EFSA: PHYSICALLY INTERPRETABLE FACET MODELING FOR HIGH-FREQUENCY RCS PREDICTION

Anonymous authors
Paper under double-blind review

Learning physically interpretable and computationally efficient representations for high-frequency radar cross section (RCS) prediction is a critical yet unresolved challenge. While mesh-based Graph Neural Networks (GNNs) provide an excellent paradigm for understanding geometric and electromagnetic interactions, they are inherently bottlenecked by rigid sub-wavelength constraints, where shrinking wavelengths lead to catastrophic graph size expansion. To fundamentally decouple neural representations from these strict physical size requirements, we introduce the Equivalent Facet Scattering Amplitude (EFSA), which models the RCS contribution of facets at arbitrary scales. EFSA combines a Physical Optics (PO) prior with a learnable complex correction factor. Specifically, the PO prior absorbs the orders-of-magnitude variations in RCS contributions across differently scaled facets within the same target, while the correction factor captures the intra-facet coherent interference of facets significantly larger than the incident wavelength. Experiments on real-world targets demonstrate that our framework enables physically interpretable RCS prediction in high-frequency regimes while outperforming existing methods in predictive accuracy.

1 INTRODUCTION

Radar cross section (RCS) is a fundamental parameter characterizing an object's macroscopic electromagnetic scattering, playing a crucial role in diverse applications such as target recognition (Safa et al., 2023; Wang et al., 2024), detection range estimation (Li & Chen, 2022), and metasurface design (Xu et al., 2022; Fang et al., 2024). To overcome the prohibitive computational costs of traditional numerical solvers, deep learning has emerged as a highly efficient alternative (Yao et al., 2019b; Zhang et al., 2023; Yao et al., 2019a; 2020; Qi et al., 2020; Guo et al., 2021; Ma et al., 2020; Li et al., 2020). Current neural approaches generally fall into two parallel paradigms: direct end-to-end RCS prediction (Kong et al., 2024; Han et al., 2025; Tang et al., 2026b;a) and the prediction of intermediate physical representations (Shan et al., 2024; Zhang et al., 2024; Shan et al., 2026). While direct prediction offers a straightforward pipeline, the latter paradigm is particularly attractive because it preserves intermediate physical quantities and provides enhanced physical interpretability. Mesh-based Graph Neural Networks (GNNs) provide a natural framework for this paradigm. By mapping graph nodes to irregular 3D faces, they effectively capture localized scattering quantities that aggregate to form the object's RCS.

Despite their representational advantages, existing mesh-based GNNs are fundamentally hindered by a severe scalability bottleneck. Existing mesh-based GNNs typically construct graph nodes on electrically small facets whose edge lengths are sufficiently smaller than the incident wavelength (Shan et al., 2024; Bakirtzis et al., 2024; Jiang et al., 2026; Shan et al., 2026). Under this rigid sub-wavelength constraint, localized intra-facet coherent interference can be neglected, allowing the network to simply predict the equivalent current response at the facet's geometric center and aggregate the total RCS via coherent superposition. However, as the incident frequency increases, preserving this sub-wavelength constraint triggers an explosive growth in graph size. The resulting graph-size growth makes existing GNN paradigms computationally expensive and memory-intensive in high-frequency electromagnetic regimes, severely limiting their practical applicability.

To overcome this limitation, we propose a framework that decouples graph nodes from electrically small facets, allowing them to represent facets at arbitrary scales and thereby alleviating the high-frequency graph-size bottleneck. To enable this scale decoupling, we introduce a novel intermediate physical representation termed the Equivalent Facet Scattering Amplitude (EFSA). As illustrated in Figure 1, EFSA represents the complex scattering contribution of a facet at an arbitrary physical scale by decomposing it into two components: an analytical Physical Optics (PO) prior and a learnable complex correction factor. The analytical prior captures the dominant scale-dependent scatter-

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

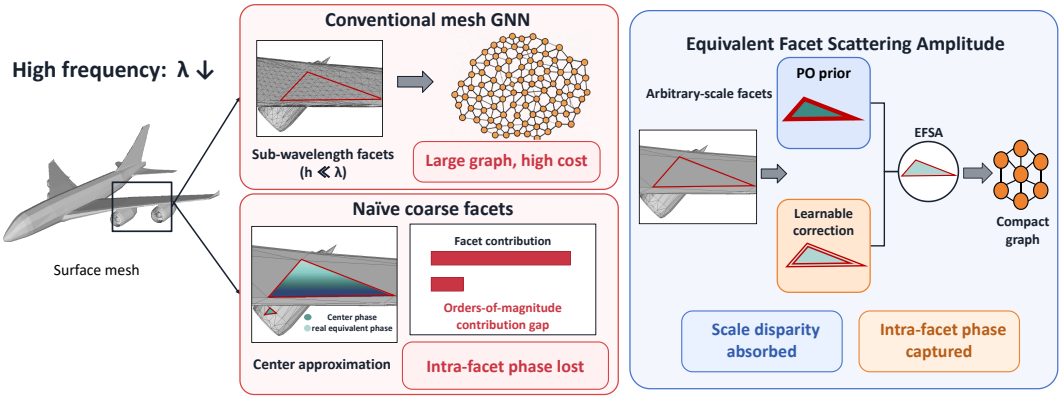

Figure 1: Concept of arbitrary-scale facet modeling for high-frequency RCS prediction. Existing mesh-based GNNs assign nodes to strict sub-wavelength facets. Maintaining this constraint at high frequencies leads to an explosive growth in graph size, whereas simply coarsening the mesh causes severe scale disparities and phase breakdown. To decouple graph resolution from the incident wavelength, we propose the EFSA. EFSA represents the scattering contribution of arbitrary-scale facets by disentangling it into an analytical Physical Optics (PO) prior which absorbs deterministic scale variations—and a flexible neural correction factor that models complex intra-facet coherent interference.

ing contribution, thereby absorbing orders-of-magnitude variations caused by highly non-uniform facet sizes and stabilizing the learning objective. The learnable correction then accounts for the residual magnitude and phase variations induced by intra-facet coherent interference on electrically large facets.

To summarize, our main contributions are as follows:

- **A Novel Intermediate Physical Representation (EFSA):** We propose the Equivalent Facet Scattering Amplitude (EFSA) to explicitly model the complex far-field contribution of facets at arbitrary physical scales.
- **Physically Interpretable High-Frequency RCS Prediction:** By using EFSA to decouple the physical representation from strict sub-wavelength discretization, our framework enables physically interpretable RCS prediction in high-frequency regimes while operating directly on targets with arbitrary mesh discretizations.
- **Effectiveness on Real-World Targets:** We validate the proposed framework through comprehensive experiments on real-world targets. Empirical results demonstrate that our approach achieves high predictive accuracy compared with existing baselines while avoiding the graph-size explosion induced by conventional sub-wavelength mesh discretization.

2 METHOD

In this section, we first formulate the EFSA, sequentially detailing its analytical PO prior and the learnable complex correction factor. We then present the overarching mesh-based GNN architecture, followed by a detailed description of the key parameters used to drive the network.

2.1 EQUIVALENT FACET SCATTERING AMPLITUDE

The core of our scale-decoupled framework lies in the definition of a novel intermediate physical representation, the Equivalent Facet Scattering Amplitude (EFSA). Unlike conventional current-based nodal representations that typically rely on electrically small discretizations, EFSA is explicitly designed to quantify the macroscopic electromagnetic scattering contribution of a physical facet at any arbitrary scale. Figure 2 shows how our framework operates on an arbitrary mesh. Conceptually, the EFSA representation for a given facet can be abstracted as the combination of an analytical baseline

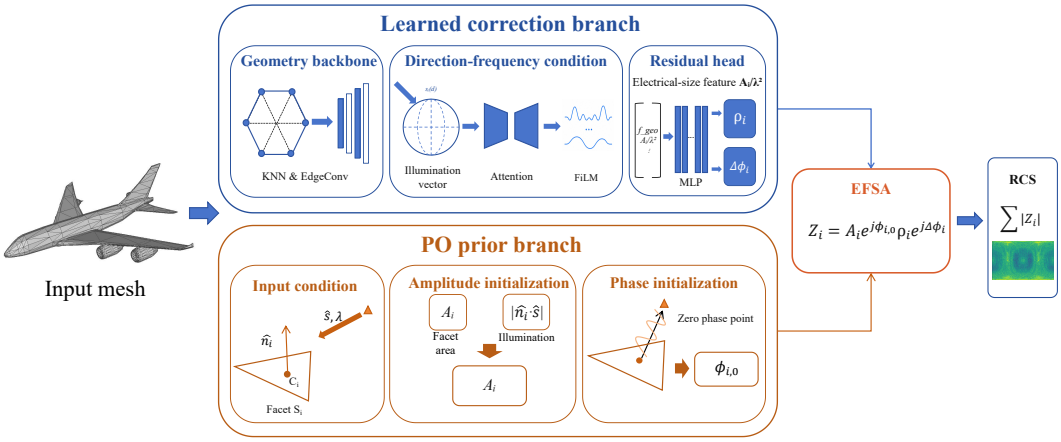

Figure 2: Overview of the proposed framework. Given an arbitrary mesh, the framework estimates the EFSA for each individual facet through two complementary pathways. First, an analytical PO prior computes a baseline amplitude and phase derived explicitly from the facet's geometry and incident conditions. Concurrently, the learnable component employs GNN, attention, FiLM and a MLP decoder to predict complex correction factors. These neural corrections modulate the physical prior to form the accurate EFSA, and the final RCS is obtained by coherently aggregating the EFSAs across all facets.

and a neural residual. In the following subsections, we detail how this representation is constructed: we first establish a deterministic physical foundation that normalizes scale disparities in Section 2.1.1, and subsequently augment it with a learnable complex correction factor to capture intra-facet coherent interference in Section 2.1.2.

2.1.1 Analytical Physical Optics Prior

Given a target surface discretized into N planar triangular facets $\{\mathbb{S}_i\}_{i=1}^{N}$, let $\boldsymbol{r}_i$ denote the geometric center of facet $\mathbb{S}_i$, and $\boldsymbol{J}_i(\boldsymbol{r}')$ denote the surface current density at a source point $\boldsymbol{r}' \in \mathbb{S}_i$. For a monostatic direction $\hat{\boldsymbol{s}}$, the incident wave propagates along $-\hat{\boldsymbol{s}}$, and the scattered field is observed in the far field along $\hat{\boldsymbol{s}}$. Following the standard surface-current radiation formulation used in computational electromagnetics (Harrington, 1968; Rao et al., 1982), its RCS σ in direction $\hat{\boldsymbol{s}}$ is calculated as:

$$\sigma = \frac{k^2\eta^2}{4\pi E_0^2} \left\| \sum_{i=1}^{N} \iint_{\mathbb{S}_i} \left[\hat{\boldsymbol{s}} \times (\hat{\boldsymbol{s}} \times \boldsymbol{J}_i(\boldsymbol{r}'))\right] e^{j\boldsymbol{k}_s \cdot \boldsymbol{r}'} \, \mathrm{d}S' \right\|_2^2, \tag{1}$$

where η is the free-space wave impedance, $E_0 > 0$ is the incident electric-field amplitude, $\boldsymbol{k}_s = k\hat{\boldsymbol{s}}$ is the scattering wave vector, and $k = 2\pi/\lambda$ is the wavenumber for wavelength λ. The dependence of $\boldsymbol{J}_i$ on the incident direction, frequency, and polarization is omitted for notational simplicity.

Existing neural approaches typically formulate the learning objective as predicting the exact surface current $\boldsymbol{J}_i(\boldsymbol{r}')$ directly (Cao et al., 2025; Shan et al., 2026). Learning such an exact current distribution requires dense element-wise supervision, for which reference currents are typically computed on finely discretized meshes using full-wave numerical solvers. However, because these numerical solvers mandate sufficiently fine discretizations to ensure computational convergence (Harrington, 1968; Rao et al., 1982), networks trained on these discrete currents inevitably inherit a rigid sub-wavelength scale restriction.

To fundamentally decouple our predictive representation from electrically small facets, we abandon numerical solver-derived currents and instead adopt the analytical surface current prior derived from Physical Optics (PO). As a widely used high-frequency asymptotic technique, PO provides a concise, scale-independent approximation of surface currents while maintaining high simulation fidelity (Perez & Catedra, 1994; Knott et al., 2004). The analytical expression for the PO current $\boldsymbol{J}_i^{\mathrm{PO}}$ is

given by:

$$\boldsymbol{J}^{\mathrm{PO}}(\boldsymbol{r}') = \begin{cases} 2\hat{\boldsymbol{n}}(\boldsymbol{r}') \times \boldsymbol{H}^{\mathrm{inc}}(\boldsymbol{r}'), & \boldsymbol{r}' \in \mathbb{S}_{\mathrm{ill}}, \\ 0, & \boldsymbol{r}' \in \mathbb{S}_{\mathrm{back}}, \end{cases} \tag{2}$$

where $\boldsymbol{H}^{\mathrm{inc}}(\boldsymbol{r}') = -\frac{E_0}{\eta}\,(\hat{\boldsymbol{s}} \times \hat{\boldsymbol{e}})\, e^{j\boldsymbol{k}_s \cdot \boldsymbol{r}'}$ and $\hat{\boldsymbol{e}}$ denotes the unit polarization vector of the incident electric field.

As indicated by Eq. 2, for a planar triangular facet under a fixed incident condition, the PO surface current has a spatially constant amplitude, while its variation within the facet arises from the propagation phase. According to Eq. 1, the scattering contribution of $\mathbb{S}_i$ is then obtained by coherently integrating $\boldsymbol{J}_i^{\mathrm{PO}}(\boldsymbol{r}')$ over the facet:

$$\mathbf{Z}_i = \iint_{\mathbb{S}_i} \left[\hat{\boldsymbol{s}} \times \left(\hat{\boldsymbol{s}} \times \boldsymbol{J}_i^{\mathrm{PO}}(\boldsymbol{r}')\right)\right] e^{j\boldsymbol{k}_s \cdot \boldsymbol{r}'}\, \mathrm{d}S'. \tag{3}$$

We consider horizontal co-polarization in this work and define the scalar EFSA as $Z_i = \hat{\boldsymbol{e}}_r^T \mathbf{Z}_i$, where $\hat{\boldsymbol{e}}_r$ is the receive-polarization unit vector. Since the normal and polarization-dependent vector terms are spatially constant over a planar facet, they can be factored out of the integral after polarization projection; the detailed derivation is provided in the Appendix A.1. The common factors involving E_0 and η are omitted in the following normalized formulation because they cancel with the corresponding RCS normalization.

For traditional sub-wavelength facets, evaluating this integral is straightforward as the source point $\boldsymbol{r}'$ can be safely approximated by the facet's geometric center, since internal phase differences are negligible. However, for electrically large facets, this center-point approximation fails and the intra-facet phase variation must be explicitly considered. Consequently, the learning target evolves from estimating point-wise currents to predicting the integrated RCS contribution of an arbitrary-sized facet after intra-facet coherent interference has been accounted for.

Simultaneously, transforming graph nodes into arbitrary-scale facets introduces a severe numerical challenge: highly non-uniform physical scales generate order-of-magnitude disparities in the scattering contributions of different facets. To isolate this deterministic scale dependence, the scalar facet contribution can be decomposed as

$$Z_i = \underbrace{\alpha_i A_i e^{j2\boldsymbol{k}_s \cdot \boldsymbol{r}_i}}_{Z_i^{(0)}} \underbrace{\frac{1}{A_i} \iint_{\mathbb{S}_i} e^{j2\boldsymbol{k}_s \cdot (\boldsymbol{r}' - \boldsymbol{r}_i)}\, \mathrm{d}S'}_{C_i}, \qquad A_i = \iint_{\mathbb{S}_i} \mathrm{d}S', \tag{4}$$

where α_i collects the facet-wise orientation and polarization terms, $Z_i^{(0)}$ denotes the analytical PO prior evaluated at the facet center, and C_i denotes the normalized intra-facet coherent factor. This decomposition separates the dominant geometric scale A_i from the residual coherent variation within the facet.

2.1.2 Learnable Complex Correction Factor

While the analytical PO prior provides a baseline estimate of the facet contribution, the magnitude variation and phase shift induced by intra-facet coherent interference still need to be determined. To account for these effects, we employ a GNN to explicitly model the residual phase and magnitude variations beyond the PO prior.

Building upon Eq. 4, we express the normalized coherent factor as

$$C_i = \rho_i e^{j\Delta\phi_i}, \qquad \rho_i = \left| \frac{1}{A_i} \iint_{\mathbb{S}_i} e^{j2\boldsymbol{k}_s \cdot (\boldsymbol{r}' - \boldsymbol{r}_i)}\, \mathrm{d}S' \right| \leq 1, \tag{5}$$

where the bound follows directly from the triangle inequality. Thus, $\rho_i \in [0, 1]$ quantifies the amplitude retention after intra-facet coherent interference, while $\Delta\phi_i$ represents the phase shift relative to the analytical center-phase prior. When the phase is uniform across the facet, $\rho_i = 1$; otherwise, destructive interference leads to $\rho_i < 1$. This decomposition separates the analytically predictable facet-level scattering response from the residual effect of intra-facet coherent interference. Consequently, the network does not need to reconstruct the complete scattering contribution from geometry alone, but only learns the magnitude and phase corrections to the physically initialized response.

Consequently, the complete mathematical formulation of our EFSA is

$$Z_i = \underbrace{Z_i^{(0)}}_{\text{PO Prior}} \cdot \underbrace{\rho_i e^{j\Delta\phi_i}}_{\text{Complex Correction Factor}} . \tag{6}$$

In this disentangled formulation, the deterministic physical prior establishes the fundamental scale and baseline phase, while the GNN predicts the amplitude retention ρ_i and residual phase shift $\Delta\phi_i$. By decoupling the scale from the neural representation, EFSA breaks the high-frequency graph size bottleneck while preserving coherent electromagnetic aggregation.

2.2 Electrical-Size-Aware Scattering Prediction Network

The correction factors derived in Section 2.1.2 depend on local geometry, illumination, wavelength, and facet electrical size. Thus, our Electrical-Size-Aware Scattering Prediction Network aims to construct a highly expressive facet representation that synergizes these inputs while preserving surface topology.

We represent the discretized mesh as a directed k-nearest-neighbor (k-NN) graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ based on Euclidean distances between facet centers, where each node $v_i \in \mathcal{V}$ corresponds to a physical facet $\mathbb{S}_i$. This defines the spatial topology for our message-passing mechanism (Pfaff et al., 2021; Battaglia et al., 2018).

Condition-Invariant Geometric Backbone. Node initialization utilizes an eight-dimensional descriptor: globally normalized center coordinates, unit normal, curvature, and edge-strength. A three-layer EdgeConv backbone (Wang et al., 2019) extracts topological features using node states, neighbor differences, and relative displacements. Each block employs a shared multilayer perceptron (MLP) for edge features, max-pooling aggregation, and residual connections with layer normalization. By decoupling direction and wavelength from this backbone, the extracted geometry representation $\boldsymbol{h}_i^{\text{geo}}$ becomes condition-invariant. This allows efficient caching and reuse across varying observation conditions for the same target, drastically reducing computational overhead.

Illumination and Wavelength Modulation. Given an incident direction $\mathbf{d}$, we construct a six-dimensional illumination descriptor using $\hat{\boldsymbol{n}}_i \cdot \mathbf{d}$, $\|\hat{\boldsymbol{n}}_i \times \mathbf{d}\|$, tangent projections $\boldsymbol{t}_{i1} \cdot \mathbf{d}$ and $\boldsymbol{t}_{i2} \cdot \mathbf{d}$, curvature, and edge-strength. A direction-conditioned attention layer maps this descriptor to queries, and projects the neighbor geometry $\boldsymbol{h}_i^{\text{geo}}$ to keys and values. The wavelength is encoded via sinusoidal basis functions into a continuous embedding. Beyond controlling the attention bias, this embedding generates FiLM (Perez et al., 2018) parameters to dynamically modulate the direction-aware features. Concurrently, the facet electrical size is encoded by applying a dedicated MLP to $\log_{10}(A_i/\lambda^2)$.

Final Representation and Readout. Finally, the modulated direction features, condition-invariant geometry, illumination descriptor, and electrical-size representation are concatenated. This comprehensive vector is fed into the prediction heads to regress the EFSA amplitude attenuation ρ_i and phase shift $\Delta\phi_i$.

3 Experiments

In this section, we conduct extensive experiments to evaluate the prediction accuracy, geometric generalization, and high-frequency scalability of our framework. We first introduce the datasets, baselines, and implementation details, and then compare our method with representative learning-based approaches and the analytical PO initialization.

3.1 Dataset Preparation

We construct datasets containing both basic geometries and complex real-world objects (Wu et al., 2015). For each target, the RCS ground truth (GT) is generated by a PO solver (Jacobs et al., 2013) over 64,800 observation angles and frequencies from 3 to 6 GHz with an interval of 0.3 GHz. To obtain accurate reference responses, the original meshes are sufficiently re-meshed according to the wavelength before PO integration. In contrast, all learning-based methods directly take the original coarse meshes as input without high-frequency re-meshing.

In addition to the full-angle RCS responses, facet-wise reference scattering quantities are recorded at 4% of the observation angles for intermediate supervision. These reference quantities are obtained by coherently integrating the surface currents over the corresponding facets; therefore, the complex EFSA can be directly derived from the surface-current integral defined in Eq. 3. This sparse facet-level supervision substantially reduces the storage cost compared with recording reference quantities over all observation angles.

For each dataset, we split the targets into training and test sets, using 80% of the targets for training and the remaining 20% for testing. Detailed mesh visualizations are provided in Appendix A.4.

3.2 Training and Inference Details

All experiments and model evaluations are conducted on a workstation equipped with a single NVIDIA GeForce RTX 4090 GPU.

Data Split and Training Strategy. During training, the network is optimized using the full-angle RCS responses, while the facet-level EFSA supervision is applied only at the sparsely sampled 4% of observation angles where facet-wise reference scattering quantities are available.

Loss Function. The network is optimized end-to-end to accurately predict the EFSA. The overall training loss $\mathcal{L}$ is formulated to minimize the discrepancy between our network's predictions and the numerical ground truths, defined as follows:

$$\mathcal{L} = \lambda_{\mathrm{l1}}\mathcal{L}_{\mathrm{L1}} + \lambda_{\mathrm{mse}}\mathcal{L}_{\mathrm{MSE}} + \lambda_{\mathrm{face}}\mathcal{L}_{\mathrm{face}} + \lambda_{\mathrm{field}}\mathcal{L}_{\mathrm{field}} \tag{7}$$

Where $\mathcal{L}_{\mathrm{L1}}$ and $\mathcal{L}_{\mathrm{MSE}}$ denote the L1 and Mean Squared Error losses between the predicted and GT RCS. $\mathcal{L}_{\mathrm{face}}$ is the facet-level complex EFSA loss. $\mathcal{L}_{\mathrm{field}}$ is the total field complex coherence loss between the predicted and true summed fields. And λ_{l1}, λ_{mse}, λ_{face} and λ_{field} are 1.0, 1.0, 0.2 and 1.0, respectively.

Inference and Evaluation During the inference phase, the network is tasked to predict the full-angle RCS for the test targets across all frequency bands. All RCS-based metrics are evaluated in the dB domain. To quantitatively assess the prediction accuracy, we adopt standard evaluation metrics, including Mean Absolute Error (MAE) and Root Mean Square Error (RMSE). To specifically evaluate the network's performance in critical strong scattering regions, we introduce a weighted MAE (wMAE) metric, formulated as:

$$\mathrm{wMAE} = \frac{1}{N}\sum_{i=1}^{N}\sqrt{\max\left(y_i, 1\right)}\left|\hat{y}_i - y_i\right|, \tag{8}$$

where y_i and $\hat{y}_i$ denote the ground truth and predicted RCS values respectively, and N is the total number of evaluated data points. The weighting term places greater emphasis on samples with larger positive RCS values while retaining a unit weight for $y_i \leq 1$.

3.3 Baseline Methods

To comprehensively evaluate the performance of our proposed framework, we compare it against two state-of-the-art neural approaches for RCS prediction. These baselines are carefully selected to represent the two dominant paradigms in the field: direct end-to-end prediction and intermediate physical representation learning. For a fair comparison, both baseline methods are trained and evaluated on the exact same target split and frequency bands as our framework.

SC-RNet: SC-RNet (Tang et al., 2026a) is a highly representative architecture that directly maps geometric inputs to the RCS without intermediate physical reasoning. We specifically pre-processed our dataset to generate the high-resolution voxel data required for its training. This baseline serves to evaluate whether our framework outperforms purely data-driven direct mappings.

GraphSolver: GraphSolver is a representative GNN-based electromagnetic solver that predicts surface currents on triangular mesh elements (Cao et al., 2025). In its original formulation, each triangular element is represented by its centroid, and the surface current at the centroid is used to approximate the current distribution over the element. For a consistent comparison, we retain its surface-current-learning formulation while applying it to the

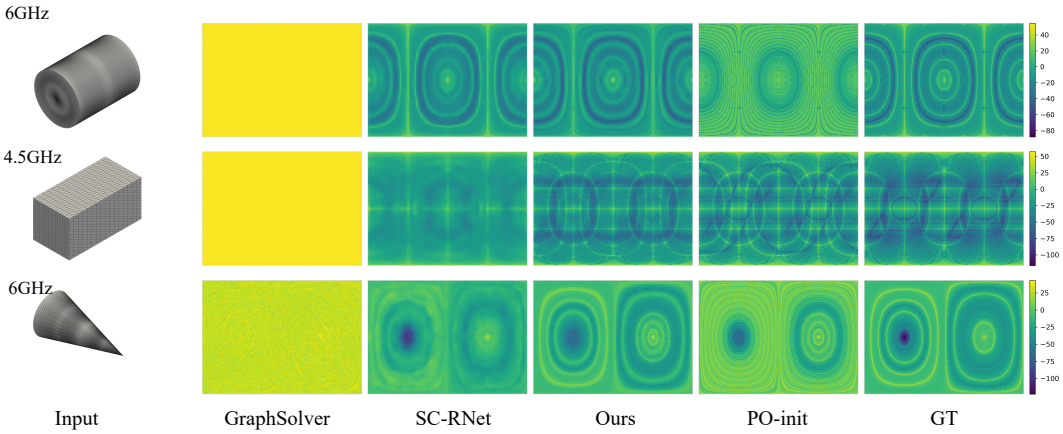

Figure 3: Visualization of prediction results of GraphSolver, SC-RNet, Ours and PO-init at different frequencies on basic geometries.

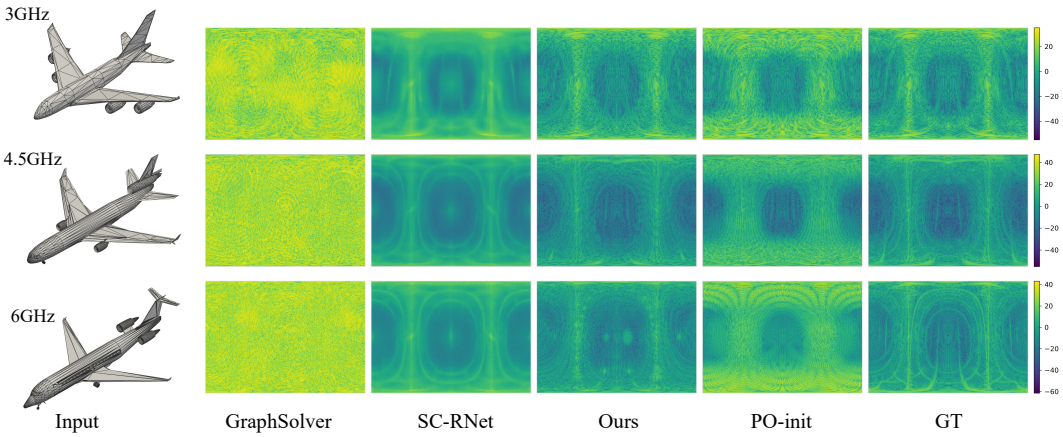

Figure 4: Visualization of prediction results of GraphSolver, SC-RNet, Ours and PO-init at different frequencies on ModelNet.

same coarse input meshes used by our method, with the corresponding facet-wise surface-current references as supervision. This setting evaluates how a current-based representation behaves when directly extended to electrically large facets without high-frequency re-meshing.

PO-init: We further evaluate the analytical PO initialization used in EFSA without the learnable complex correction factor. Specifically, PO-init directly uses $Z_i^{(0)}$ in Eq. 4 as the facet contribution and coherently aggregates all facets to obtain the final RCS. Comparing PO-init with our complete model directly evaluates the contribution of the learned correction for intra-facet coherent interference.

Table 1: Results of different discretizations.

#Facets	$\overline{d_{\max}/\lambda}$	Mem.	MAE	wMAE	RMSE
3k	3.3	2	6.88	8.57	9.71
1k*	6.5	/	19.48	20.97	23.03
1k	6.5	0.5	8.76	11.02	11.47
0.6k*	8.7	/	23.57	25.53	26.81
0.6k	8.7	0.3	9.63	12.57	12.47

3.4 Experimental Results

In this section, we conducted qualitative and quantitative analysis of the performance of all methods on various datasets. In addition, we also organized accuracy experiments of our framework on different mesh discretization on the basic geometry dataset.

Table 2: Quantitative comparison using MAE, wMAE, and RMSE. Lower values are better.

Frequency		3 GHz	4.5 GHz	6 GHz	Avg
Dataset	Method	MAE / wMAE / RMSE	MAE / wMAE / RMSE	MAE / wMAE / RMSE	MAE / wMAE / RMSE
Basic	PO-init	9.90 / 10.81 / 13.33	13.28 / 14.09 / 17.69	16.29 / 17.25 / 20.31	13.15 / 14.05 / 17.11
	GraphSolver	51.32 / 62.84 / 53.28	55.59 / 66.87 / 57.34	58.90 / 69.63 / 60.50	55.27 / 66.45 / 57.04
	SC-RNet	7.57 / 9.67 / 9.78	**7.29** / 10.41 / 10.71	7.80 / 10.00 / 9.95	7.88 / 10.03 / 10.15
	Ours	**6.51** / **7.74** / **9.30**	7.42 / **9.46** / **10.52**	**6.71** / **8.50** / **9.31**	**6.88** / **8.57** / **9.71**
Airplane	PO-init	7.00 / 9.89 / 9.21	8.79 / 12.90 / 11.02	10.38 / 15.34 / 12.66	8.72 / 12.71 / 10.97
	GraphSolver	27.56 / 37.94 / 29.97	30.13 / 42.69 / 32.35	32.30 / 46.48 / 34.26	30.00 / 42.37 / 32.19
	SC-RNet	5.99 / 8.46 / 7.64	5.97 / 8.68 / 7.62	6.06 / 8.90 / 7.70	6.01 / 8.68 / 7.65
	Ours	**4.51** / **7.01** / **6.27**	**4.91** / **8.16** / **6.73**	**5.22** / **8.91** / **7.07**	**4.88** / **8.03** / **6.69**

3.4.1 Performance on Basic Geometry

We first evaluate the predictive performance on basic geometric shapes. Although these targets possess relatively uniform mesh discretizations, their electrically large facet sizes present significant modeling challenges. Figure 3 and Table 2 show the qualitative and quantitative results.

SC-RNet: As a direct RCS prediction method, SC-RNet captures the general RCS envelope and several dominant scattering peaks. However, its predictions are considerably smoother than the GT and fail to reproduce the rapid angular oscillations, particularly as the frequency increases. This discrepancy indicates that directly regressing the final RCS makes it difficult to recover the fine angular variations associated with coherent scattering on electrically large targets.

GraphSolver: Under the coarse-mesh setting, GraphSolver produces nearly angle-independent predictions and fails to recover the characteristic angular variations of the GT. This behavior becomes particularly evident when the input facets are electrically large, where a single facet may contain substantial spatial phase variations that cannot be represented by a single facet-level current value. These results indicate the difficulty of directly extending its current-based representation to coarse meshes without high-frequency re-meshing.

PO-init: To evaluate the isolated contribution of our analytical physical prior, we test the PO-init baseline. As shown in the qualitative visualizations, PO-init successfully captures the general contour of coherent superposition across varying incident angles. However, quantitative metrics reveal that as the incident frequency increases, the performance of PO-init deteriorates rapidly, resulting in significantly overestimated RCS values. This agrees with our theoretical analysis: as the frequency increases, the facets become electrically larger and exhibit stronger intra-facet phase variations, which cannot be represented by the facet-center PO initialization alone.

Ours: In contrast, our framework accurately reconstructs these high-frequency variations. By introducing the learnable complex correction factor, our network effectively captures the destructive intra-facet interference, seamlessly compensating for the amplitude attenuation on electrically large facets that PO-init misses.

3.4.2 Performance on ModelNet-Airplane

Quantitative metrics in Table 2 and visualizations in Figure 4 demonstrate the effectiveness and generalization ability of our method on complex unseen targets.

SC-RNet: While SC-RNet captures the general trend of high-amplitude RCS peaks, its predictions exhibit highly similar patterns across geometrically distinct targets. This suggests limited sensitivity to target-specific scattering characteristics on unseen complex structures. Moreover, the discrepancy becomes more evident in regions with rapid angular variations, where the predicted responses are noticeably smoother than the GT. Such behavior indicates that directly mapping geometric features to the final RCS can recover the dominant response envelope, but remains challenging for resolving fine-grained angular variations associated with coherent scattering. Consequently, SC-RNet fails to reproduce some of the target-dependent high-frequency RCS variations observed in the GT.

GraphSolver: GraphSolver substantially overestimates the overall RCS on these complex targets, and the predicted angular responses deviate considerably from the GT across different frequencies. One important challenge under this setting is the highly heterogeneous facet scale of ModelNet meshes, which introduces large variations in facet-level scattering contributions. Unlike the

relatively uniform discretizations typically used for current prediction, the coarse and non-uniform facets here require the representation to accommodate both large contribution disparities and increasingly significant intra-facet phase variations. The observed prediction errors therefore highlight the difficulty of directly extending a current-based representation to electrically large and non-uniform mesh discretizations without high-frequency re-meshing.

PO-init: Similar to its performance on basic geometries, PO-init provides a reasonable approximate contour for the angular RCS distribution, indicating that the analytical prior captures a substantial part of the deterministic facet-level scattering response. Nonetheless, its predictions remain consistently higher than the ground truth, and its accuracy decreases at higher frequencies. This shows that the analytical prior alone is insufficient to account for the increasing intra-facet coherent interference on electrically large facets.

Ours: As shown in the visualization result, our method not only captures the overall trend of the RCS patterns, but also, benefiting from its physically interpretable modeling, accurately reproduces the coherent superposition of scattering contributions from different facets. Compared with PO-init, the predicted results further demonstrate that the learnable complex correction factor effectively accounts for the impact of intra-facet coherent interference on each facet's complex far-field contribution. By combining the deterministic PO prior with the learnable complex correction factor, our method maintains accurate predictions across different frequencies and unseen targets while effectively modeling intra-facet coherent interference.

3.4.3 Robustness to Mesh Discretization

We further investigate the robustness of our framework to different mesh discretizations on the Basic Geometry dataset. As summarized in Table 1, we progressively coarsen the input meshes from approximately 3k to 1k and 0.6k facets, increasing the average electrical facet size $d_{\max}/\lambda$ at 3GHz from 3.3 to 6.5 and 8.7, respectively. Notably, all results are obtained directly using the checkpoint trained exclusively on the 3k-facet discretization, without any retraining or fine-tuning.

As the mesh becomes coarser, the prediction error gradually increases, which is expected because larger facets exhibit stronger intra-facet phase variations and therefore require more substantial corrections to the analytical PO initialization. Nevertheless, our method maintains an MAE of 8.76 and 9.63 dB on the unseen 1k- and 0.6k-facet discretizations, respectively. In comparison, the corresponding PO-init errors increase to 19.48 and 23.57 dB. The persistent improvement over PO-init indicates that the learned complex correction remains effective even when the electrical facet size substantially exceeds that encountered during training, rather than merely fitting a particular mesh resolution. Meanwhile, reducing the facet count from 3k to 1k and 0.6k decreases the peak inference memory from 2 GB to 0.5 GB and 0.3 GB, respectively. These results demonstrate that EFSA enables a practical trade-off between mesh resolution, prediction accuracy, and memory consumption while retaining its correction capability across different discretizations. Detailed experimental settings and additional results are provided in Appendix A.3.

4 Conclusion

In this paper, we present a physically interpretable, scale-decoupled framework for high-frequency RCS prediction that addresses the scalability bottleneck of mesh-based GNNs. Our core contribution is the Equivalent Facet Scattering Amplitude (EFSA), an intermediate physical representation that decouples graph nodes from rigid sub-wavelength discretization constraints. By decomposing the facet-level scattering contribution into an analytical PO prior and a learnable complex correction factor, EFSA absorbs scale-induced disparities while capturing intra-facet coherent interference, enabling direct prediction on arbitrarily scaled physical facets. Experiments on both basic geometries and unseen real-world targets demonstrate the effectiveness of our framework in predictive accuracy, geometric generalization, and high-frequency scalability. Compared with direct prediction and current-based GNN baselines, our method accurately reproduces high-frequency angular RCS variations without requiring high-frequency re-meshing. These results support EFSA as a practical intermediate representation for learning high-frequency scattering responses without coupling the graph resolution to the incident wavelength. Future research will extend this scale-decoupled paradigm to more complex scenarios, such as targets with dielectric coatings or deep cavities involving intense multi-bounce interactions.

REFERENCES

Stefanos Bakirtzis, Marco Fiore, and Ian Wassell. Towards physics-informed graph neural network-based computational electromagnetics. In *2024 IEEE International Symposium on Antennas and Propagation and INC/USNC-URSI Radio Science Meeting*, pp. 673–674, 2024. doi: 10.1109/AP-S/INC-USNC-URSI52054.2024.10686000. URL https://doi.org/10.1109/AP-S/INC-USNC-URSI52054.2024.10686000.

Peter W. Battaglia, Jessica B. Hamrick, Victor Bapst, Alvaro Sanchez-Gonzalez, Vinicius Zambaldi, Mateusz Malinowski, Andrea Tacchetti, David Raposo, Adam Santoro, Ryan Faulkner, et al. Relational inductive biases, deep learning, and graph networks. *arXiv preprint arXiv:1806.01261*, 2018. doi: 10.48550/arXiv.1806.01261. URL https://arxiv.org/abs/1806.01261.

Huilin Cao, Haoran Sun, Jingwei Liang, Tao Shan, Yuxiang Ren, and Maokun Li. Physics-data of hybrid-driven edge-featured graph attention networks for surface currents learning 3-d PEC targets. *IEEE Transactions on Antennas and Propagation*, 73(11):9141–9153, 2025. doi: 10.1109/TAP.2025.3602467. URL https://doi.org/10.1109/TAP.2025.3602467.

Xinyu Fang, Mengmeng Li, Ziyang Lai, Davide Ramaccia, Alessandro Toscano, Filiberto Bilotti, and Dazhi Ding. Multifunctional space–time-modulated metasurface for direction of arrival estimation and rcs manipulation in a single system. *IEEE Trans. Microw. Theory Tech.*, 72(6): 3797–3808, 2024. doi: 10.1109/TMTT.2023.3330898.

Rui Guo, Zhichao Lin, Tao Shan, Maokun Li, Fan Yang, Shenheng Xu, and Aria Abubakar. Solving combined field integral equation with deep neural network for 2-d conducting object. *IEEE Antennas and Wireless Propagation Letters*, 20(4):538–542, 2021. doi: 10.1109/LAWP.2021.3056460. URL https://doi.org/10.1109/LAWP.2021.3056460.

Jin-Yang Han, Kun-Yi Guo, Bi-Yi Wu, and Xin-Qing Sheng. A physics-driven stacked ensemble network for efficient target RCS prediction. *IEEE Transactions on Antennas and Propagation*, 73(10):8116–8129, 2025. doi: 10.1109/TAP.2025.3582441. URL https://doi.org/10.1109/TAP.2025.3582441.

Roger F. Harrington. *Field Computation by Moment Methods*. Macmillan, New York, 1968. ISBN 9780780310148. URL https://www.wiley.com/en-us/Field+Computation+by+Moment+Methods-p-9780780310148.

Ockert B. Jacobs, Johann W. Odendaal, and Johan Joubert. Quad-ridge horn antenna with elliptically shaped sidewalls. *IEEE Transactions on Antennas and Propagation*, 61(6):2948–2955, 2013. doi: 10.1109/TAP.2013.2254436.

Xiaotian Jiang, Huan Huan Zhang, Kaitai Guo, Yang Zheng, and Jimin Liang. Mesh-based deep learning surrogate model for rapid wideband 3-d electromagnetic scattering computation. *IEEE Transactions on Antennas and Propagation*, 74(5):4546–4561, 2026. doi: 10.1109/TAP.2026.3661600. URL https://doi.org/10.1109/TAP.2026.3661600.

Eugene F. Knott, John F. Shaeffer, and Michael T. Tuley. *Radar Cross Section*. SciTech Publishing, Raleigh, NC, 2 edition, 2004. ISBN 9781891121258. doi: 10.1049/SBRA026E. URL https://doi.org/10.1049/SBRA026E.

De-Hua Kong, Jia-Ning Cao, Wen-Wei Zhang, Wen-Chi Huang, Xiao-Yang He, Lu Liu, and Ming-Yao Xia. An AI predictor: From point clouds to scattered far fields for 3-d PEC targets. *IEEE Transactions on Antennas and Propagation*, 72(6):5179–5190, 2024. doi: 10.1109/TAP.2024.3390615. URL https://doi.org/10.1109/TAP.2024.3390615.

Xiayu Li and Peng Chen. Distributed radar interval distance estimation based on deep neural network. In *Int. Conf. Wireless Commun. Signal Process.*, pp. 1–5. IEEE, 2022.

Yongzhong Li, Yinpeng Wang, Shutong Qi, Qiang Ren, Lei Kang, Sawyer D Campbell, Pingjuan L Werner, and Douglas H Werner. Predicting scattering from complex nano-structures via deep learning. *IEEE Access*, 8:139983–139993, 2020.

Zhenchao Ma, Kuiwen Xu, Rencheng Song, Chao-Fu Wang, and Xudong Chen. Learning-based fast electromagnetic scattering solver through generative adversarial network. *IEEE Trans. Antennas Propag.*, 69(4):2194–2208, 2020.

Ethan Perez, Florian Strub, Harm de Vries, Vincent Dumoulin, and Aaron Courville. FiLM: Visual reasoning with a general conditioning layer. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, pp. 3942–3951, 2018. doi: 10.1609/aaai.v32i1.11671. URL https://doi.org/10.1609/aaai.v32i1.11671.

Jesus Perez and Manuel F. Catedra. Application of physical optics to the RCS computation of bodies modeled with NURBS surfaces. *IEEE Transactions on Antennas and Propagation*, 42(10):1404–1411, 1994. doi: 10.1109/8.320747. URL https://doi.org/10.1109/8.320747.

Tobias Pfaff, Meire Fortunato, Alvaro Sanchez-Gonzalez, and Peter W. Battaglia. Learning mesh-based simulation with graph networks. In *International Conference on Learning Representations*, 2021. URL https://openreview.net/forum?id=roNqYL0_XP.

Shutong Qi, Yinpeng Wang, Yongzhong Li, Xuan Wu, Qiang Ren, and Yi Ren. Two-dimensional electromagnetic solver based on deep learning technique. *IEEE J. Multiscale Multiphys. Comput. Techniques*, 5:83–88, 2020.

Sadasiva M. Rao, Donald R. Wilton, and Allen W. Glisson. Electromagnetic scattering by surfaces of arbitrary shape. *IEEE Transactions on Antennas and Propagation*, 30(3):409–418, 1982. doi: 10.1109/TAP.1982.1142818. URL https://doi.org/10.1109/TAP.1982.1142818.

Ali Safa, Federico Corradi, Lars Keuninckx, Ilja Ocket, André Bourdoux, Francky Catthoor, and Georges G. E. Gielen. Improving the accuracy of spiking neural networks for radar gesture recognition through preprocessing. *IEEE Trans. Neural Netw. Learn. Syst.*, 34(6):2869–2881, 2023. doi: 10.1109/TNNLS.2021.3109958.

Tao Shan, Maokun Li, Fan Yang, and Shenheng Xu. Solving combined field integral equations with physics-informed graph residual learning for EM scattering of 3-d PEC targets. *IEEE Transactions on Antennas and Propagation*, 72(1):733–744, 2024. doi: 10.1109/TAP.2023.3331262. URL https://doi.org/10.1109/TAP.2023.3331262.

Tao Shan, Xin Zhang, and Di Wu. A fully data-driven solver for combined field integral equations of 3-d conducting bodies based on graph neural networks. *IEEE Transactions on Antennas and Propagation*, 74(7):6623–6633, 2026. doi: 10.1109/TAP.2026.3686620. URL https://doi.org/10.1109/TAP.2026.3686620.

Zehao Tang, Yuxuan Luo, Shuo Zhang, and Biao Leng. Scattering center guided mono-static radar cross section prediction. *Neural Networks*, 203:109061, 2026a. doi: 10.1016/j.neunet.2026.109061. URL https://doi.org/10.1016/j.neunet.2026.109061.

Zehao Tang, Kaiming Wu, Yuxuan Luo, Hao Wang, Shuo Zhang, and Biao Leng. ECHO: Estimating continuous radar cross section via hierarchical features of objects. *IEEE Transactions on Antennas and Propagation*, 74(8):7780–7791, 2026b. doi: 10.1109/TAP.2026.3686018. URL https://doi.org/10.1109/TAP.2026.3686018.

Mou Wang, Shunjun Wei, Zichen Zhou, Jun Shi, Xiaoling Zhang, and Yongxin Guo. Ctv-net: Complex-valued tv-driven network with nested topology for 3-d sar imaging. *IEEE Trans. Neural Netw. Learn. Syst.*, 35(4):5588–5602, 2024. doi: 10.1109/TNNLS.2022.3208252.

Yue Wang, Yongbin Sun, Ziwei Liu, Sanjay E. Sarma, Michael M. Bronstein, and Justin M. Solomon. Dynamic graph CNN for learning on point clouds. *ACM Transactions on Graphics*, 38(5):146:1–146:12, 2019. doi: 10.1145/3326362. URL https://doi.org/10.1145/3326362.

Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, and Jianxiong Xiao. 3d shapenets: A deep representation for volumetric shapes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1912–1920, 2015.

Yinfang Xu, Jiaqi Liu, Lu Gao, Xia Ai, Zhongjie Zhang, and Huan Liang. 1-bit coding metasurface for wideband rcs reduction. In *Int. Appl. Comput. Electromagn. Soc. Symp.*, pp. 1–3, 2022. doi: 10.1109/ACES-China56081.2022.10065103.

He Ming Yao, Min Li, and Lijun Jiang. Applying deep learning approach to the far-field subwavelength imaging based on near-field resonant metalens at microwave frequencies. *IEEE Access*, 7: 63801–63808, 2019a.

He Ming Yao, EI Wei, and Lijun Jiang. Two-step enhanced deep learning approach for electromagnetic inverse scattering problems. *IEEE Antennas Wireless Propag. Lett.*, 18(11):2254–2258, 2019b.

He Ming Yao, Lijun Jiang, and EI Wei. Enhanced deep learning approach based on the deep convolutional encoder–decoder architecture for electromagnetic inverse scattering problems. *IEEE Antennas Wireless Propag. Lett.*, 19(7):1211–1215, 2020.

Wen-Wei Zhang, De-Hua Kong, Xiao-Yang He, and Ming-Yao Xia. A machine learning method for 2-d scattered far-field prediction based on wave coefficients. *IEEE Antennas and Wireless Propagation Letters*, 22(5):1174–1178, 2023. doi: 10.1109/LAWP.2023.3235928. URL https://doi.org/10.1109/LAWP.2023.3235928.

Zhao Zhang, Danping He, Ke Guan, Ben Chen, Jianwu Dou, Wei E. I. Sha, and Zhangdui Zhong. A deep learning based surface current generation method for scattering modeling at terahertz band. In *2024 18th European Conference on Antennas and Propagation (EuCAP)*, pp. 1–5, 2024. doi: 10.23919/EuCAP60739.2024.10501597. URL https://doi.org/10.23919/EuCAP60739.2024.10501597.

5 AI USE STATEMENT

In this work, we used generative AI tools for help develop theoretical models or conceptual frameworks, implement methods, assist with translation, We have not used generative AI tools for Generate synthetic data sets, formulate mathematical claims, provide critical ingredients for proving mathematical claims, assist in the writing of proofs, propose or refine hypotheses, design or provide feedback on research methodology or experiments, clean and reformat dataset, support qualitative and thematic data analysis, interpret results, and the rest of the required disclosure tasks are not applicable to this work. Additionally, we used generative AI tools for create or modify scientific figures or images, suggest experimental parameters, create or edit software code, draft parts of a research paper, transcribe recordings of research material, summarize or analyses existing literature, sourcing/searching for information, edit a research paper to improve readability, identify relevant literature, format references, propose a title or keywords for a research paper. We have reviewed all AI-assisted work. We take responsibility for the final content of this work, including text, claims or artifacts produced with the aid of generative AI.