-modulation envelope is $\Delta n(z)=\Delta n_{\text{max}} g(z/L)$, where g is the generalized apodization family described above. The design vector is $\theta=(L, \Delta n_{\text{max}}, \Delta\lambda_c, p, a)$. Within the optimization bounds used here, g remains smooth and nonnegative after clipping. This avoids discontinuous phase steps and keeps the number of free parameters small enough for global optimization without a surrogate model.

$$\lambda_B(z) = \lambda_0 + \Delta\lambda_c (z/L - 1/2) \quad (1)$$

$$\Lambda(z) = \lambda_B(z) / (2 n_{\text{eff}}) \quad (2)$$

$$g(x) = \sin(\pi x)^p [1 + a \cos(2\pi(x - 1/2))], \quad x=z/L \quad (3)$$

$$\Delta n(z) = \Delta n_{\max} g(z/L) \quad (4)$$

3.2 Coupled-mode equations

Let $A(z, \lambda)$ and $B(z, \lambda)$ denote the slowly varying forward and backward modal amplitudes. For a lossless first-order Bragg grating with real coupling coefficient, the coupled-mode system can be written in matrix form. The detuning δ measures the mismatch between the modal propagation constant β and the local Bragg spatial frequency π/Λ . The coupling coefficient κ is proportional to the local index modulation. A unit fringe visibility is assumed; polarization dependence, cladding-mode coupling, and material loss are neglected in the base model.

The formulation is energy conserving for real δ and κ . It is also adequate for the present objective because the study concerns relative effects of chirp, apodization, and bounded parameter perturbations rather than radiation loss or vector-mode conversion. The index-modulation range remains within the order of magnitude commonly considered in FBG modeling, and the optimized length stays in the several-centimeter range used in dispersion-compensating gratings.

$$d/dz [A; B] = i [[\delta(z, \lambda), \kappa(z, \lambda)], [-\kappa(z, \lambda), -\delta(z, \lambda)]] [A; B] \quad (5)$$

$$\beta(\lambda) = 2\pi n_{\text{eff}} / \lambda, \quad \delta(z, \lambda) = \beta(\lambda) - \pi/\Lambda(z) \quad (6)$$

$$\kappa(z, \lambda) = \pi \Delta n(z) / \lambda \quad (7)$$

3.3 Segment transfer matrix

The grating is partitioned into N segments of equal length $h=L/N$. Within segment j , δ and κ are evaluated at the segment midpoint and treated as constant. Define $\gamma_j = (\kappa_j^2 - \delta_j^2)^{1/2}$, with the complex square root used automatically when $|\delta_j| > |\kappa_j|$. Exact exponentiation of the constant-coefficient 2×2 coupled-mode matrix gives a segment propagator. This eliminates the local truncation error associated with an explicit ODE step; remaining error is dominated by representing the continuously varying grating with piecewise-constant coefficients.

If T_j maps the mode amplitudes across segment j , the total transfer matrix is the ordered product $T_N \dots T_2 T_1$. The boundary condition for a reflection calculation sets the forward incident amplitude at $z=0$ to one and the backward incident amplitude at $z=L$ to zero. If the total matrix maps $[A(0), B(0)]^T$ to $[A(L), B(L)]^T$, the complex reflection coefficient follows directly from the second row. This approach is numerically compact and naturally supplies the reflection phase needed for group-delay estimation.

$$\gamma_j = (\kappa_j^2 - \delta_j^2)^{1/2} \quad (8)$$

$$T_j = \begin{bmatrix} \cosh(\gamma_j h) + i\delta_j \sinh(\gamma_j h)/\gamma_j & i\kappa_j \sinh(\gamma_j h)/\gamma_j \\ -i\kappa_j \sinh(\gamma_j h)/\gamma_j & \cosh(\gamma_j h) - i\delta_j \sinh(\gamma_j h)/\gamma_j \end{bmatrix} \quad (9)$$

$$T = T_N T_{N-1} \dots T_1 \quad (10)$$

$$r(\lambda) = -T_{21}(\lambda)/T_{22}(\lambda), \quad R(\lambda) = |r(\lambda)|^2 \quad (11)$$

3.4 Group delay and spectral metrics

The reflection phase $\phi(\lambda) = \arg r(\lambda)$ is unwrapped over the wavelength grid. With angular frequency $\omega = 2\pi c/\lambda$, the reflection group delay is $\tau_g = -d\phi/d\omega$. Because an arbitrary reference delay does not affect dispersion compensation, figures display group delay after subtracting its pass-band mean, while optimization uses a percentile-based ripple measure. A 5th-to-95th percentile span is preferred to a raw max-minus-min value because numerical differentiation is sensitive at isolated phase-transition points near weak-reflection regions.

The target pass band is 1548-1552 nm. Stop-band penalties are evaluated in 1544-1546.4 nm and 1553.6-1556 nm for optimization, leaving transition intervals on either side of the pass band. Metrics include mean in-band

reflectivity, minimum in-band reflectivity, pass-band ripple in decibels, maximum stop-band reflectivity, mean stop-band reflectivity, and the 5th-to-95th percentile group-delay span. No single metric is sufficient: an optimizer can trivially reduce stop-band leakage by narrowing the grating response, but this can destroy edge reflectivity; alternatively, very strong coupling can flatten the pass band but enlarge transition leakage.

$$\tau_g(\lambda) = -d\phi/d\omega, \quad \omega = 2\pi c/\lambda \quad (12)$$

$$\rho_R = \max_{\{\lambda \in P\}} 10 \log_{10} R(\lambda) - \min_{\{\lambda \in P\}} 10 \log_{10} R(\lambda) \quad (13)$$

3.5 Conservation structure and model verification

The lossless coupled-mode matrix has a structure that provides useful internal checks on the computation. Its trace is zero, and the exact segment propagator therefore has determinant one. The determinant of the product of all segments is consequently also one up to floating-point roundoff. This property is independent of whether γ is real or imaginary and provides a direct check that matrix multiplication has not introduced an unphysical gain or loss term. In a lossless grating with the adopted boundary conditions, reflected and transmitted powers should also satisfy power conservation when both amplitudes are normalized consistently.

The determinant property can be seen directly from the segment eigenvalues. The coefficient matrix iH has eigenvalues $\pm\gamma$, so the eigenvalues of $\exp(iHh)$ are $\exp(+\gamma h)$ and $\exp(-\gamma h)$; their product is one. In the propagating-detuned regime γ is imaginary and the pair is phase conjugate. In the stop-band-like regime γ is real and the pair is reciprocal. Large intermediate matrix elements can occur in the latter case, but the final physical reflection remains bounded. This is one reason to use double precision and to avoid extremely long or extremely strong gratings without a stabilized scattering-matrix formulation.

The present parameter range is moderate enough that ordinary 2×2 transfer-matrix multiplication remains stable. Numerical checks performed during model development verified finite matrix elements over the full wavelength grid and reflectivities no larger than unity to ordinary floating-point tolerance. The mesh-convergence results reported later provide an independent check on the piecewise-constant approximation. Together, determinant structure, bounded reflectivity, and convergence against an 1800-segment reference reduce the chance that the optimized spectra are numerical artifacts.


The model also has clear asymptotic limits. If the modulation amplitude approaches zero, κ tends to zero and the reflection coefficient vanishes. If chirp approaches zero while the envelope remains fixed, the device approaches an apodized uniform Bragg grating whose response is concentrated around a single Bragg wavelength. Increasing chirp maps a wider wavelength interval onto different axial reflection positions. These limits were used as qualitative sanity checks before optimization: zero-coupling test cases produced negligible reflection, while strong uniform-grating cases reproduced the expected high-reflectivity resonance near the Bragg wavelength.

A second useful check concerns symmetry. The chosen apodization family is symmetric about $z=L/2$, while the Bragg wavelength is linearly chirped about the same center. Reversing the grating direction changes the sign of the chirp but does not change the set of local Bragg wavelengths. For a reciprocal lossless scalar model, the reflection-power spectrum should therefore be closely related under the corresponding wavelength mapping even though reflection phase and reference delay can differ. This reciprocity perspective is consistent with the smooth spectral envelopes obtained here and helps distinguish amplitude design from absolute phase referencing.

$$\det(T_j)=1 \Rightarrow \det(T)=\prod_j \det(T_j)=1 \quad (14)$$

Table 1 Computational model and fixed assumptions

Quantity	Value / definition	Rationale
Central wavelength	1550 nm	Telecommunications C-band reference
Effective index n_{eff}	1.447	Representative silica-fiber modal index
Target pass band	1548–1552 nm	4 nm flat-reflection target
Optimization spectral grid	121 points over 1544–1556 nm	Balances global-search cost and spectral resolution
High-resolution verification	1201 points; 800 segments	Used for reported nominal spectra
Loss and polarization	Neglected; scalar single-mode model	Isolates chirp/apodization physics
Fringe visibility	1	Absorbed into effective modulation amplitude
Phase derivative	Unwrapped reflection phase vs. angular frequency	Defines group delay

*Fig. 1 Normalized apodization envelopes for the baseline, nominally optimized, and uncertainty-aware robust designs.*

4 Inverse-design formulation

4.1 Design variables and bounds

The optimization bounds were chosen to keep the search within a physically plausible, compact design region while allowing substantial spectral reshaping. Grating length varies from 25 to 58 mm; peak index modulation from 1.0×10^{-4} to 5.8×10^{-4} ; total chirp span from 4.8 to 8.0 nm; apodization exponent p from 0.45 to 3.5; and shape coefficient a from -0.45 to 0.45. The upper bounds are not asserted to be universal fabrication limits. They define the design box investigated in this computational study. The fact that both optimized solutions approach the length and modulation upper bounds is therefore discussed explicitly as a design-space limitation rather than hidden as an unconstrained physical optimum.

A conventional baseline is included to provide a reference independent of the optimizer. It uses $L=45$ mm, $\Delta n_{\text{max}}=4.0 \times 10^{-4}$, chirp span 5.2 nm, $p=2$, and $a=0$. This corresponds to a smoothly tapered $\sin^2$ envelope and

provides excellent remote stop-band suppression, but the high-resolution calculations below show that it does not sustain a flat 4 nm reflection plateau.

Table 2 Design variables and search bounds

Variable	Meaning	Lower	Upper
L (mm)	Grating length	25	58
$\Delta n_{\text{max}} (\times 10^{-4})$	Peak effective-index modulation	1.0	5.8
$\Delta \lambda_{\text{c}}$ (nm)	Total linear chirp span	4.8	8.0
p	Sinusoidal apodization exponent	0.45	3.5
a	Cosine shoulder-shaping coefficient	-0.45	0.45

4.2 Nominal objective

The nominal loss is a weighted sum of physically interpretable penalties. The dominant terms reward high mean and high minimum pass-band reflectivity. A separate stop-band term penalizes the maximum leakage outside the transition zones. Ripple is penalized only above a small tolerance, preventing the optimizer from sacrificing large gains in edge reflectivity for numerically insignificant flattening. A mild group-delay penalty activates only above a threshold. The weights are engineering choices and are reported to make the optimization reproducible rather than presented as uniquely optimal.

The use of both mean and minimum reflectivity is important. Mean reflectivity alone permits narrow edge failures, whereas a minimum-only objective can overemphasize one grid point. Combining them encourages a broad plateau. Similarly, maximum stop-band reflectivity suppresses isolated sidelobes that could be missed by an average leakage measure.

$$J_{\text{nom}} = 8(1 - \bar{R}_P) + 3(1 - R_{\text{min},P}) + 5R_{\text{max},S} + 0.12[\rho_R - 0.2]_+ + 0.0008[\rho_\tau - 100]_+ \quad (15)$$

4.3 Uncertainty-aware objective

Fabrication uncertainty is represented through three multiplicative or additive perturbations: an overall Bragg-wavelength shift ϵ_{λ} , a coupling-strength scale ϵ_{κ} , and a chirp-scale factor ϵ_c . The seven deterministic scenarios include the nominal point, two aligned extreme corners, two crossed corners, and two moderate mixed perturbations. The extreme wavelength shift is ± 0.25 nm, coupling-scale error reaches $\pm 10\%$, and chirp-scale error reaches approximately $\pm 8\%$. These values are not a statistical claim about a particular inscription platform; they define a stress-test set for comparing design philosophies.

For each scenario, the same spectral metrics are calculated. The robust objective averages the nominal-style loss over all scenarios and adds explicit worst-case penalties for the smallest scenario-average reflectivity, the largest stop-band leakage, and the smallest pointwise in-band reflectivity. The design is thus discouraged from improving the nominal spectrum at the expense of a severe edge collapse under one corner. Importantly, the Monte Carlo distribution used later is not part of this objective. It is an independent test with Gaussian-like perturbations truncated inside a wider practical range.

$$\theta_{\text{rob}} = \arg \min_{\theta} \{ \text{mean}_q J(\theta, q) + 4[1 - \min_q \bar{R}_P(q)] + 2 \max_q R_{\text{max},S}(q) + 1.5[1 - \min_q R_{\text{min},P}(q)] \} \quad (16)$$

Table 3 Deterministic uncertainty scenarios used by the robust objective

Scenario	Bragg shift (nm)	Coupling scale	Chirp scale
Nominal	0.00	1.00	1.00
Corner A	-0.25	0.90	0.92
Corner B	+0.25	1.10	1.08

Scenario	Bragg shift (nm)	Coupling scale	Chirp scale
Crossed C	-0.25	1.10	1.08
Crossed D	+0.25	0.90	0.92
Mixed E	-0.15	0.95	1.06
Mixed F	+0.15	1.05	0.94

4.4 Optimization algorithm

The five-dimensional box-constrained problem is solved with differential evolution followed by Nelder-Mead refinement. Differential evolution is useful here because the objective contains max, percentile, and threshold operators that make analytic gradients inconvenient and because the transfer-matrix response can be highly nonconvex near band edges. Fixed pseudorandom seeds are used for repeatability. The global stage uses a population size of seven times the number of variables and 18 generations in the final reported run; the local stage then refines the best candidate until parameter and objective changes become small or the iteration budget is reached.

During optimization the grating is represented by 130 segments and the spectrum by 121 wavelength samples. This deliberately inexpensive model is used only to navigate the design space. Final metrics are recomputed with substantially finer resolution. The convergence study in Sect. 6.6 verifies that the final trends are not consequences of the coarse optimization mesh.

4.5 Why corner robustness is separated from probability modeling

Robust photonic design is sometimes described as though there were a single mathematically preferred treatment of uncertainty. In reality, the correct risk model depends on what is known about fabrication. If only hard tolerances are available, a minimax or corner-based formulation is natural. If a statistically calibrated process distribution is available, expected loss, chance constraints, or conditional value at risk may be more informative. If even the distribution is uncertain, distributionally robust optimization can protect against a family of plausible probability laws. The present study intentionally uses a deterministic scenario set for optimization and a separate probability model for evaluation so that these concepts are not conflated.

The corner set is small enough to keep every objective evaluation transparent. Each scenario has a direct physical interpretation: a spectral shift, a global coupling error, and a chirp-scale error. Mixed and crossed scenarios prevent the optimizer from assuming that all favorable changes occur together. Because the objective includes both average scenario loss and explicit extremes, a candidate cannot hide one severe edge failure behind six good scenarios. This is particularly important for flat-top filters because the mean reflectivity may remain close to unity even when a narrow edge region drops appreciably.

The Monte Carlo stage asks a different question: if perturbations are more densely concentrated near the nominal point, what distribution of performance is observed? Using common random numbers for all designs reduces comparison noise because each design experiences exactly the same sampled perturbations. The fixed seed makes the experiment deterministic for reproducibility even though the samples emulate a stochastic manufacturing population. Quantiles are emphasized alongside means because yield-sensitive photonic components are often limited by lower-tail behavior.

No claim is made that the chosen uncertainty ranges correspond to a particular commercial fabrication line. A future experimental study should estimate these parameters from repeated inscriptions, including correlations between index modulation and chirp caused by writing speed, pulse energy, mask positioning, or thermal history. The robust framework can then be recalibrated without changing the optical solver. This separation between device physics and process statistics is a practical advantage of the formulation.




Fig. 2 Differential-evolution objective histories for the nominal and robust searches. Local Nelder-Mead refinement follows the displayed global stage.

5 Computational implementation and verification protocol

All calculations were performed with a reproducible Python workflow using double-precision NumPy arrays and SciPy optimization routines. For each wavelength, the segment matrices were multiplied in propagation order. The implementation vectorizes the wavelength dimension while looping over grating segments, which reduces the cost of objective evaluation without altering the mathematical model. The complex square root in γ is evaluated on the principal branch; for $|\gamma|$ close to zero, the ratio $\sinh(\gamma h)/\gamma$ is replaced by its limiting value h to avoid a removable numerical singularity.

The high-resolution nominal spectra use 1201 wavelength points over 1544-1556 nm and 800 grating segments. The deterministic corner study uses 161 wavelengths and 360 segments, which is sufficient for the comparatively smooth worst-case metrics. The Monte Carlo validation uses 180 samples. Bragg shift is drawn from a zero-mean normal distribution with standard deviation 0.10 nm and truncated to ± 0.30 nm. Coupling scale has standard deviation 0.045 and is truncated to [0.88, 1.12]. Chirp scale has standard deviation 0.035 and is truncated to [0.90, 1.10]. A fixed seed (20260821) makes the sample set reproducible.

The validation logic is intentionally separated into four layers. First, nominal high-resolution spectra test whether optimization actually creates the intended plateau. Second, deterministic corners probe specified edge cases and supply guaranteed performance over the chosen finite scenario set. Third, Monte Carlo sampling tests average and quantile behavior under a different perturbation distribution. Fourth, mesh convergence compares robust-design reflectivity at five representative wavelengths against an 1800-segment reference. Agreement across these layers is stronger evidence than a single optimized spectrum.

6 Results

6.1 Optimized design parameters

Table 4 lists the baseline and optimized parameter vectors. Both optimized solutions move toward longer, stronger gratings and substantially flatter apodization envelopes than the $p=2$ baseline. The nominal solution uses $p \approx 0.450$ and $a \approx -0.175$, while the robust solution drives the shoulder coefficient farther negative, $a \approx -0.292$. Because the cosine term is centered symmetrically, the more negative coefficient redistributes coupling away from the center relative to the nominal optimum. This apparently small profile change is the main structural difference between the two optimized solutions.

The tendency toward the upper length and modulation bounds should be interpreted carefully. Within a passive lossless coupled-mode model, stronger integrated coupling is advantageous for maintaining high reflection across a chirped band. Fabrication loss, UV-dose limits, nonlinear photosensitivity, and device-size cost are not included. The optimizer therefore reveals that these omitted costs would be necessary to produce an interior optimum in L and $\Delta n_{\max}$. The meaningful comparison here is between profile shapes and robustness at a common constrained design scale, not a claim that still longer or stronger gratings would never be preferable.

Table 4 Baseline and optimized design parameters

Design	L (mm)	$\Delta n_{\max}$ ($\times 10^{-4}$)	Chirp span (nm)	p	a
Baseline	45.00	4.000	5.200	2.000	0.000
Nominal optimized	57.80	5.781	5.563	0.450	-0.175
Robust optimized	57.34	5.799	5.563	0.450	-0.292

Figure 1 shows the normalized envelopes. The baseline has strong edge suppression and a narrow central maximum. Both optimized profiles remain smooth at the ends but are much flatter through the interior, allowing a larger fraction of the grating length to contribute significant coupling. The robust solution slightly depresses the central region and strengthens the shoulders relative to the nominal solution. This redistributes spectral coupling in a way that improves the weakest pass-band edge under the prescribed uncertainty corners.

6.2 Nominal reflection spectra

Figure 3 compares the high-resolution nominal power reflectivity. The conventional $p=2$ baseline is spectrally clean outside the reflection band but does not provide a flat 4 nm plateau: its mean in-band reflectivity is only 0.7275 and its minimum falls to 0.0730 near the target-band edge. The nominal optimization changes the character of the device. Mean in-band reflectivity rises to 0.999705 and minimum in-band reflectivity to 0.998874. The robust design remains close, with mean 0.999123 and minimum 0.997430. Thus robustness is purchased with a very small nominal penalty.

The dB-scale view in Fig. 4 exposes the associated stop-band trade-off. The baseline's aggressive taper produces extremely low remote reflection in the selected stop windows. The nominal and robust optimizations broaden the useful reflection plateau, but their maximum stop-band reflectivities rise to approximately 5.75×10^{-4} and 7.06×10^{-4} , respectively. In decibels these remain below approximately -32 dB in the defined remote windows, but they are orders of magnitude larger than the baseline. The design problem therefore cannot be summarized by pass-band reflectivity alone: the optimized devices exchange some remote spectral cleanliness for a much more uniform wide reflection band.

The pass-band ripple metric reinforces this conclusion. The baseline ripple is 11.33 dB, indicating that the nominal target band extends far beyond its flat high-reflection region. The nominal optimum reduces ripple to

0.00418 dB; the robust optimum gives 0.01041 dB. These are both effectively flat on the scale of many device applications, though the exact acceptability threshold depends on system requirements.




Fig. 3 High-resolution nominal reflection spectra. The shaded region marks the 1548–1552 nm optimization pass band.




Fig. 4 Reflection spectra on a logarithmic scale, emphasizing the pass-band/stop-band trade-off.

Table 5 High-resolution nominal spectral metrics

Design	Mean R (pass)	Min R (pass)	Ripple (dB)	Max R (stop)	GD ripple (ps)
Baseline	0.727481	0.072974	11.3264	3.315e-10	363.3

Design	Mean R (pass)	Min R (pass)	Ripple (dB)	Max R (stop)	GD ripple (ps)
Nominal optimized	0.999705	0.998874	0.0042	5.753e-04	354.7
Robust optimized	0.999123	0.997430	0.0104	7.059e-04	350.2

6.3 Group-delay behavior

Figure 5 plots group delay after subtracting the pass-band mean. All three gratings exhibit wavelength-dependent delay, as expected for linearly chirped structures. The optimized devices produce a high-reflectivity plateau without eliminating delay structure. Using the high-resolution 5th-to-95th percentile measure, the baseline, nominal, and robust designs yield spans of 363.3, 354.7, and 350.2 ps, respectively. The robust design therefore provides a modest improvement in this metric, but group-delay ripple remains appreciable compared with the pass-band amplitude ripple.

This result is important for interpretation. The objective gives group delay a relatively low weight and activates the penalty only above a threshold. The optimizer accordingly prioritizes reflectivity and band coverage. A dispersion-compensator design intended for a specific fiber link should replace the generic group-delay-ripple term with a target phase or target second-order dispersion curve. The present result should be viewed as a robust spectral reflector with chirped delay characteristics, not as a fully optimized end-to-end dispersion compensator.




Fig. 5 Mean-centered reflection group delay across the target pass band.

6.4 Deterministic fabrication-corner robustness

The deterministic corner test is where the nominal and robust solutions separate most clearly. Across the seven prescribed scenarios, the nominal design retains a worst scenario-average pass-band reflectivity of 0.99783, while the robust design gives 0.99748. On this average metric the nominal solution is marginally better. However, the pointwise minimum reveals a different picture: the nominal design falls to 0.96613 at its weakest in-band wavelength and corner, whereas the robust design stays above 0.98199. The robust solution therefore improves the corner-wise minimum reflectivity by about 1.59 percentage points.

The effect is stronger when expressed as amplitude ripple. Worst-corner ripple falls from 0.1476 dB for the nominal design to 0.0768 dB for the robust design, a reduction of approximately 48%. The robust design also

slightly reduces the worst group-delay percentile span, from 110.63 to 110.01 ps on the corner-analysis grid. These improvements are accompanied by a stop-band cost: worst maximum stop-band reflectivity rises from 0.00648 to 0.00755. The trade-off is therefore explicit and physically reasonable—the robust solution protects the weakest pass-band locations by accepting somewhat greater leakage near remote-band boundaries.

Figure 6 visualizes the range of spectra generated by the deterministic corners. The envelopes overlap across most of the central pass band, but the robust design constrains the lower edge of the envelope more tightly near the pass-band shoulders. The shape difference is subtle in a linear-scale plot because all central reflectivities are high; the tabulated minimum and ripple metrics provide the more sensitive comparison.




Fig. 6 Reflection envelopes over the seven deterministic fabrication scenarios. Filled regions span the minimum-to-maximum reflectivity at each wavelength.

Table 6 Worst metrics over deterministic fabrication corners

Design	Worst mean R	Worst min R	Worst ripple (dB)	Worst max stop R	Worst GD ripple (ps)
Baseline	0.640258	0.000530	32.6829	1.329e-05	111.8
Nominal optimized	0.997831	0.966127	0.1476	6.478e-03	110.6
Robust optimized	0.997479	0.981991	0.0768	7.546e-03	110.0

6.5 Monte Carlo uncertainty analysis

The Monte Carlo analysis provides an intentionally different perspective because its perturbations are concentrated near the nominal point instead of placing equal emphasis on stress-test corners. Across 180 common random samples, the nominal design has an expected mean pass-band reflectivity of 0.999594 and a 5th-percentile mean of 0.999141. The robust design gives 0.998993 and 0.997853, respectively. The 5th percentile of the pointwise minimum reflectivity is also slightly higher for the nominal design (0.99483 versus 0.99428). Thus the robust solution does not dominate in expected performance under this particular probability model.

This apparent reversal is not a contradiction. Robust optimization and stochastic expected-value optimization solve different problems. The robust objective was constructed to protect specified corners, and it succeeds on the most vulnerable corner-wise minimum and ripple metrics. The Monte Carlo distribution places relatively little

probability mass near the simultaneous extremes, so the small nominal cost of robustness is visible in the average statistics. For applications in which process data are available, a natural next step would be distributionally informed optimization that balances expected loss with a conditional-tail-risk term rather than using hand-selected corners.

Figures 7 and 8 show the Monte Carlo results. The baseline distribution is far below both optimized designs and much broader. The optimized designs cluster near unity mean reflectivity. The empirical cumulative distribution of minimum in-band reflectivity illustrates that low-reflectivity failures are eliminated over the sampled perturbation range for both optimized solutions, while the conventional baseline frequently exhibits severe edge loss.




Fig. 7 Monte Carlo distribution of mean in-band reflectivity for 180 common fabrication-uncertainty samples.




Fig. 8 Empirical cumulative distribution of the minimum in-band reflectivity under Monte Carlo perturbations.

Table 7 Monte Carlo summary (N = 180)

Design	E[mean R]	P05 mean R	P05 min R	Median ripple (dB)	P95 ripple (dB)	P95 max stop R
Baseline	0.715896	0.677308	0.010330	13.5552	19.8267	9.010e-05
Nominal optimized	0.999594	0.999141	0.994830	0.0073	0.0215	1.510e-03
Robust optimized	0.998993	0.997853	0.994275	0.0106	0.0229	1.773e-03

6.6 Discretization convergence

A computational paper must show that performance differences are not comparable to numerical error. Figure 9 and Table 8 report reflectivity errors for the robust design at five wavelengths, using an 1800-segment result as reference. With 80 segments, the maximum absolute error is approximately 1.6×10^{-3} . At 260 segments it falls below 1.6×10^{-4} , and at 800 segments below 3.2×10^{-5} . The nonmonotonic step between 120 and 180 segments reflects cancellation differences at selected wavelengths rather than divergence; the overall envelope decreases strongly with refinement.

The final nominal spectra use 800 segments, so their reflectivity discretization error at the sampled check wavelengths is roughly two orders of magnitude smaller than the difference between the baseline and optimized designs and smaller than the robust-versus-nominal corner effect. The corner and Monte Carlo studies use coarser meshes to control runtime, but those analyses compare the designs on common grids and their qualitative conclusions are stable.




Fig. 9 Transfer-matrix segmentation convergence for the robust design relative to an 1800-segment reference.

Table 8 Mesh-convergence error for the robust design

Segments	Max $ \Delta R $	RMSE in R
80	1.600e-03	9.524e-04
120	2.953e-04	1.462e-04
180	7.000e-04	3.379e-04
260	1.588e-04	7.336e-05
400	1.275e-04	7.471e-05
600	5.558e-05	2.697e-05
800	3.085e-05	1.448e-05
1200	1.060e-05	4.879e-06

6.7 Local design landscape

The local robustness landscape in Fig. 10 varies chirp span and apodization exponent while holding the other robust-design variables fixed. The lowest-loss region is elongated rather than point-like, indicating that chirp and taper can partially compensate one another. The optimum lies at the lower p boundary of the investigated slice, consistent with the full optimization. Moving to stronger tapering rapidly increases objective value because the pass-band shoulders lose coupling. Changes in chirp span shift the spectral edges relative to the fixed 1548-1552 nm target and therefore produce a pronounced valley.

This landscape also explains why a finite catalog of standard apodization windows can be restrictive. A named window fixes a specific relationship between center and shoulder coupling. The continuous p - a family permits the optimizer to move along a spectrum of taper strengths and shoulder redistributions. Even with only two shape parameters, that freedom is enough to create a high-reflectivity plateau while explicitly managing uncertainty.




Fig. 10 Local robust-objective landscape versus chirp span and apodization exponent with the remaining variables fixed at the robust optimum.

6.8 One-factor tolerance sweeps

To identify which uncertainty source drives each design, one-factor-at-a-time sweeps were performed over the full Monte Carlo truncation ranges while holding the other perturbations at nominal values. This analysis is not used in optimization and therefore serves as another diagnostic. Table 9 reports the worst minimum in-band reflectivity, worst ripple, and worst mean reflectivity encountered in each sweep.

For Bragg-wavelength shifts between -0.30 and $+0.30$ nm, the robust design is clearly more stable: its worst minimum reflectivity is 0.99648 versus 0.99212 for the nominal design, and its worst ripple is 0.0145 dB versus 0.0337 dB. This is consistent with the deterministic corner result because wavelength shift primarily threatens pass-band edge coverage. The stronger shoulder redistribution in the robust apodization profile appears to provide additional edge margin when the complete spectrum moves relative to the fixed target window.

The situation is different for coupling-strength uncertainty. When the global coupling scale is swept from 0.88 to 1.12, the nominal design retains a worst minimum reflectivity of 0.99486, whereas the robust design falls to 0.99023. The robust design also has larger worst ripple in this one-dimensional sweep. This confirms that the robust objective did not simply make the device insensitive to every individual perturbation. Instead, it optimized a multi-parameter scenario set in which simultaneous chirp and wavelength errors were important. Robustness is therefore directional in uncertainty space.

For isolated chirp-scale variation from 0.90 to 1.10, both designs remain strong. The nominal design reaches a worst minimum reflectivity of 0.99778 and the robust design 0.99549. Again, the nominal design is somewhat better when the uncertainty is restricted to a single parameter. The combined evidence suggests that the main

benefit of the robust solution is protection against compound edge-moving perturbations rather than universal desensitization.

This decomposition is useful for process engineering. If measured fabrication data later show that Bragg-wavelength placement dominates and coupling amplitude is tightly controlled, the robust profile is attractive. If coupling-dose variability dominates while spectral placement is accurate, the nominal solution or a reweighted robust objective may be preferable. A manufacturing-aware optimizer should therefore weight uncertainty dimensions using measured process capability rather than assuming equal importance.

Table 9 One-factor-at-a-time tolerance sweep summary

Design	Perturbation sweep	Worst min R	Worst ripple (dB)	Worst mean R
Nominal	Bragg shift (nm)	0.992124	0.0337	0.999449
Nominal	Coupling scale	0.994857	0.0175	0.998177
Nominal	Chirp scale	0.997781	0.0094	0.999479
Robust	Bragg shift (nm)	0.996481	0.0145	0.999025
Robust	Coupling scale	0.990232	0.0374	0.995991
Robust	Chirp scale	0.995486	0.0179	0.998305

6.9 Relation to recent computational FBG studies

The reported numbers should be compared with prior computational studies cautiously because target bandwidths, grating lengths, modulation amplitudes, system architectures, and performance definitions differ. Abdallah et al. (2026), for example, optimized hybrid apodization in cascaded gratings and reported approximately 99.12% reflectivity with a very narrow FWHM under their parameter set. The present study instead targets a four-nanometre high-reflectivity plateau and evaluates uncertainty around that plateau. A wider target inherently changes the appropriate chirp and coupling requirements, so direct percentage-improvement claims across the two studies would be misleading.

Similarly, Sayed et al. (2020) and Mustafa et al. (2022) evaluated CFBG or cascaded-FBG structures inside specific WDM dispersion-compensation systems, using Q-factor and BER as downstream metrics. Their results establish that apodization choice can materially affect link performance. The present work stops one level earlier in the modeling hierarchy and optimizes the grating response itself. This avoids coupling the claimed device improvement to one transmitter power, modulation format, amplifier noise figure, or fiber span, but it also means that system-level BER superiority is not demonstrated here.

Naguib et al. (2021) emphasized the influence of grating length and apodization on reflectivity, sidelobes, and dispersion. The current optimizer likewise moves toward longer and stronger gratings, reinforcing the physical importance of integrated coupling. The additional contribution is to place length, modulation, chirp, and continuous apodization-shape parameters in one optimization vector and then test that vector against explicitly defined uncertainty. In other words, the novelty lies more in the decision framework and robustness analysis than in proposing a new physical mechanism for Bragg reflection.

The transfer-matrix method itself is not new. Its value here is transparency and computational efficiency for a one-dimensional inverse problem. Morozov (2023) discusses transfer-matrix representations for one-dimensional optical structures; in the present application, exact segment propagation makes it inexpensive to evaluate thousands of candidate gratings. This permits robust optimization without replacing the governing physics by a learned surrogate. For higher-dimensional integrated photonics or strongly vectorial gratings, finite-element, eigenmode-expansion, or FDTD solvers would be required, and the same optimization philosophy could be retained at greater computational cost.

7 Discussion

The first major result is the scale of improvement obtained by joint continuous optimization relative to the conventional $p=2$ baseline. Both optimized gratings transform a sloped, edge-limited reflection response into a nearly unity 4 nm plateau. This result is not surprising in the sense that stronger integrated coupling and flatter apodization increase broadband reflectivity, but the inverse-design formulation quantifies how those ingredients should be combined within the defined bounds. More importantly, the framework exposes the accompanying stop-band and delay trade-offs instead of reporting peak reflectivity alone.

The second result is that robustness depends on the performance notion. Under deterministic corners, the uncertainty-aware design substantially improves the weakest point and worst ripple. Under the central Monte Carlo distribution, the nominal design has a small advantage in expected metrics. This distinction matters for engineering decisions. A network component whose failure threshold is driven by the worst channel edge may benefit from the robust design even if its mean reflectivity is slightly lower. Conversely, a device produced by a tightly controlled process with well-characterized normal variability may prefer an expectation-optimized design. Reporting only a single 'robustness percentage' would conceal this choice.

The third result concerns the optimized apodization shape. The robust solution uses a more negative cosine coefficient than the nominal one. In physical terms, this modifies the coupling shoulders while keeping the profile symmetric and smooth. The deterministic-corner improvement suggests that spectral edge stability is sensitive not only to total coupling strength but also to where that coupling is distributed along a chirped grating. Because different wavelengths reflect predominantly at different positions, spatial redistribution of coupling acts as a way to equalize margin across wavelengths. This is one reason chirped gratings are naturally suited to position-dependent inverse design.

Compared with prior work based on predefined apodization functions, the present approach shifts emphasis from ranking windows to optimizing a low-dimensional continuous family. Abdallah et al. (2026) showed that hybrid window combinations can improve FBG characteristics, and Mustafa et al. (2022) similarly demonstrated performance sensitivity to apodization choice. The current framework is complementary: it preserves an analytic, easily fabricated envelope but treats its shape parameters as optimization variables. The approach could readily incorporate Nuttall, Bartlett, Hamming, or spline-based basis functions as additional families, provided smoothness and fabrication constraints are enforced.

The device-level focus is another distinction. Recent optical-link studies often report BER, Q-factor, and dispersion compensation for a complete simulated transmission system (Sayed et al. 2020; Shaalan et al. 2025). Such metrics are useful when a specific link architecture is fixed. Here the grating is characterized through reflection amplitude and phase, so the result is portable: a system designer can insert the device response into a separate link simulator, coherent receiver model, or sensing architecture. The price is that this paper does not claim an end-to-end BER improvement.

The convergence analysis supports the numerical credibility of the transfer-matrix calculations. Morozov (2023) emphasized that transfer-matrix formulations provide an efficient representation of one-dimensional optical propagation. In the present chirped grating, piecewise-constant segmentation converges rapidly once several hundred segments are used. Because the global optimization is far more expensive than final verification, it is computationally efficient to search on a coarser mesh and then validate the selected candidates at high resolution. This search-then-verify pattern is broadly applicable to photonic inverse design.

From a fabrication perspective, the optimized peak modulation of roughly 5.8×10^{-4} and length of roughly 57-58 nm are not inherently implausible for strongly written gratings, but feasibility depends on fiber composition and inscription platform. More importantly, both solutions press against the selected bounds. A manufacturing-oriented extension should introduce explicit penalties for grating length, maximum UV dose, modulation slew rate, or thermal stability. Such costs would turn the present spectral optimum into a design that trades optical performance against fabrication burden.

The robust objective is also extensible. Temperature and axial strain can be incorporated as correlated Bragg-wavelength shifts; nonuniform dose errors can be represented by low-order spatial perturbation modes; and effective-index uncertainty can be wavelength dependent. In that setting, deterministic corners could be generated from a calibrated process covariance or from principal components measured across fabricated devices. Risk measures such as conditional value at risk could then bridge worst-case and expected-value design.

A further opportunity is phase-targeted optimization. For dispersion compensation, the desired reflection phase is not merely smooth; it should approximate a quadratic spectral phase that cancels the accumulated propagation phase of a known fiber span. The current group-delay term is intentionally generic. Replacing it with a least-squares or minimax phase-error functional would produce a design more directly comparable with link-level compensation studies. The same transfer-matrix framework and uncertainty machinery would remain applicable.

Finally, the transparent low-dimensional formulation provides a useful benchmark for more complex machine-learning or adjoint methods. A learned surrogate or neural optimizer should outperform this baseline in speed or design quality while preserving physical constraints and uncertainty calibration. Shaalan et al. (2025) demonstrate growing interest in machine learning for dispersion-compensation problems. A strong future comparison would use identical coupled-mode physics and uncertainty distributions to test whether learned models discover genuinely better profiles or merely reproduce the smooth solutions accessible with standard global optimization.

8 Limitations

This is a computational device study, not an experimental validation. The model assumes a scalar single guided mode, real refractive-index modulation, unit fringe visibility, no polarization dependence, no cladding-mode coupling, no radiation loss, and no wavelength dependence of the background effective index beyond the explicit $\beta = 2\pi n_{\text{eff}}/\lambda$ relation. These assumptions are appropriate for isolating first-order Bragg coupling but can be insufficient for a fabricated device with strong vector, loss, or higher-order effects.

The uncertainty magnitudes are stress-test assumptions rather than statistics measured from a specific phase-mask, point-by-point, or femtosecond writing process. The deterministic corner conclusion is therefore conditional: the robust design is more protected over the defined set, not over every conceivable fabrication error. Similarly, the Monte Carlo distribution is illustrative and should not be interpreted as a measured manufacturing yield distribution.

The objective weights and target spectral windows encode one design preference. Other applications may value stop-band rejection, grating length, group-delay linearity, or transition steepness more strongly. The optimized solutions also lie near upper bounds for length and modulation, demonstrating that omitted manufacturing cost influences the result. The numerical optimizer was run with fixed finite budgets, so global optimality is not proven.

Finally, the group-delay metric is computed from numerical differentiation of unwrapped reflection phase. Although percentile spans reduce sensitivity to isolated phase noise and high-resolution calculations are used for the final nominal values, a dispersion-focused follow-up should validate phase response with an independent solver and directly optimize the desired dispersion coefficient.

9 Conclusions

A five-parameter inverse-design framework has been developed for linearly chirped, continuously apodized fiber Bragg gratings. The approach combines coupled-mode theory, exact segment transfer matrices, global-plus-local derivative-free optimization, deterministic uncertainty corners, Monte Carlo validation, and mesh-convergence testing. Within the specified design box, joint optimization increases nominal mean reflectivity over 1548-1552 nm from 0.7275 for a conventional $\sin^2$ -apodized baseline to 0.9997 for the nominal optimum and 0.9991 for the robust optimum.

The uncertainty-aware objective does not improve every statistic. Instead, it changes the risk profile. Under the prescribed deterministic corners, it increases worst minimum in-band reflectivity from 0.9661 to 0.9820 and

reduces worst ripple from 0.1476 dB to 0.0768 dB relative to the nominal optimum, while increasing worst stop-band leakage. Under a centrally concentrated 180-sample Monte Carlo distribution, the nominal optimum retains a small advantage in average reflectivity. This contrast demonstrates why robustness should be defined through an explicit risk criterion rather than inferred from nominal performance.

The results support uncertainty-aware continuous apodization as a practical direction for computational FBG design. The next steps are to introduce measured fabrication distributions, wavelength-dependent material dispersion and loss, explicit device-cost constraints, and phase targets tied to a specified optical link. Experimental inscription and spectral characterization would then be required before claims about manufacturing yield or system-level communication performance could be made.

References

- Abdallah, R.M., Elbahi, F.M., Soliman, H.Y., Shaalan, I.E.: Dispersion mitigation using cascaded FBGs with hybrid apodization profiles. *Opt. Quantum Electron.* 58, 123 (2026). <https://doi.org/10.1007/s11082-026-08690-x>
- Chaoui, F., Aghzout, O., Chakkour, M., Mounir, Y.: Apodization optimization of FBG strain sensor for quasi-distributed sensing measurement applications. *Act. Passive Electron. Compon.* 2016, 6523046 (2016). <https://doi.org/10.1155/2016/6523046>
- Erdogan, T.: Fiber grating spectra. *J. Lightwave Technol.* 15(8), 1277–1294 (1997). <https://doi.org/10.1109/50.618322>
- Hill, K.O., Meltz, G.: Fiber Bragg grating technology fundamentals and overview. *J. Lightwave Technol.* 15(8), 1263–1276 (1997). <https://doi.org/10.1109/50.618320>
- Irawan, D., Azhar, A., Ramadhan, K.: High-performance compensation dispersion with apodization chirped fiber Bragg grating for fiber communication system. *J. Penelit. Pendidik. IPA* 8(2) (2022). <https://doi.org/10.29303/jppipa.v8i2.1521>
- Kashyap, R.: *Fiber Bragg Gratings*, 2nd edn. Academic Press, Oxford (2010). <https://doi.org/10.1016/C2009-0-16830-7>
- Kogelnik, H.: Coupled wave theory for thick hologram gratings. *Bell Syst. Tech. J.* 48(9), 2909–2947 (1969). <https://doi.org/10.1002/j.1538-7305.1969.tb01198.x>
- Lv, Z., Du, J., He, Z.: On-chip multichannel dispersion compensation and wavelength division MUX/DeMUX using chirped-multimode-grating-assisted counter-directional coupler. *Photonics* 11(2), 110 (2024). <https://doi.org/10.3390/photonics11020110>
- Morozov, G.V.: Two forms of transfer matrix for one-dimensional optical structures. *Opt. Quantum Electron.* 55, 1120 (2023). <https://doi.org/10.1007/s11082-023-05413-4>
- Mustafa, F.M., Abdelhalim, M.M., Aly, M.H., et al.: Dispersion compensation analysis of optical fiber link using cascaded apodized FBGs hybrid with maximum time division multiplexing transmission technique. *Opt. Quantum Electron.* 53 (2021). <https://doi.org/10.1007/s11082-021-03006-7>
- Mustafa, F.M., Sayed, A.F., Zaky, A.A.M., Khalaf, A.A.M., Aly, M.H.: Dispersion compensation in silica doped fiber using soliton transmission technique over cascaded FBG. *Opt. Quantum Electron.* 53 (2021). <https://doi.org/10.1007/s11082-021-02890-3>
- Mustafa, F.M., Sayed, A.F., Aly, M.H.: A reduced power budget and enhanced performance in a WDM system: a new FBG apodization function. *Opt. Quantum Electron.* 54, 471 (2022). <https://doi.org/10.1007/s11082-022-03876-5>
- Naguib, B.A., Ata, M.M., Alzalabani, M.M., Yousif, B.B.: Performance evaluation and enhancement of apodized fiber Bragg grating for dispersion compensation. *AIP Adv.* 11, 015231 (2021). <https://doi.org/10.1063/5.0026190>
- Othonos, A.: Fiber Bragg gratings. *Rev. Sci. Instrum.* 68(12), 4309–4341 (1997). <https://doi.org/10.1063/1.1148392>

- Sayed, A.F., Mustafa, F.M., Khalaf, A.A.M., Aly, M.H.: An enhanced WDM optical communication system using a cascaded fiber Bragg grating. *Opt. Quantum Electron.* 52 (2020). <https://doi.org/10.1007/s11082-020-02305-9>
- Sayed, A.F., Mustafa, F.M., Khalaf, A.A.M., Aly, M.H.: Apodized chirped fiber Bragg grating for postdispersion compensation in wavelength division multiplexing optical networks. *Int. J. Commun. Syst.* 33(14), e4551 (2020). <https://doi.org/10.1002/dac.4551>
- Sayed, A.F., Mustafa, F.M., Khalaf, A.A.M., Aly, M.H.: Spectral width reduction using apodized cascaded fiber Bragg grating for post-dispersion compensation in WDM optical networks. *Photon. Netw. Commun.* 41(3), 231–241 (2021). <https://doi.org/10.1007/s11107-021-00926-y>
- Sayed, A.F., Mustafa, F.M., Khalaf, A.A.M., Aly, M.H.: Symmetrical and post dispersion compensation in WDM optical communication systems. *Opt. Quantum Electron.* 53 (2021). <https://doi.org/10.1007/s11082-020-02663-4>
- Shaan, I.E., Selmy, M.I., Aly, M.H., et al.: Optical fiber dispersion compensation: supervised machine learning with regression approach. *Opt. Quantum Electron.* 57, 303 (2025). <https://doi.org/10.1007/s11082-025-08169-1>
- Sreeragi, G.R., Vijayakumar, N., Pradeep, R.: Modified CFBG-based approach for chromatic dispersion compensation in high speed DWDM links. *Appl. Opt.* 63(4) (2024). <https://doi.org/10.1364/AO.510360>

Statements and Declarations

Funding. The author declares that no funds, grants, or other support were received during the preparation of this manuscript.

Competing interests. The author has no relevant financial or non-financial interests to disclose.

Author contributions. This is a single-author manuscript. The named author is responsible for independently verifying the computational methods, numerical results, references, and final submitted text.

Data availability. No external datasets were used. All numerical data reported in this study were generated from the coupled-mode/transfer-matrix model and parameter values specified in the manuscript. The tabulated results provide the principal numerical outputs needed to reproduce the reported comparisons.

Code availability. No public code repository is associated with this manuscript draft. The computational workflow, equations, parameter bounds, grids, random seed, uncertainty model, and optimization settings are reported in sufficient detail to support an independent reimplement.

Ethics approval. Not applicable.

Consent to participate. Not applicable.

Consent for publication. Not applicable.

Appendix A: Transfer-matrix derivation and numerical checks

For completeness, consider a segment in which δ and κ are constant. The coupled-mode system has coefficient matrix $H = \begin{bmatrix} \delta & \kappa \\ -\kappa & -\delta \end{bmatrix}$. Direct multiplication gives $H^2 = (\delta^2 - \kappa^2)I = -\gamma^2 I$, with $\gamma^2 = \kappa^2 - \delta^2$. The matrix exponential $\exp(iHh)$ therefore separates into even and odd powers. Summing the even powers produces $\cosh(\gamma h)I$ and summing the odd powers produces $iH \sinh(\gamma h)/\gamma$, which yields Eq. (9). The expression remains valid for imaginary γ because hyperbolic functions analytically continue to trigonometric functions.

At $\gamma=0$ the quotient $\sinh(\gamma h)/\gamma$ has the limit h . The implementation evaluates that limit explicitly for sufficiently small $|\gamma|$. This prevents division by a near-zero complex number at the local transition $|\kappa| \approx |\delta|$. The transfer matrices are multiplied from the input toward the output so that the final matrix maps the input amplitudes to the far-end amplitudes.

With no backward incident wave at $z=L$, $B(L)=0$. Writing $[A(L), B(L)]^T = T[A(0), B(0)]^T$ and setting $A(0)=1$ gives $0 = T_{21} + T_{22} r$, hence $r = -T_{21}/T_{22}$. In a lossless formulation, numerical solutions remain bounded by unit reflectivity to floating-point tolerance. The reported spectra were additionally checked for finite values and continuity over the wavelength grid.

Group delay is obtained by unwrapping $\arg r$ and differentiating with respect to angular frequency. Since ω decreases as wavelength increases, using the explicit ω array avoids sign mistakes that can occur when differentiating phase directly with respect to wavelength. Only the relative variation across the pass band is used in the plots and objective, so constant reference delay is immaterial.

$$H^2 = (\delta^2 - \kappa^2)I = -\gamma^2 I \quad (17)$$

$$\exp(iHh) = I \cosh(\gamma h) + iH \sinh(\gamma h)/\gamma \quad (18)$$

Appendix B: Reproducibility recipe

A reproducing implementation can follow the sequence below. (1) Define the five design variables and bounds from Table 2. (2) For each candidate, construct the local Bragg wavelength and grating period at segment midpoints. (3) For every wavelength sample, compute β , δ , κ , and the exact segment transfer matrices. (4) Multiply segment matrices and calculate complex reflection. (5) Evaluate pass-band, stop-band, and group-delay metrics. (6) Apply Eq. (14) for nominal optimization or Eq. (15) across the seven Table 3 scenarios for robust optimization. (7) Run differential evolution with the fixed seeds and follow with Nelder-Mead. (8) Recompute the selected candidates on the 1201-point/800-segment verification grid. (9) Perform the independent Monte Carlo test using seed 20260821 and the distributions given in Sect. 5. (10) Confirm mesh convergence against a higher-segment reference.

The intentionally low dimension means independent implementations should recover profiles with the same qualitative characteristics even if stochastic optimization terminates at slightly different parameter values. Exact numerical equality requires the same optimization library version, population initialization, floating-point behavior, and stopping criteria. The scientific conclusions rely on large spectral differences between baseline and optimized devices and on explicitly reported robust-versus-nominal trade-offs rather than on the last optimization decimal.

Appendix C: Interpretation of optimization-bound activity

Both optimized solutions approach the upper bounds on grating length and peak index modulation. This appendix clarifies why that behavior is expected under the present objective. For a uniform grating at exact Bragg resonance, the idealized peak reflectivity scales approximately as $\tanh^2(\kappa L)$. Increasing either κ or L therefore monotonically increases peak reflection toward unity in the lossless model. Chirp spreads resonant

interaction across wavelength and reduces the effective interaction length available to each spectral component, so strong integrated coupling becomes even more valuable when a broad high-reflectivity plateau is requested.

The optimization includes no explicit cost proportional to physical length, UV exposure, writing time, induced loss, or peak modulation. Consequently, once stop-band and ripple penalties can be controlled by apodization and chirp, the objective has little reason to prefer shorter or weaker gratings. The active upper bounds are thus evidence about the objective model rather than evidence of a universal physical optimum. A practical manufacturing optimization should add at least one cost or constraint that increases with L and $\Delta n_{\max}$.

One simple extension is a regularized objective $J_{\text{total}} = J_{\text{spectral}} + w_L(L/L_{\text{ref}})^2 + w_n(\Delta n_{\max}/\Delta n_{\text{ref}})^2$. Another is to constrain the integrated modulation dose, proportional in a first approximation to the integral of $\Delta n(z)$ over z . A third approach is multi-objective optimization that returns a Pareto front between spectral performance and fabrication burden. These formulations would allow an engineer to select a point that meets a reflectivity specification with the shortest or lowest-dose grating rather than automatically saturating the search bounds.

Bound activity also affects the interpretation of the apodization exponent. Because the optimizer can obtain large total coupling from L and $\Delta n_{\max}$, it is free to flatten the apodization profile to improve edge coverage. If length or modulation were tightly constrained, stronger central weighting might re-emerge as useful for controlling sidelobes. The optimal p and a values reported here should therefore be understood jointly with the chosen fabrication-scale bounds, not as universal constants for CFBGs.

This observation does not invalidate the robust comparison because the nominal and robust solutions share the same bounds and similar overall scale. The main difference between them is the shape coefficient and resulting distribution of coupling along the grating. Nevertheless, future experimental work should repeat the optimization with platform-specific constraints before fabricating the exact parameter values reported in Table 4.

Appendix D: Algorithmic complexity and implementation pseudocode

For M sampled wavelengths and N grating segments, one spectral evaluation requires $O(MN)$ operations because each segment updates four scalar transfer-matrix elements for every wavelength. Memory use is $O(M)$ when the wavelength dimension is vectorized and segment matrices are accumulated in place. The cost of robust optimization is multiplied by the number of uncertainty scenarios. This simple scaling explains the two-resolution strategy used in the study: coarse grids are efficient for global search, whereas high-resolution grids are reserved for final verification and figures.

A practical implementation begins by allocating complex arrays $M_{00}=M_{11}=1$ and $M_{01}=M_{10}=0$ for all wavelengths. At each segment midpoint, compute the local Bragg wavelength, period, detuning, and coupling vector. Evaluate γ and the stable factor $s = \sinh(\gamma h)/\gamma$, replacing s by h when $|\gamma|$ is sufficiently small. Form the four elements of the segment propagator and left-multiply them into the running transfer-matrix arrays. After the final segment, calculate $r = -M_{10}/M_{11}$ and $R = |r|^2$. Unwrap phase only when group delay is required, because phase processing is unnecessary for objectives that depend solely on reflectivity.

The optimization wrapper should validate bounds before calling the optical solver, compute pass-band and stop-band masks once for each wavelength grid, and reuse the same deterministic uncertainty scenarios for every candidate. For the robust objective, loop over the scenario list, collect the metric vector for each scenario, and then apply the average and worst-case penalties. Differential evolution can operate directly on this scalar objective. The best global candidate is then passed to Nelder-Mead for local refinement without assuming differentiability.

For reproducibility, implementation metadata should record the wavelength grid, segment count, optimization bounds, pseudorandom seed, objective weights, stopping tolerances, and final parameter vector. The code used for this study followed that practice, and the manuscript reports the values that materially affect the numerical result. Independent implementations may differ at the final optimization decimal because stochastic population

methods and library versions initialize and terminate differently; agreement should instead be assessed through spectra, reported metrics, and the qualitative robust-versus-nominal trade-off.

Appendix E: Dimensional audit and replication diagnostics

A useful replication check is dimensional consistency before any optimization is attempted. Wavelength, grating period, and longitudinal position must be expressed in the same length unit inside the phase terms. The relation $\lambda_B = 2 n_{\text{eff}} \Lambda$ fixes the local period, $\beta = 2 \pi n_{\text{eff}} / \lambda$ and π / Λ both have units of inverse length, and therefore detuning δ has units of inverse length. The coupling coefficient $\kappa = \pi \Delta n / \lambda$ has the same units. Multiplication by segment length h then produces dimensionless arguments for $\cosh$, $\sinh$, and the matrix exponential. Mixing nanometres for wavelength with metres for z without explicit conversion changes both δh and κh by orders of magnitude and is a particularly easy way to obtain a numerically smooth but physically meaningless spectrum.

A second diagnostic is the zero-coupling limit. Setting Δn_{max} to zero must make κ vanish everywhere and therefore drive reflected power toward zero across the simulated wavelength interval. A complementary check is a uniform, unchirped grating: when the detuning is zero and the coupling is constant, the numerical transfer-matrix result should approach the standard $\tanh$ -squared dependence on integrated coupling. These two limits exercise opposite ends of the solver without depending on the optimization routine. They are more informative than checking only that a plotted spectrum looks plausible.

The uncertainty analysis can also be replicated independently of the optimizer. Take the final nominal and robust parameter vectors from Table 4, perturb only the Bragg wavelength, coupling scale, and chirp scale according to Sect. 5, and recompute the same pass-band masks. The deterministic test should recover the qualitative ordering reported in Table 6: the robust profile has the stronger corner-wise minimum and smaller worst ripple, while its worst remote stop-band leakage is modestly larger. The Monte Carlo test should recover the different qualitative result in Table 7: because probability is concentrated near the nominal point, the nominal optimum retains a slight advantage in expected reflectivity. Agreement with both orderings is a stronger diagnostic than matching a single decimal value.

Finally, numerical resolution should be treated as part of the model specification rather than as a plotting choice. The optimizer uses a deliberately coarse segmentation to reduce search cost, but all headline nominal spectra are recomputed at 800 segments and 1201 wavelengths. Table 8 shows that the remaining reflectivity discrepancy relative to an 1800-segment reference is small compared with the baseline-to-optimized spectral change and with the robust-versus-nominal deterministic-corner difference. A replication that uses a different mesh should therefore report its own convergence check instead of assuming that identical optimizer parameters alone guarantee identical spectra.