

Article

Physics-Verified Spectral Dreaming Enables Interpretable and Manufacturable Inverse Design of Multilayer Radiative Coolers

Jiajun Wang  and Xiuye Liu * 

College of Applied Physics, Chengdu University of Technology, Chengdu 610059, China;
202320030121@stu.cdut.edu.cn

* Correspondence: liuxiuye@cdut.edu.cn

Abstract

Optical inverse design faces a dilemma: neural surrogates enable fast, differentiable search but can yield physically unreliable pseudo-optima, whereas solver-in-the-loop optimization is reliable yet costly. Most surrogate methods also trust the surrogate throughout the search, train separate models for performance prediction and structure optimization, and remain largely black-box. We propose Physics-Verified Spectral Dreaming (PVSD), a unified framework for forward prediction, inverse design, and physical interpretability: a frozen differentiable spectral surrogate “dreams” structural mutations by input-gradient ascent to explore the design space, while a physical solver adjudicates every accepted update—the surrogate proposes, physics decides. We instantiate it as PVSD-TMM for one-dimensional multilayer radiative coolers. The forward predictor attains $R^2 = 0.9936/0.9964/0.9828$ for net cooling power, solar reflectance, and primary-window emissivity; neural dreaming lifts the population-mean net cooling power of 1000 random seeds from -466.7 to 65.8 W m^{-2} (91.4% reaching net cooling), and continuous-thickness refinement with 5 nm rounding yields a 14-layer manufacturable final design. Independent COMSOL finite-element and analytic TMM cross-validation converge to $P_{\text{cool}} \approx 172 \text{ W m}^{-2}$, $R_{\text{solar}} \approx 0.970$, and $\epsilon_{\text{win}} = 0.9252$. This is a full-spectrum radiative-balance result for an idealized radiative-only case ($h_{\text{conv}} = 0$), not a window-emittance-only metric; PVSD thus achieves high simulated broadband radiative-cooling performance under the stated assumptions, without claiming global optimality.

Keywords: radiative cooling; inverse design; deep learning; multilayer thin films; transfer-matrix method; surrogate model; neural dreaming; cross-solver validation

1. Introduction

Passive daytime radiative cooling dissipates heat to outer space through the 8–13 μm atmospheric transparency window while reflecting incident sunlight in the 0.3–2.5 μm solar band, achieving sub-ambient cooling with no energy input. It is valuable for building energy savings, outdoor electronics thermal management, photovoltaic cooling, and low-carbon thermal control [1,2]. A high-performance daytime cooler must simultaneously satisfy two seemingly conflicting spectral goals: near-total reflectance in the solar band to suppress solar heating, and high emittance in the atmospheric window to maximize net radiative loss to the sky.

Multilayer optical thin films are a classic and practical route to selective radiative cooling [3]. Compared with complex metasurfaces, randomized photonic metamaterials [4], and radiative-cooling structural materials [5], multilayers offer a simple one-dimensional



Received: 8 June 2026

Revised: 9 July 2026

Accepted: 15 July 2026

Published: 21 July 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

geometry, mature deposition, a broad material palette, and spectra rapidly computable by the transfer-matrix method (TMM); they are thus both a realistic device platform and an ideal model system for validating optical inverse-design methods.

However, multilayer inverse design is intrinsically a high-dimensional mixed optimization problem over material identity, layer count, layer order, and thickness. For M candidate materials, K discrete thicknesses, and N layers, the discrete combinatorial space alone reaches $(MK)^N$; admitting continuous thickness makes it larger and strongly non-convex. Traditional black-box optimizers such as genetic algorithms (GAs) and particle-swarm optimization (PSO) are reliable because they call the physical solver directly, and have been applied to radiative-cooling multilayers [6–8], but they typically require large evaluation budgets and rarely explain *why* a structure works.

Deep learning offers new tools for optical inverse design [9]: tandem networks [10], generative models (GANs/VAEs/diffusion), reinforcement learning, and surrogate-plus-optimizer schemes can all propose candidates quickly; recent examples span machine-learning inverse design of radiative-cooling metasurface emitters [11] and machine-learning-guided optimization of photonic spectral structures [12]. Yet most do not unify performance prediction and structure optimization in a single module, remain largely black-box with limited physical interpretability, and trust the neural surrogate throughout optimization—scoring or backpropagating against surrogate predictions, at best re-checking a few final candidates—at the well-known cost of surrogate pseudo-optima: high predicted performance but poor true physics. An underappreciated point is that this “trust-without-verification” gap is benign when the true solver is cheap, as for one-dimensional multilayers, whose TMM gradient is itself exact and inexpensive. It becomes unacceptable, however, once inverse design moves to expensive three-dimensional electromagnetics (FDTD/RCWA) or experimental loops, where trusting a possibly biased surrogate is untenable. What is needed, therefore, is not merely a “better surrogate” but a paradigm that places the physical solver at the core of the optimization loop, letting physics—not the surrogate—adjudicate every update.

DeepDream/neural inceptionism [13] offers exactly such a mechanism on the proposal side: rather than training a separate generator, one freezes a trained forward model and optimizes its *input* by gradient ascent so the output moves toward a target. This suits multilayers especially well, because a film maps naturally to a material/thickness token sequence with clear physical meaning per token.

We propose Physics-Verified Spectral Dreaming (PVSD), which makes this proposal mechanism part of a closed loop between a neural surrogate and a physical solver: the neural model proposes candidate changes, and the solver verifies every accepted update, so the surrogate only proposes and physics decides. We instantiate PVSD as PVSD-TMM for multilayer radiative coolers (Figure 1). The “physics-verified” principle runs throughout—not only in the TMM-in-the-loop design stage, but also in a post-design cross-solver validation of the final design (Section 2.10), where COMSOL 6.3 finite-element (FEM) modeling and an independently implemented analytic TMM re-check the same structure.

Three lines of work are most closely related, which we distinguish to delimit our contribution. (i) *Frozen-surrogate, input-gradient inverse design*: Peurifoy et al. [14] train a forward surrogate and backpropagate to the input to design layered nanoparticle spectra—mechanistically the same as our “dreaming”—and Kim and Kim [15] likewise take surrogate input-gradients for one-dimensional multilayers (a neural-adjoint scheme). Our distinction is not the gradient mechanism but placing a hard physical accept/reject in every loop iteration and auditing the final structure with an independent solver, so physics, not the surrogate, governs every accepted update. (ii) *Deep-learning multilayer radiative coolers*: Zhou et al. [16] inverse-design multilayer coolers by deep learning—closely related to

our work—and machine-learning- or genetic-algorithm-based designs of radiative-cooling multilayers also exist [6,8]; together with ours these form a “machine learning + optimization for radiative-cooling multilayers” family, and our increment is the physics-verification loop, population-level evidence (1000 structures), interpretability, and cross-solver auditing rather than another set of design numbers. (iii) *Dreaming as a principled, named method*: it originates in feature visualization and activation maximization [13,17,18] and has been used for inverse design in molecules [19] and metal–organic frameworks [20]. Input-gradient ascent itself is not new; to our knowledge, however, the optics and thin-film inverse-design literature has not commonly framed this strategy as neural dreaming. The contribution here is its integration with hard solver-based acceptance and post-design cross-solver validation, demonstrated end-to-end on radiative cooling.

Contributions.

- PVSD, a physics-verified neural dreaming framework for optical inverse design that couples differentiable input optimization with closed-loop solver verification.
- PVSD-TMM for 1D multilayer radiative coolers, where every accepted proposed update is verified by three-angle TMM.
- Evidence of population-level optimization, multi-solution diversity, and interpretability over 1000 random seeds, same-seed restarts, and attention analysis.
- A 14-layer manufacturable final design obtained via Super-DD+Refine, continuous-thickness polishing, and 5 nm rounding, with post-design COMSOL FEM and analytic TMM validation giving $P_{\text{cool}} \approx 172 \text{ W m}^{-2}$ under the radiative-only assumption $h_{\text{conv}} = 0$.




Figure 1. PVSD-TMM framework: a design space of material/thickness tokens; a frozen spectral surrogate learner; a neural dreaming optimizer; a physics verifier (TMM); and a manufacturable-refinement stage. The workflow separates neural proposal generation from solver-based acceptance; the gray dashed arrow denotes the verifier’s accept/reject feedback to the dreaming optimizer.

2. Results

2.1. Concept: Neural Gradient Exploration, Physical-Solver Adjudication

PVSD resolves the central tension of surrogate inverse design: relying solely on a neural predictor invites pseudo-optima from predictor error, while relying solely on a solver makes search costly and undirected. PVSD divides labor: the differentiable surrogate supplies update *directions*; the physical solver retains final adjudication. PVSD-TMM has five modules: (i) a spectral surrogate learner mapping material/thickness token sequences to spectral KPIs; (ii) a neural dreaming optimizer that freezes the surrogate and optimizes the input token distribution to mutate structures toward target performance; (iii) a physics verifier that scores each decoded candidate by TMM; (iv) a verified accept–reject loop that accepts only true-value improvements; and (v) manufacturable refinement (continuous-thickness polishing, rounding, robustness). The verifier is replaceable—TMM for multilayers, RCWA/FDTD/FEM for gratings or metasurfaces, even measured feedback in an experimental loop. Section 2.10 extends “physics verification” beyond design completion, to an independent FEM cross-check.

2.2. The Forward Predictor Builds a Reliable Differentiable Landscape

The forward predictor provides a reliable, differentiable performance landscape for the subsequent neural dreaming: on 10,000 held-out structures, the parity plots for P_{cool} , R_{solar} , and ε_{win} cluster tightly on the diagonal against TMM ground truth (Table 1). To obtain it, we generate 100,000 TMM-labeled multilayer samples, each alternating material and thickness tokens on a fixed Ag substrate (material pool SiO_2 , TiO_2 , Si_3N_4 , MgF_2 , Al_2O_3 , HfO_2 , ZnS , ZnSe , SiC , Ge , Si ; thickness tokens 10–2000 nm); the predictor uses an LSTM with scalar attention pooling and an MLP head, jointly predicting net cooling power P_{cool} , solar reflectance R_{solar} , and primary-window emissivity ε_{win} on an 80k/10k/10k train/val/test split.

Table 1. Forward-predictor accuracy on the held-out test set.

Target	R^2	MAE	RMSE
P_{cool}	0.9936	10.85 W m^{-2}	14.69 W m^{-2}
R_{solar}	0.9964	0.00904	0.01253
ε_{win}	0.9828	0.02502	0.03642

2.3. Neural Dreaming Optimizes Low-Performing Films into Net Coolers

To show that updates are not black-box jumps, we trace a single low-performing seed. The initial $\text{SiC}(200)/\text{ZnSe}(1500)/\text{Al}_2\text{O}_3(20)/\text{Ag}$ has $P_{\text{cool}} = -418.1 \text{ W m}^{-2}$, $R_{\text{solar}} = 0.506$, $\varepsilon_{\text{win}} = 0.224$. During dreaming, the frozen predictor supplies input-gradient directions and TMM verifies each decoded proposal. The final $\text{TiO}_2(2000)/\text{SiO}_2(1500)/\text{SiO}_2(20)/\text{Ag}$ reaches $P_{\text{cool}} = 69.9 \text{ W m}^{-2}$, $R_{\text{solar}} = 0.904$, and $\varepsilon_{\text{win}} = 0.700$ —a 487.9 W m^{-2} gain. Neural gradients steer the structure toward high-performance regions while TMM blocks surrogate-error-driven unreliable updates (Figure 2).




Figure 2. Forward-predictor parity and single-seed optimization trajectory: (a–c) predicted vs. true P_{cool} , R_{solar} , and ε_{win} for 10,000 held-out structures (point color indicates local point density); (d) the P_{cool} dreaming trajectory of a single low-performing seed; (e) the same seed's initial vs. final spectra, with the solar band and atmospheric window shaded.

2.4. Same-Seed Dreaming Yields Diverse High-Performing Solutions

A useful inverse-design framework should return many distinct, comparably good candidates rather than one answer. From five low-performing seeds, each with 12 independent DD+Refine restarts (TMM budget 500), all 5/5 seeds yield 12/12 distinct final structures, and all final structures from these seeds satisfy $P_{\text{cool}} > 80 \text{ W m}^{-2}$ (Figure 3). Multi-solution output enables secondary screening on cost, deposition window, thermal stability, or environmental compatibility.




Figure 3. Same-seed diversity: from a single low-performing seed, PVSD-TMM finds many distinct high-performing structures. Footnote values are optimization-time KPIs; the final design's independently verified values are reported in Section 2.8.

2.5. 1000 Random Structures Show Population-Level Performance Migration

Running the pipeline on 1000 random seeds and evaluating all initial and final structures by three-angle TMM, the population mean and median P_{cool} both migrate from strongly negative to net-positive, with comparable mean and median gains (full values in Table 2, Figure 4). PVSD-TMM does not merely tune one structure; it pushes an entire random population from the net-heating regime into net cooling.

Table 2. Population-level performance migration over 1000 random seeds.

Metric	Seed	Final	Change
mean P_{cool} (W m ⁻²)	−466.711	65.830	+532.541
median P_{cool} (W m ⁻²)	−470.546	77.675	+534.638
mean R_{solar}	0.4245	0.9320	+0.5074
mean ε_{win}	0.5132	0.7391	+0.2259
$P_{\text{cool}} > 0$ success	0.0%	91.4%	+91.4%
$P_{\text{cool}} > 80$ success	0.0%	46.2%	+46.2%

The population KPIs are optimization-time values from the original TMM pipeline, characterizing the population-level migration; the relative migration trends and rankings

are robust. The final design's performance is established by post-design cross-solver validation (Sections 2.8 and 2.10).




Figure 4. Population shift over 1000 random seeds: before/after P_{cool} distributions, paired improvement, migration in the $(R_{\text{solar}}, \epsilon_{\text{win}})$ plane, and success-rate statistics. Dashed vertical lines in (a) mark the population means; in (b) the dashed diagonal is the $y = x$ no-change line and the dotted horizontal line marks the 5 nm champion.

2.6. Attention Reveals a Focus on Thickness Control and Key Materials

From 1000 optimized final structures we extract 23,500 active token attention records. Thickness tokens receive moderately higher mean attention than material tokens ($1.25\times$; Table 3, Figure 5). This is consistent with multilayer interference physics: materials set index contrast and absorption/emission, while thickness sets phase accumulation, cavity conditions, and peak positions.

Table 3. Attention statistics over 1000 optimized structures.

Item	Value
Total structures	1000
Total tokens	23,500
Attention range	0.0007–0.2076
Mean attention, material tokens	0.0163
Mean attention, thickness tokens	0.0205
Thickness/material ratio	$1.25\times$




Figure 5. Scalar attention interpretability: material and thickness token attention, token-type ratio, and token-frequency statistics. Colors distinguish materials in (a,d) and token types in (c).

2.7. Same-Seed, Equal-Budget Benchmark ($n = 20$)

We compare DD+Refine against genetic-algorithm (GA) and particle-swarm (PSO) baselines under a same-seed, equal-budget, same-evaluator protocol: all methods start from the same 20 low-performing seeds, use the identical three-angle TMM evaluator, and are capped at equal numbers of TMM calls (500–5000). We report the mean best P_{cool} over the 20 seeds with its standard error and test each pairwise difference with a Wilcoxon signed-rank test (paired by seed; a paired t -test gives the same conclusions), as shown in Figure 6 and Table 4.




Figure 6. Same-seed, equal-budget benchmark vs. GA/PSO ($n = 20$ seeds): mean best P_{cool} ($\pm$ SEM) across TMM-call budgets. DD+Refine is significantly better than PSO at every budget ($p < 10^{-3}$); the DD-vs-GA advantage is modest and reaches significance only at the largest budget. Grey labels above each budget give the DD-vs-GA Wilcoxon signed-rank outcome (n.s. = not significant, $p \geq 0.05$; ** $p < 0.01$; *** $p < 0.001$); DD-vs-PSO is significant ($p < 10^{-3}$) at every budget.

DD+Refine attains the highest mean at every budget. Its advantage over PSO is large and significant at all budgets (+24 to +48 W m⁻², $p < 10^{-3}$). Its advantage over the stronger GA baseline is smaller and consistent (+5 to +9 W m⁻²) but, given seed-to-seed variability, reaches statistical significance only at the largest budget ($p = 4 \times 10^{-3}$ at 5000 calls; not significant at 500–3000). We therefore state the result conservatively: *DD+Refine is competitive with GA—consistently ahead on average but not significantly so at low-to-moderate budgets—and clearly superior to PSO, at equal TMM budget.* (PSO here optimizes a continuous re-encoding of the discrete material/thickness space, which likely disadvantages it; the more informative comparison is against GA.) Consistent with our framing throughout, the claim is the budget efficiency of the neural proposal mechanism as a competitive optimizer, not absolute superiority over all optimizers.

For runtime (Table 4), each run’s wall-clock time is dominated by the TMM evaluations and is therefore comparable across methods at fixed budget: on a single CPU core of the compute cluster (AMD EPYC), a 500-call run takes ≈13–15 s and a 5000-call run ≈130–190 s (DD+Refine ≈15 s and ≈186 s), with the number of TMM calls equal to the budget by construction. The per-call solver cost is thus the same for all methods; PVSD’s efficiency lies in converting a given call budget into higher net cooling power, not in a cheaper per-call cost.

Table 4. Same-seed benchmark ($n = 20$): mean best P_{cool} ($\pm$ SEM), DD-vs-GA significance (Wilcoxon signed-rank), and mean runtime. DD-vs-PSO is significant at all budgets ($p < 10^{-3}$). n.s. = not significant ($p \geq 0.05$); ** $p < 0.01$.

Budget (TMM Calls)	DD+Refine (W m ⁻²)	GA (W m ⁻²)	PSO (W m ⁻²)	DD vs GA (Wilcoxon p)	Runtime (DD) (s)
500	124.8 $\pm$ 2.3	119.1 $\pm$ 3.7	101.0 $\pm$ 4.5	0.23 (n.s.)	14.7
1000	137.9 $\pm$ 2.4	132.1 $\pm$ 3.4	100.2 $\pm$ 6.6	0.076 (n.s.)	30.4
2000	147.9 $\pm$ 1.6	143.1 $\pm$ 2.5	100.2 $\pm$ 8.8	0.23 (n.s.)	66.5
3000	150.4 $\pm$ 1.8	143.4 $\pm$ 2.1	105.5 $\pm$ 9.5	0.053 (n.s.)	107.4
5000	152.5 $\pm$ 1.1	144.0 $\pm$ 2.3	111.6 $\pm$ 5.5	0.0042 (**)	185.7

2.8. PVSD-TMM Finds a High-Performance Manufacturable Final Design

After the main validation, we perform final-design refinement to obtain a representative manufacturable structure. Super-DD+Refine identifies a high-performing discrete structure in the token space; further discrete search yields diminishing improvement. Fixing the material sequence, we relax the 14 thicknesses to continuous variables and polish against the official three-angle TMM, then round to a 5 nm manufacturable grid (Table 5). The 14-layer stack (≈9.34 μ m total on a ≥ 200 nm Ag substrate) has the converged, independently cross-validated performance of Table 6. We use this 14-layer design as a representative demonstration structure throughout the paper; how the net cooling power depends on the number of layers—and why 14 is a practical rather than an optimal choice—is examined in Section 2.14.

The energy balance is $P_{\text{cool}} = P_{\text{rad}} - P_{\text{atm}} - P_{\text{solar}} - P_{\text{conv}} \approx 282.0 - 79.9 - 30 - 0 \approx 172$ W m⁻². This value is a full-spectrum radiative-balance result. The final design reflects nearly all sunlight in 0.3–2.5 μ m and emits strongly in the primary atmospheric window, with $\varepsilon_{\text{win}} = 0.9252$ over 8–13 μ m; additional long-wave emission outside the primary window also contributes to P_{rad} and is accounted for through the same atmospheric-transmittance-weighted balance. Thus, ε_{win} is an important window-selectivity metric, but the reported P_{cool} is established by the complete P_{rad} , P_{atm} , and P_{solar} decomposition rather than by the window emittance alone. Because $P_{\text{conv}} = 0$ here, this value should be interpreted as an idealized radiative-only upper-bound performance

at $T_{\text{surface}} = T_{\text{ambient}} = 300$ K, not as a prediction of open-air operating cooling power without a heat-transfer model (Figure 7).

Table 5. Final 5 nm rounded manufacturable design.

Layer	Material	Thickness (nm)
1	ZnSe	160
2	HfO ₂	570
3	Si ₃ N ₄	160
4	TiO ₂	2235
5	SiO ₂	65
6	TiO ₂	140
7	SiO ₂	45
8	TiO ₂	1200
9	HfO ₂	1015
10	Si ₃ N ₄	1200
11	TiO ₂	135
12	Si ₃ N ₄	1425
13	TiO ₂	935
14	MgF ₂	55

Table 6. Converged final-design performance (cross-solver validation).

Metric	Converged Value	Note
P_{cool}	$\approx 172 \text{ W m}^{-2}$	COMSOL FEM 172.16; analytic TMM ≈ 172.8
R_{solar}	≈ 0.970	COMSOL 0.9698/TMM 0.9705
ϵ_{win}	0.9252	primary-window metric; same three-solver value
P_{rad}	$\approx 282.0 \text{ W m}^{-2}$	three-way agreement (mid-IR grid-robust)
P_{atm}	$\approx 79.9 \text{ W m}^{-2}$	three-way agreement
P_{solar}	$\approx 30 \text{ W m}^{-2}$	from converged R_{solar}
P_{conv}	0	idealized radiative-only case; $h_{\text{conv}} = 0$




Figure 7. Final 5 nm rounded design: (a) 14-layer stack and KPIs; (b) three-angle spectrum with solar and atmospheric-window backdrops; (c) layer-wise sensitivity (worst of $\pm 1/2/5\%$); (d) Monte Carlo robustness ($n = 300$ per tolerance).

2.9. Robustness Analysis: Manufacturing Tolerance and Sensitive Layers

To check that the final response is not a non-manufacturable numerical spike, we analyze rounding, random fabrication error, and single-layer sensitivity, all recomputed with the same official three-angle TMM. Rounding the continuous final design to the 5 nm grid changes P_{cool} by only 0.2 W m^{-2} ($171.83 \rightarrow 171.61$), indicating that the structure is not a fractional-thickness artifact. Monte Carlo independent thickness perturbation (Table 7) indicates strong tolerance: even at $\pm 10\%$, P_{cool} stays at $170.6 \pm 0.9 \text{ W m}^{-2}$ (worst of 300 samples 167.6, i.e., 97.7% of nominal). Single-layer sensitivity (Figure 7c): each single-layer thickness perturbation ($\pm 1/2/5\%$) reduces P_{cool} by less than 1 W m^{-2} (converged worst case $\approx 0.78 \text{ W m}^{-2}$), with the ultra-thin $\text{SiO}_2/\text{MgF}_2$ boundary layers being the least sensitive.

Table 7. Monte Carlo thickness-error robustness ($n = 300$).

Perturbation	Mean P_{cool}	Std	Worst of 300
$\pm 1\%$	171.44	0.23	170.88
$\pm 2\%$	171.37	0.24	170.71
$\pm 5\%$	171.25	0.46	169.92
$\pm 10\%$	170.55	0.87	167.61

2.10. Independent Cross-Solver Validation

Physics verification should not end with the in-loop TMM accept/reject but should continue after design completion using solver implementations that are as independent as practical. We therefore cross-validate the final design using a COMSOL 6.3 Wave Optics periodic finite-element (FEM) model and a separately implemented analytic Byrnes TMM, both using the same frozen optical constants (n, k) and the discipline “fix the target metrics first, solve independently, never tune toward any original KPI.” The FEM validation was re-checked on a licensed COMSOL installation, with consistent results.

The three solvers give the same reported primary-window emittance ($\varepsilon_{\text{win}} = 0.9252$), supporting consistency of the key 8–13 μm optical metric. The two independent solvers also corroborate the full-spectrum net cooling power and solar reflectance: $P_{\text{cool}} \approx 172 \text{ W m}^{-2}$ (COMSOL 172.16, analytic TMM ≈ 172.8) and $R_{\text{solar}} \approx 0.970$ (COMSOL 0.9698/analytic TMM 0.9705). The energy balance $P_{\text{cool}} = P_{\text{rad}} - P_{\text{atm}} - P_{\text{solar}} - P_{\text{conv}} \approx 282.0 - 79.9 - 30 - 0 \approx 172 \text{ W m}^{-2}$ closes across the three solvers. This cross-solver closure applies to the same $h_{\text{conv}} = 0$ radiative-only benchmark used in Table 6. The final design’s performance is thus supported by the combination of the window optical metric and the complete energy-balance decomposition rather than by ε_{win} alone, extending PVSD’s physics-verification principle beyond the design loop to post-design verification.

Cross-validation of additional high-performance candidates. To confirm that the framework’s solution diversity (Section 2.4) is physically real and not a pipeline-specific artifact, we cross-check five structurally distinct high-performing final structures—spanning six to eight layers with different material sequences—drawn from the same-seed diversity runs. Each is evaluated both by the official pipeline and by the independently implemented analytic Byrnes TMM (here the back reflector is modeled as a finite 200 nm Ag layer with an air exit, a different modeling choice from the pipeline’s semi-infinite Ag substrate). The two solvers agree to within 0.01 W m^{-2} in P_{cool} and to four decimal places in R_{solar} and ε_{win} for every candidate (Table 8), with recomputed net cooling powers of 137–149 W m^{-2} . Because the transfer-matrix method is analytically exact for planar multilayers, two correct independent implementations are expected to agree to numerical precision; the value of the check is therefore not the size of the (vanishing) discrepancy but that the reported KPIs of these diverse candidates reproduce exactly under an independently coded solver and an

independent back-reflector model, ruling out an implementation-specific bias and supporting the diversity claim with independent physics. This complements the deeper three-way validation of the final champion (pipeline TMM, COMSOL FEM, and analytic TMM).

Table 8. Five diverse high-performing candidates: official pipeline vs. independent analytic TMM (three-angle, fine grid).

Candidate	Layers	P_{cool} (W m^{-2}) Pipeline/Indep.	R_{solar} Pipeline/Indep.	ε_{win} Pipeline/Indep.
C1	7	148.84/148.85	0.9593/0.9593	0.8993/0.8993
C2	7	142.22/142.23	0.9622/0.9622	0.9036/0.9036
C3	8	146.42/146.42	0.9574/0.9574	0.9000/0.9000
C4	8	136.43/136.44	0.9547/0.9547	0.8578/0.8578
C5	6	137.03/137.04	0.9583/0.9583	0.9245/0.9245

2.11. Ablation: What the Physics Verifier and the Neural Proposer Each Contribute

PVSD couples two components—a neural *proposer* (dreaming) and a physics *verifier* (TMM accept/reject). We ablate each to quantify its role, holding the other fixed (Table 9).

The physics verifier catches surrogate pseudo-optima. Over 1313 dreamed proposals across five independent runs, the in-loop TMM rejects 65.6% (861/1313) as not truly improving. More tellingly, among the 511 proposals the surrogate itself *predicted* to be improvements, 25.6% (131/511) are pseudo-optima—predicted better but verified worse by TMM—and are blocked by the verifier (mean surrogate optimism -0.14 W m^{-2} , with individual runs ranging from -1.6 to $+1.1$). Without the verifier, roughly one in four surrogate-endorsed updates would inject a physically unreliable structure; the accept/reject step is therefore not a formality but the mechanism that keeps the search anchored to true physics.

Neural dreaming proposes better than random under equal budget. Replacing the dreaming proposer with random proposals, while keeping the same TMM verifier and the same evaluation budget, lowers the mean best P_{cool} (five seeds) from 126.4 to 93.1 W m^{-2} at a budget of 500 TMM calls ($+33.3 \text{ W m}^{-2}$, $+36\%$ for dreaming). Random search saturates near 93 W m^{-2} by 500 calls, whereas DD+Refine continues to improve with budget ($126.4 \rightarrow 147.3 \text{ W m}^{-2}$ from 500 to 5000 calls). The gradient-guided proposals thus direct the verifier’s limited calls into higher-value regions of the design space. Consistent with the benchmark of Section 2.7, the claim is budget efficiency of the proposal mechanism, not absolute superiority at unlimited budget.

Table 9. Ablation of the two PVSD components (five same-seed runs).

Ablation	Metric	Value
Verifier off (accept all)	surrogate proposals rejected by TMM	65.6% (861/1313)
Verifier off (accept surrogate-claimed)	pseudo-optima caught among claimed	25.6% (131/511)
Proposer $\rightarrow$ random	mean best P_{cool} @500 calls	93.1 vs 126.4 W m^{-2}
Dreaming, budget scaling	mean best P_{cool} @5000 calls	147.3 W m^{-2}

2.12. Physical Mechanism: How the 14-Layer Stack Cools

To interpret *why* the design cools—not merely that it does—we decompose its response layer by layer. Using an independently implemented analytic TMM [21], we compute the fraction of incident power absorbed in each layer, $A_j(\lambda)$; by Kirchhoff’s law each layer’s spectral absorptance equals its emittance, so the same $A_j(\lambda)$ that quantifies parasitic solar absorption also quantifies thermal emission. The decomposition is internally

consistent: summed over layers it reproduces the reported totals, $\sum_j \varepsilon_{\text{win},j} = 0.925$ and $\sum_j P_{\text{rad},j} = 282.2 \text{ W m}^{-2}$ (Figure 8).




Figure 8. Per-layer mechanism of the 14-layer representative design (three-angle, fine grid): (a) bare 200 nm Ag vs. the full stack in R_{solar} and ε_{win} (the stack raises emission, not reflection); (b) per-layer window-emittance contribution $\varepsilon_{\text{win},j}$, colored by material; (c) material-grouped emittance spectra; (d) depth-resolved absorbed-power density at 0.5 and 10 μm . Colors denote materials as in the legend of (b); MgF_2 and ZnSe contribute negligibly (near-zero bars), and their legend entries are retained for completeness.

Solar reflection is provided mainly by the silver substrate. A bare 200 nm Ag mirror already reflects the solar band well ($R_{\text{solar}} = 0.961$) but emits almost nothing in the thermal infrared ($\varepsilon_{\text{win}} = 0.012$, $P_{\text{rad}} = 3.9 \text{ W m}^{-2}$) and therefore *net-heats* under sunlight ($P_{\text{cool}} = -36.1 \text{ W m}^{-2}$). Adding the 14 dielectric layers raises the solar reflectance only marginally ($0.961 \rightarrow 0.970$), acting as a weak antireflection/interference coating that trims residual solar absorption from 38.7 to 30.6 W m^{-2} ; over 80% of that residual absorption occurs in the Ag itself. The dielectric stack is thus *not* a dedicated solar mirror—silver already fills that role.

The dielectric stack’s essential job is to turn the low-emissivity mirror into a broadband emitter. The same layers lift the window emittance from 0.012 to 0.925 ($\sim 79\times$) and the total thermal emission from 3.9 to 282.2 W m^{-2} ($\sim 72\times$), converting a net-heating mirror into a $+172 \text{ W m}^{-2}$ cooler. Depth-resolved profiles are consistent with this division of labor: at 0.5 μm the absorbed-power density peaks inside the silver, whereas at 10 μm it peaks in the dielectric stack (layer 3, Si_3N_4).

Window emission is a distributed, multilayer effect—and not the textbook SiO_2 phonon route. The 8–13 μm emittance is shared across several layers rather than localized in one emitter: layer 3 Si_3N_4 (30.8%), layer 4 TiO_2 (28.6%), layer 10 Si_3N_4 (16.9%), layer 12 Si_3N_4 (8.6%), and layer 8 TiO_2 (7.1%) together supply $\sim 92\%$ of ε_{win} . By material, Si_3N_4 contributes 56% and TiO_2 37%, while SiO_2 —the Si–O phonon emitter most commonly invoked in radiative-cooling coatings—appears here only as ultra-thin (45–65 nm) spacer

layers and contributes just 3.6%. The design instead emits through Si_3N_4 lattice absorption and TiO_2 mid-infrared absorption spread over multiple layers. This also explains the robustness result of Section 2.9: the ultra-thin $\text{SiO}_2/\text{MgF}_2$ boundary layers are the least sensitive precisely because they are thin spacers, not primary emitters.

The non-obviousness of this mechanism is itself evidence for automated search. The resulting stack does not follow the intuitive “metal mirror plus a thick SiO_2 phonon emitter” motif that a designer would be inclined to hand-pick; it distributes emission collaboratively across Si_3N_4 and TiO_2 layers and relegates SiO_2 to a minor role. That a high-performance design sits in this counter-intuitive region of the space is a concrete reason why a guided search such as PVSD is useful: had the optimum coincided with the textbook motif, an automated proposal-and-verification loop would add little. We emphasize that the per-layer figures are an absorbed-power decomposition of a coherently interfering stack, not an assertion that each layer acts in isolation; they identify where absorption and emission are concentrated, not independent sub-devices.

2.13. Angular Robustness

Daytime cooling requires performance to hold as the sun and sky span a range of angles. We evaluate the champion’s directional response from 0 to 80° on the same fine grid (unpolarized, since sunlight and thermal radiation are unpolarized); Figure 9 and Table 10 summarize the result.




Figure 9. Angular response of the 14-layer representative design (unpolarized, fine grid): (a) $R_{\text{solar}}(\theta)$ and $\epsilon_{\text{win}}(\theta)$ from 0 to 80° ; (b) directional net cooling vs. angle (diagnostic; the operating value is the hemispherical 171.6 W m^{-2}).

Two behaviors emerge. The *solar reflectance is essentially angle-independent*: R_{solar} changes by less than 0.005 from normal to 80° ($0.969 \rightarrow 0.966$), so oblique sunlight is rejected as effectively as normal sunlight—a direct consequence of the silver mirror, which reflects the solar band well at all angles (Section 2.12). The *window emittance is stable through the operationally important range and then declines at grazing angles*, as expected for a planar thin-film emitter: ϵ_{win} drops by only 0.012 out to 40° ($0.954 \rightarrow 0.942$) and by 0.055 out to 60° (0.899), before the usual rolloff beyond 60° (0.82 at 70° , 0.60 at 80°).

We report the full curve, including the grazing-angle decline, rather than truncating it: the decline is the physically expected behavior of a planar multilayer and does not affect the hemispherically integrated metric. As an internal consistency check, the three-angle ($0/45/75^\circ$, weights 0.5/0.4/0.1) average of the directional emittance is 0.92, matching the hemispherical $\epsilon_{\text{win}} = 0.9252$ used throughout the paper; the angular analysis and the headline metric are therefore mutually consistent. In summary, the design is strongly angle-insensitive in solar reflection and retains high window emission across the practically relevant angular range.

Table 10. Champion directional optical response vs. incidence angle (unpolarized).

Angle	0°	40°	60°	70°	80°
R_{solar}	0.9692	0.9676	0.9655	0.9648	0.9658
ϵ_{win}	0.9537	0.9421	0.8988	0.8164	0.6001

2.14. The Role of Layer Count and the Choice of a Representative Design

The 14-layer stack studied throughout this paper is a *representative demonstration structure*, not a claimed optimal layer count. To make this explicit we map how net cooling power depends on N . Starting from the 14-layer design (re-verified at 172.5 W m^{-2} on the fine grid), we greedily reduce N by removing the least-important layer and re-polishing, and greedily increase N by inserting the best material/position and re-polishing; each point is re-optimized and re-verified with the official three-angle TMM (Figure 10, Table 11).




Figure 10. Net cooling power vs. layer count N (anchored greedy prune/grow from the representative 14-layer design; fine grid, three-angle). P_{cool} increases with N (to $\approx 180 \text{ W m}^{-2}$ at $N = 20$); the 14-layer design is a representative demonstration case, not the optimal count.

Net cooling power *increases* with N across the range studied—from 143 W m^{-2} at $N = 4$ to about 180 W m^{-2} at $N = 20$ —and does not level off to a single optimum within $N \leq 20$. Below roughly ten layers the gain per added layer is large (each early layer is essential); above $N \approx 12$ the improvement is gradual, of order 1 W m^{-2} per layer, partly reflecting the additional optimization freedom that more thicknesses provide.

We therefore do *not* claim that fourteen layers is the globally optimal count: designs with more layers reach a somewhat higher simulated P_{cool} (for example $\approx 180 \text{ W m}^{-2}$ at $N = 20$). Fourteen layers was chosen as a practical balance—it already lies in the high-efficiency part of the curve, while further layers yield diminishing returns at the price of more deposition steps and tighter tolerance stack-up. This positioning is consistent with the rest of the paper, which makes no global-optimality claim for the specific structure (Sections 2.8 and 3); the reported $\approx 172 \text{ W m}^{-2}$ is the verified performance of *this* representative 14-layer design, not an upper bound over all layer counts. Read the other way, the steady improvement of P_{cool} with N is itself evidence that PVSD finds strong structures at *every* layer count we tried, not only at $N = 14$ —the framework is not tuned to one fortunate size. Extending to still larger N is left as future work, bounded in practice by the number of layers a deposition process can reliably realize.

Table 11. Net cooling power vs. layer count N (anchored greedy prune/grow from the representative 14-layer design; fine grid, three-angle TMM).

N	4	8	12	14	16	20
$P_{\text{cool}} (\text{W m}^{-2})$	143.1	168.2	173.1	172.5	174.4	179.7

3. Discussion

PVSD-TMM changes the relationship between neural models and physical solvers in optical inverse design. Conventional surrogate inverse design uses the network as a performance evaluator, which is easily misled by model error; PVSD casts the network as a proposal generator and the solver as the verifier, so final performance is determined entirely by TMM and the model only improves candidate-generation efficiency. Multilayers are a useful testbed because the structure is naturally a token sequence with one physical layer per token, TMM is fast enough to embed in every accept/reject step, and attention weights and optimization paths map directly back to materials and thicknesses. The framework extends to RCWA/FDTD/FEM or experimental-loop solvers by swapping the verifier. The positioning is efficiency at low-to-moderate evaluation budgets, not absolute replacement of all optimizers at unlimited budget; for expensive 3D electromagnetic or experimental loops, the solver-in-the-loop direction is more valuable.

The reported $\approx 172 \text{ W m}^{-2}$ should be read as a cross-solver-verified simulation in an idealized radiative-only boundary condition ($h_{\text{conv}} = 0$)—COMSOL FEM (re-checked on a licensed installation) and an independent analytic TMM—not as a measurement or as a direct outdoor cooling-power prediction. In real devices, finite conduction/convection with the mounting, air enclosure, wind shield, and surrounding air would add $P_{\text{conv}} = h_{\text{conv}} (T_{\text{surface}} - T_{\text{ambient}})$ or an equivalent nonradiative heat-transfer term, and the operating temperature or useful cooling power must be recomputed for the chosen system-level h_{conv} .

Figure 11 situates PVSD-TMM among representative machine-learning and optimization approaches to radiative-cooling multilayer films. Because these works differ in material set, layer count, spectral and atmospheric models, and solver, their simulation-reported values are not directly comparable; we therefore draw no superiority claim and present the figure only to locate our method within this family.

Non-Ideal Convective Boundary and Sub-Ambient Operation


The $\approx 172 \text{ W m}^{-2}$ result assumes the idealized radiative-only boundary $h_{\text{conv}} = 0$. To connect it to realistic operation we add a nonradiative heat-transfer term and sweep h_{conv} from 0 to $20 \text{ W m}^{-2}\text{K}^{-1}$, using the champion's own hemispherical emittance and the standard sub-ambient balance

$$P_{\text{net}}(T_s) = P_{\text{rad}}(T_s) - P_{\text{atm}}(T_a) - P_{\text{solar}} - h_{\text{conv}} (T_a - T_s), \quad (1)$$

where the nonradiative term $h_{\text{conv}}(T_a - T_s)$ is a parasitic heat gain when the surface is below ambient (Figure 12).

Two points follow. First, the net cooling power at ambient temperature ($T_s = T_a$) is independent of h_{conv} and equals the reported $\approx 172 \text{ W m}^{-2}$; the convective term vanishes at $\Delta T = 0$, so the headline value does not depend on the choice $h_{\text{conv}} = 0$. Second, what h_{conv} controls is the achievable sub-ambient temperature: solving $P_{\text{net}} = 0$ gives a stagnation $\Delta T = T_a - T_s$ of $\approx 55 \text{ K}$ in the idealized radiative-only limit ($h_{\text{conv}} = 0$)—an upper bound that is not physically realizable, reported only as a limiting case—falling to $\approx 20 \text{ K}$ at $h_{\text{conv}} = 5$, $\approx 12 \text{ K}$ at 10, and $\approx 7 \text{ K}$ at $20 \text{ W m}^{-2}\text{K}^{-1}$.

Realistic passive setups with wind shielding correspond to $h_{\text{conv}} \approx 6\text{--}12 \text{ W m}^{-2}\text{K}^{-1}$, i.e., $\Delta T \approx 10\text{--}15 \text{ K}$, in the range of typical measured daytime sub-ambient drops. A full building- or apparatus-scale coupling (for example an EnergyPlus model with a specific enclosure, wind shield, and weather file) is beyond the scope of this optical-design study and is left as future work; the physical h_{conv} sweep here is the appropriate, self-consistent way to place the idealized 172 W m^{-2} in operational context and to make explicit that it is a radiative-power upper bound, not an open-air temperature prediction.



Heterogeneous-condition landscape, not a controlled benchmark. Shape: simulation = hollow, experiment = filled, star = this work, gray triangle = reference. Color: broadband = warm, selective = cool. ALL selective coolers (So & Rho 84.8 sim / 70.6 exp, LiF 60 exp, Raman 40.1 exp, Rephaeli 105 theory) lie below the 139.7 selective ceiling — the selective route is capped at ~ 140 by design, not inferiority. Our broadband design (172) enters the broadband-only regime (140–224), at 77% of the 224 broadband ceiling. Selective emitters optimize for sub-ambient temperature reduction (low T_{eq}) at the cost of lower at-ambient power; broadband emitters (this work) optimize at-ambient net cooling power. This figure compares the at-ambient net-power metric; the two represent different application objectives, not a performance hierarchy. Notes: Raman 2014 includes $h_c = 6.9 \text{ W/m}^2\text{K}$ convective loss (30° tilt), hence lower; Rephaeli 2013 is theoretical (RCWA, $h=0$) but its emitter is a 2D photonic crystal (not pure 1D) $\rightarrow$ reference only. Dry atmosphere (Mauna Kea PWV 1.6 mm) sets these ceilings; standard-atm single-faced limit (Gan 2020) ≈ 157 .

Figure 11. Positioning of PVSD-TMM within the machine-learning-and-optimization family of radiative-cooling multilayer designs. Sources: Raman et al. [1]; Rephaeli et al. [3]; So et al. [6]; Yang et al. [22] (LiF); and Zhou et al. [23] (“Gan 2020”). Heterogeneous methods and conditions preclude direct comparison; no superiority is claimed.




Figure 12. Non-ideal convective boundary for the 14-layer representative design: (a) equilibrium sub-ambient ΔT vs. h_{conv} (the $h = 0$, 55 K value is an idealized upper bound); (b) net cooling power held 5 K below ambient. The at-ambient power ($\Delta T = 0$) is 171.6 W m^{-2} , independent of h_{conv} .

Broadband versus spectrally selective emission. The reported $\approx 172 \text{ W m}^{-2}$ is a *broadband* net cooling power evaluated at $T_{\text{surface}} = T_{\text{ambient}}$: the design emits strongly across the whole thermal infrared, not only in the 8–13 μm window (see the per-layer decomposition of Section 2.12, where HfO_2 also emits beyond 13 μm). This is the natural figure of merit when the goal is to reject the most heat at ambient temperature. It is a different objective from that of spectrally *selective* emitters, which deliberately suppress emission outside the atmospheric window so that, once the surface cools below ambient, they absorb less returning atmospheric radiation and can reach a lower stagnation temperature. The 172 W m^{-2} value should therefore not be read as uniformly “better” than selective designs: a broadband emitter maximizes net cooling power near ambient, whereas a selective emitter trades peak power for a larger sub-ambient temperature drop. This is consistent with our equilibrium-temperature analysis above: the broadband design delivers high power at $\Delta T = 0$, but its achievable sub-ambient ΔT shrinks once a realistic nonradiative heat-transfer coefficient is included. Which objective is preferable is application-dependent—net heat rejection at ambient (e.g., a surface held near air temperature) versus maximum sub-ambient cooling (e.g., a passively chilled enclosure)—and PVSD can target either by changing the KPI it optimizes, since the framework places no restriction on the objective the solver scores.

The work most closely related to ours is Zhou et al. [16], who also use deep learning to inverse-design multilayer thin films for radiative cooling. To make our incremental contribution explicit, Table 12 contrasts the two approaches along the dimensions of proposal mechanism, acceptance criterion, solver role, interpretability, validation depth, and reported outcomes. The central distinction is where physics enters: in [16] a forward neural surrogate trained on an FDTD dataset predicts spectra and the design is selected from those predictions, whereas PVSD keeps a physical solver *inside* the optimization loop—every accepted update is verified by three-angle TMM—and additionally audits the final design with independent solvers (COMSOL FEM and a separately implemented analytic TMM). The two contributions are complementary rather than competing; our added value is the per-candidate physics verification, the population-scale and cross-solver evidence, and the token attention interpretability, not another single high-performing design.

Table 12. Systematic comparison of PVSD-TMM (this work) with Ref. [16]. “Not reported” denotes a dimension the cited paper does not explicitly address.

Dimension	PVSD-TMM (This Work)	Zhou et al. [16]
Proposal mechanism	Frozen forward surrogate optimized by input-gradient “dreaming” (neural-adjoint-style)	Forward MLP surrogate predicts spectra; selection of the optimum from predictions
Acceptance criterion	Every accepted update verified in-loop by three-angle TMM (propose $\rightarrow$ solve $\rightarrow$ accept/reject)	Surrogate-prediction-driven; a per-candidate in-loop solver check is not reported
Solver role	TMM in the loop; plus post-design COMSOL FEM and independent analytic TMM audit	FDTD used to generate the offline training dataset
Target application	1D multilayer daytime radiative cooler (11-material pool, 14 layers)	1D multilayer radiative cooler ($\text{W/SiO}_2/\text{Si}_3\text{N}_4/\text{Ge}$, ~ 6 layers)
Interpretability	Scalar token attention analysis + per-layer absorption mechanism decomposition	Physics-based field/power-dissipation analysis (no ML-interpretability method)
Validation depth	1000-seed population migration; 5 extra candidates cross-checked; champion validated by three solvers	8000-sample hold-out ($R^2 \approx 0.999$); independent cross-solver validation of the final design not reported
Reported outcome	Full energy balance: $P_{\text{cool}} \approx 172$, $R_{\text{solar}} \approx 0.97$, $\varepsilon_{\text{win}} = 0.925$ (idealized $h_{\text{conv}} = 0$)	$\varepsilon_{\text{win}} = 0.855$ (8–13 μm), 3–8 μm band ≈ 0.23 ; net cooling power/ ΔT not reported

4. Methods

4.1. Performance Metrics and TMM Evaluation

Spectra are computed by the official TMM pipeline at $T_{\text{surface}} = T_{\text{ambient}} = 300 \text{ K}$, $h_{\text{conv}} = 0$, Ag substrate, and three-angle integration. The choice $h_{\text{conv}} = 0$ defines a radiative-only benchmark used to compare optical designs and to estimate an upper-bound ambient-temperature cooling power. It does not model a particular outdoor apparatus. Net cooling power is $P_{\text{cool}} = P_{\text{rad}} - P_{\text{atm}} - P_{\text{solar}} - P_{\text{conv}}$, where P_{rad} and P_{atm} are evaluated from the full thermal-infrared spectrum used in the energy balance. For practical devices, the same spectra should be combined with a finite nonradiative heat-transfer model, e.g., $P_{\text{conv}} = h_{\text{conv}}(T_{\text{surface}} - T_{\text{ambient}})$ or a measured system heat-transfer coefficient, before predicting the operating temperature or delivered cooling power. R_{solar} is the AM1.5G-weighted normal reflectance; ε_{win} is the blackbody- and transmittance-weighted hemispherical emittance over the primary 8–13 μm atmospheric window and is reported as a window-selectivity metric, not as a substitute for P_{cool} . The AM1.5G spectrum is ASTM G-173-03 (SMARTS2, global-tilt, $\approx 1000 \text{ W m}^{-2}$) [22]. The atmospheric transmittance is generated by ATRAN [23] (Gemini Observatory; Mauna Kea, PWV $\approx 1.6 \text{ mm}$); it enters only P_{atm} , with a three-way-confirmed effect of $\leq 0.13 \text{ W m}^{-2}$.

4.2. Dataset, Predictor, Dreaming, Refinement

The dataset has 100,000 TMM structure–performance samples (80k/10k/10k), each encoded as an alternating material/thickness sequence with PAD tokens. The predictor uses embedding/noisy-one-hot projection, an LSTM, scalar attention pooling, and an MLP head, jointly outputting $P_{\text{cool}}/R_{\text{solar}}/\varepsilon_{\text{win}}$.

Token representation and differentiable relaxation. Each structure is encoded as an alternating sequence of material and thickness tokens on a fixed Ag substrate, $[m_1, t_1, m_2, t_2, \dots, m_N, t_N]$, zero-padded to a maximum of $N_{\text{max}} = 20$ layers. Material tokens are drawn from the 11-element dielectric pool; thickness tokens are quantized to 12 discrete levels, $\{10, 20, 50, 100, 150, 200, 300, 500, 750, 1000, 1500, 2000\} \text{ nm}$, giving a token vocabulary of $11 + 12 + 1 = 24$ classes (the last is a PAD token). During dreaming the discrete input is relaxed to a differentiable form: every token position carries a class-probability vector—a noisy one-hot distribution over the 24 classes—and input-gradient ascent updates these probabilities, while a projection step keeps each position on the valid probability simplex and suppresses the PAD columns of occupied positions. For TMM evaluation the relaxed distribution is decoded to a concrete structure by per-position argmax . The 12-level thickness grid is used *only* for this discrete proposal stage. After the discrete design is fixed, its material sequence is frozen and the N thicknesses are optimized as *continuous* variables against the official three-angle TMM (Section 2.8), which moves them off the 12-level grid, before a final rounding to a 5 nm manufacturable grid. This two-stage scheme—coarse differentiable proposal on 12 bins, then continuous physics-in-the-loop refinement—is why the final thicknesses (Table 5) are not multiples of the coarse dreaming grid.

Dreaming freezes the weights and optimizes only the input token distribution: seed $\rightarrow$ tokenize $\rightarrow$ differentiable noisy input $\rightarrow$ frozen forward pass $\rightarrow$ input-gradient update $\rightarrow$ decode $\rightarrow$ TMM evaluation $\rightarrow$ accept only if TMM improves. After the final discrete design’s material sequence is fixed, only continuous thicknesses are optimized against the official three-angle TMM, then rounded to a 5 nm grid and re-verified. Robustness uses rounding, Monte Carlo independent perturbation, and layer-wise sensitivity, all re-evaluated by the same pipeline. AI-assisted tools were used for code implementation, manuscript formatting, and language editing (see Acknowledgments); all numerical results were computed and independently cross-verified by the authors’ physics solvers (the

in-house TMM, an independent analytic Byrnes TMM, and COMSOL Multiphysics), not by AI.

4.3. Cross-Solver Validation Protocol

Two solver paths check the final design: (i) a COMSOL 6.3 Wave Optics 2D periodic FEM model with Floquet boundaries and a PeriodicStructure port, meshed to 5 nm at the thinnest metal layer (overall $h_{\max} = 50$ nm) [24]; (ii) a from-scratch analytic Byrnes TMM [21]. Both use the same frozen (n, k) data [25,26]. For the validation audit, target metrics were fixed before re-solving, and no model parameters were tuned toward the original KPIs.

5. Conclusions

We presented Physics-Verified Spectral Dreaming (PVSD), which couples a differentiable surrogate's directional search with solver-based acceptance, turning inverse design from unconstrained surrogate search into a physics-checked proposal-and-verification workflow. For multilayer radiative coolers, PVSD-TMM is effective in forward prediction, single-structure optimization, same-seed diversity, 1000-structure population migration, attention interpretability, and a fair benchmark. Via Super-DD+Refine, continuous-thickness polishing, and 5 nm rounding, we obtained a 14-layer final design and cross-validated it using COMSOL FEM and analytic TMM to $P_{\text{cool}} \approx 172 \text{ W m}^{-2}$ for the radiative-only case $h_{\text{conv}} = 0$, $R_{\text{solar}} \approx 0.970$, and $\varepsilon_{\text{win}} = 0.9252$. The window metric is confirmed three ways, while the high simulated cooling power is assessed from the full-spectrum energy balance. Finite nonradiative heat exchange must be added for apparatus- or application-specific operating predictions. This provides a generalizable workflow combining neural proposal generation, fast physical solvers, and interpretable optical design, with post-design cross-solver validation supporting the final structure's reliability.

Author Contributions: Conceptualization, J.W. and X.L.; methodology, J.W.; software, J.W.; validation, J.W.; formal analysis, J.W.; investigation, J.W.; data curation, J.W.; writing—original draft preparation, J.W.; writing—review and editing, X.L.; visualization, J.W.; supervision, X.L.; funding acquisition, X.L. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by the Scientific Research Starting Foundation of Chengdu University of Technology (Grant No. KYQD2025-10944).

Data Availability Statement: The data underlying the results presented in this paper are not publicly available at this time but may be obtained from the authors upon reasonable request.

Acknowledgments: The authors thank colleagues for discussions on cross-solver validation. During the preparation of this manuscript, the authors used Claude (Anthropic) for assistance with code implementation, LaTeX formatting, and language editing. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the content of the publication.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Raman, A.P.; Anoma, M.A.; Zhu, L.; Rephaeli, E.; Fan, S. Passive radiative cooling below ambient air temperature under direct sunlight. *Nature* **2014**, *515*, 540–544. [CrossRef] [PubMed]
2. Zhao, D.; Aili, A.; Zhai, Y.; Xu, S.; Tan, G.; Yin, X.; Yang, R. Radiative sky cooling: Fundamental principles, materials, and applications. *Appl. Phys. Rev.* **2019**, *6*, 021306. [CrossRef]
3. Rephaeli, E.; Raman, A.; Fan, S. Ultrabroadband photonic structures to achieve high-performance daytime radiative cooling. *Nano Lett.* **2013**, *13*, 1457–1461. [CrossRef] [PubMed]
4. Zhai, Y.; Ma, Y.; David, S.N.; Zhao, D.; Lou, R.; Tan, G.; Yang, R.; Yin, X. Scalable-manufactured randomized glass-polymer hybrid metamaterial for daytime radiative cooling. *Science* **2017**, *355*, 1062–1066. [CrossRef] [PubMed]

5. Li, T.; Zhai, Y.; He, S.; Gan, W.; Wei, Z.; Heidarinejad, M.; Dalgo, D.; Mi, R.; Zhao, X.; Song, J.; et al. A radiative cooling structural material. *Science* **2019**, *364*, 760–763. [CrossRef] [PubMed]
6. So, S.; Yang, Y.; Son, S.; Lee, D.; Chae, D.; Lee, H.; Rho, J. Highly suppressed solar absorption in a daytime radiative cooler designed by genetic algorithm. *Nanophotonics* **2022**, *11*, 2107–2115. [CrossRef] [PubMed]
7. You, P.; Li, X.; Huang, Y.; Ma, X.; Pu, M.; Guo, Y.; Luo, X. High-Performance Multilayer Radiative Cooling Films Designed with Flexible Hybrid Optimization Strategy. *Materials* **2020**, *13*, 2885. [CrossRef] [PubMed]
8. Li, S.; An, M.; Zheng, Z.; Gou, Y.; Lian, W.; Yu, W.; Zhang, P. Daytime radiative cooling multilayer films designed by a machine learning method and genetic algorithm. *Appl. Opt.* **2023**, *62*, 4359–4369. [CrossRef] [PubMed]
9. Wiecha, P.R.; Arbouet, A.; Girard, C.; Muskens, O.L. Deep learning in nano-photonics: Inverse design and beyond. *Photonics Res.* **2021**, *9*, B182–B200. [CrossRef]
10. Liu, D.; Tan, Y.; Khoram, E.; Yu, Z. Training deep neural networks for the inverse design of nanophotonic structures. *ACS Photonics* **2018**, *5*, 1365–1369. [CrossRef] [PubMed]
11. Zhang, Y.; Chen, Y.; Wang, T.; Zhu, Q.; Gu, M. Ultrahigh performance passive radiative cooling by hybrid polar dielectric metasurface thermal emitters. *Opto-Electron. Adv.* **2024**, *7*, 230194. [CrossRef]
12. Kim, S.; Jung, S.; Bobbitt, A.; Lee, E.; Luo, T. Wide-angle spectral filter for energy-saving windows designed by quantum annealing-enhanced active learning. *Cell Rep. Phys. Sci.* **2024**, *5*, 101847. [CrossRef]
13. Mordvintsev, A.; Olah, C.; Tyka, M. Inceptionism: Going Deeper into Neural Networks. Google Research Blog. 2015. Available online: <https://research.google/blog/inceptionism-going-deeper-into-neural-networks/> (accessed on 8 June 2026).
14. Peurifoy, J.; Shen, Y.; Jing, L.; Yang, Y.; Cano-Renteria, F.; DeLacy, B.G.; Joannopoulos, J.D.; Tegmark, M.; Soljačić, M. Nanophotonic particle simulation and inverse design using artificial neural networks. *Sci. Adv.* **2018**, *4*, eaar4206. [CrossRef] [PubMed]
15. Kim, S.; Kim, J. Interpretable inverse design of optical multilayer thin films based on extended neural adjoint and regression activation mapping. *arXiv* **2025**, arXiv:2507.18644. [CrossRef]
16. Zhou, G.; Zhu, J.; Zhou, F.; Zhang, W.; Yi, Z.; Cheng, S. Deep-learning-based multilayer nano-film structure reverse design for efficient passive radiative cooling. *Chin. Opt. Lett.* **2026**, *24*, 042201. [CrossRef]
17. Erhan, D.; Bengio, Y.; Courville, A.; Vincent, P. *Visualizing Higher-Layer Features of a Deep Network*; Technical Report 1341; University of Montréal: Montreal, QC, Canada, 2009.
18. Simonyan, K.; Vedaldi, A.; Zisserman, A. Deep inside convolutional networks: Visualising image classification models and saliency maps. In Proceedings of the International Conference on Learning Representations (ICLR) Workshop, Banff, AB, Canada, 14–16 April 2014. [CrossRef]
19. Shen, C.; Krenn, M.; Eppel, S.; Aspuru-Guzik, A. Deep molecular dreaming: Inverse machine learning for de-novo molecular design and interpretability with surjective representations. *Mach. Learn. Sci. Technol.* **2021**, *2*, 03LT02. [CrossRef]
20. Cleeton, C.; Sarkisov, L. Inverse design of metal-organic frameworks using deep dreaming approaches. *Nat. Commun.* **2025**, *in press*. [CrossRef] [PubMed]
21. Byrnes, S.J. Multilayer optical calculations. *arXiv* **2016**, arXiv:1603.02720. [CrossRef]
22. ASTM G173-03; Standard Tables for Reference Solar Spectral Irradiances: Direct Normal and Hemispherical on 37° Tilted Surface. ASTM International: West Conshohocken, PA, USA, 2020. [CrossRef]
23. Lord, S.D. *A New Software Tool for Computing Earth's Atmospheric Transmission of Near- and Far-Infrared Radiation*; Technical Report NASA Technical Memorandum 103957; NASA Ames Research Center: Mountain View, CA, USA, 1992.
24. COMSOL Multiphysics, v.6.3. Wave Optics Module. COMSOL AB: Stockholm, Sweden, 2024. Available online: <https://www.comsol.com> (accessed on 8 June 2026).
25. Kischkat, J.; Peters, S.; Gruska, B.; Semtsiv, M.; Chashnikova, M.; Klinkmüller, M.; Fedosenko, O.; Machulik, S.; Aleksandrova, A.; Monastyrskiy, G.; et al. Mid-infrared optical properties of thin films of aluminum oxide, titanium dioxide, silicon dioxide, aluminum nitride, and silicon nitride. *Appl. Opt.* **2012**, *51*, 6789–6798. [CrossRef] [PubMed]
26. Malitson, I.H. Interspecimen comparison of the refractive index of fused silica. *J. Opt. Soc. Am.* **1965**, *55*, 1205–1209. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.