

Electromagnetic Sculptor: a differentiable geometric optimization framework to manipulate electromagnetic fields

Received: 7 July 2025

Accepted: 5 March 2026

Cite this article as: Yang, K., Liu, C., Yu, W. *et al.* Electromagnetic Sculptor: a differentiable geometric optimization framework to manipulate electromagnetic fields. *Commun Eng* (2026). <https://doi.org/10.1038/s44172-026-00642-3>

Kaiqiao Yang, Che Liu, Wenming Yu & Tie Jun Cui

We are providing an unedited version of this manuscript to give early access to its findings. Before final publication, the manuscript will undergo further editing. Please note there may be errors present which affect the content, and all legal disclaimers apply.

If this paper is publishing under a Transparent Peer Review model then Peer Review reports will publish with the final article.

Electromagnetic Sculptor: A Differentiable Geometric Optimization

Framework to Manipulate Electromagnetic Fields

Kaiqiao Yang^{1,2,3}, Che Liu^{1,2,3}, Wenming Yu^{1,2}, and Tie Jun Cui^{1,2*}

1 The State Key Laboratory of Millimeter Waves, Southeast University, Nanjing 210096, China

2 Institute of Electromagnetic Space, Southeast University, Nanjing 210096, China

3 These authors contributed equally.

Email: tjcui@seu.edu.cn

Abstract

Electromagnetic fields are commonly controlled through geometric design, but existing approaches often lack efficient and differentiable modeling tools for complex shapes. Here we introduce Electromagnetic Sculptor, a differentiable geometric optimization framework for manipulating electromagnetic fields on arbitrarily meshed structures. The framework combines a numerical electromagnetic model based on shooting and bouncing rays with a gradient-based geometric optimizer that stabilizes mesh deformation through spatial filtering. To avoid excessive shape distortion during optimization, a shape-preserving regularization strategy is incorporated. The method is demonstrated using radar cross section reduction as a representative application. Numerical and experimental results show pronounced field suppression at both single frequencies and across a broadband range, while maintaining geometric smoothness and manufacturability. The framework enables fast optimization for models containing thousands of vertices, with simulated results consistent with experimental measurements. These results illustrate how differentiable computation can be integrated with physically grounded electromagnetic modeling and practical design constraints.

Introduction

Just as a sculptor carves raw material into a desired shape based on a reference image—a process

demanding exceptional skill and precision—electromagnetic engineers face a similar challenge: designing scatterers or devices to mold electromagnetic wave energy into predefined patterns. This class of problems, termed geometry optimization, is pivotal in modern electromagnetic engineering and has a wide range of applications, including microwave device optimization^{1,2}, stealth design^{4,5}, and electromagnetic imaging^{6,7}. In modern wireless communications, geometry optimization is also indispensable: while forward channel modeling and performance optimization have been extensively studied^{3,8}, the geometric configuration of scatterers ultimately determines channel quality. However, solving electromagnetic geometry optimization problems remains challenging due to the immense variability of object shapes and the highly nonlinear nature of electromagnetic field governing equations⁹. On the one hand, it is difficult to derive concise and general design theories that are directly applicable to complex geometries. On the other hand, achieving sufficient modeling generality often requires a large number of design parameters, which readily leads to the curse of dimensionality. As a result, despite over a century of advances in electromagnetic field theory, there is still a lack of direct and powerful theoretical tools for systematic geometric optimization, making numerical optimization techniques an inevitable choice in practice. In this context, developing efficient and reliable geometric optimization frameworks is essential to bridge the gap between theoretical design and engineering implementation.

Heuristic optimization methods have been widely adopted to address complex electromagnetic geometry optimization problems by mimicking collective or evolutionary behaviors in nature, such as genetic algorithms⁹, particle swarm optimization¹⁰, and grey wolf optimizer¹¹. Owing to their strong global search capability, these methods are well suited for nonlinear and non-convex problems and have been successfully applied to inverse scattering^{12,13}, stealth design^{14,15}, and topology optimization^{16–18}. Despite these successes, heuristic methods suffer from inherent limitations in practical electromagnetic applications. Their reliance on repeated fitness evaluations results in high computational cost^{15,19–21}, while surrogate modeling strategies only partially alleviate this burden^{15,19–21}. Moreover, the stochastic nature of heuristic

algorithms and their dependence on empirical parameter tuning lead to instability in convergence behavior and solution quality²². As a result, existing studies are typically limited to problems with a few hundred design parameters^{12–18}. Although such methods may be adequate for simple two-dimensional (2D) cases, they struggle to scale to complex three-dimensional (3D) geometries with fine geometric variations. These limitations significantly restrict the applicability of heuristic optimization methods in modern electromagnetic system design, particularly for high-dimensional and high-precision geometric optimization tasks.

In recent years, advances in artificial intelligence (AI) have introduced new opportunities to alleviate the computational bottlenecks of traditional geometric optimization methods. Consequently, intelligent geometric optimization has emerged, encompassing surrogate model optimization, generative optimization, and reinforcement learning–based approaches. Intelligent surrogate model optimization accelerates geometric optimization by replacing expensive electromagnetic simulations with learned approximations. Compared with traditional surrogate models^{13,15,19–21}, deep learning–based surrogates exhibit stronger nonlinear modeling capability and have been successfully applied to antenna and metamaterial design^{23–26}. However, both traditional and intelligent surrogate models still suffer from limited generalization performance. Generative intelligent optimization attempts to bypass iterative optimization by directly synthesizing geometries from design objectives^{2,27–29}, achieving high efficiency and even sub-second response time²⁸. Nevertheless, this approach typically requires large amounts of high-quality simulation data, resulting in high training cost, and remains challenged by convergence and model complexity issues despite recent efforts incorporating physical constraints^{30,31}. Reinforcement learning–based optimization has also been explored for high-dimensional geometric optimization problems^{32–34}, demonstrating improved efficiency over heuristic methods^{32–34}, but still facing challenges in search efficiency and convergence stability. Overall, although AI-based methods provide promising alternatives for geometric optimization, their data-driven nature conflicts with the limited availability of electromagnetic data. Moreover, most existing AI frameworks are developed for image-based tasks and are effective primarily for two-

dimensional structures, while their applicability to complex three-dimensional geometries remains insufficient. These limitations continue to hinder the practical deployment of intelligent geometric optimization in modern electromagnetic engineering.

Gradient descent-based optimization provides a third alternative for addressing complex geometric optimization problems involving high-dimensional, nonlinear, and small-sample scenarios. Compared with gradient-free methods such as heuristic optimization, gradient descent can exploit local sensitivity information to guide the optimization direction and avoid blind search and redundant evaluations. Unlike data-driven intelligent optimization methods, gradient descent does not require large-scale training data, making it particularly suitable for small-sample and stochastic electromagnetic optimization problems. As a result, gradient-based approaches have achieved notable success in a wide range of applications, including waveguides^{35,36}, antennas^{37,38}, metamaterials^{1,39}, and low-scattering structures^{4,5}. The key challenge of gradient-based optimization lies in efficient and accurate gradient computation. Early studies primarily relied on finite-difference methods⁴⁰, whose computational cost scales poorly with increasing parameter dimension. This limitation has been largely addressed by the adjoint method, which enables gradient evaluation for all parameters using only one forward and one adjoint simulation^{1,4,5,36,37}, thereby supporting high-dimensional free-form optimization. More recently, advances in automatic differentiation, driven by modern AI frameworks such as TensorFlow⁴¹ and PyTorch⁴², have further reduced the complexity of gradient computation and facilitated the emergence of differentiable electromagnetics^{7,44,45}. Despite these advances, gradient descent methods still face critical challenges in practical electromagnetic optimization. Large variations in parameter sensitivity can lead to optimization instability and excessive geometric distortion⁵, while repeated full-wave simulations impose prohibitive computational costs in three-dimensional scenarios¹⁵. Moreover, the greedy nature of gradient-based strategies makes them susceptible to local minima⁴⁶ and saddle-point stagnation⁴⁷. These limitations indicate that further algorithmic and computational innovations are required to fully exploit the potential of gradient descent for large-scale electromagnetic geometric optimization.

To address the geometric optimization of three-dimensional perfectly electric conducting (PEC) objects, we propose Electromagnetic Sculptor, a differentiable geometric optimization framework for electromagnetic field manipulation. Built upon mesh-based representations, the framework exhibits excellent scalability and supports arbitrary 3D geometries with thousands of degrees of freedom, overcoming the limitations of traditional parametric models.

To tackle key challenges in conventional differentiable electromagnetic optimization—namely optimization oscillation, premature stagnation, and the high computational cost of full-wave simulations—Electromagnetic Sculptor introduces three core innovations. First, the shooting and bouncing ray (SBR) method⁴⁸, combined with physical optics (PO)⁵², is adopted to replace full-wave solvers, enabling efficient and differentiable high-frequency electromagnetic simulation that naturally aligns with polygonal meshes. Second, a spatially filtered geometric optimizer based on Fibonacci lattices is proposed to stabilize gradient-based optimization by diffusing local sensitivity information across neighboring surface regions, thereby preventing excessive geometric distortion and improving convergence robustness^{5,49,50}. Third, the framework is validated through radar cross-section (RCS) optimization and compact range experiments. By incorporating shape-preserving constraints inspired by optical design⁵¹, the optimization maintains global contour consistency while allowing localized geometric adaptation, ensuring manufacturability. Beyond performance validation, we further investigate the practical applicability boundary of the proposed optimization by analyzing the relationship between angular sampling density and optimization consistency. The results indicate that stable and consistent RCS reduction requires the sampling strategy to be compatible with the intrinsic spatial-frequency characteristics of the target's scattering response. Experimental results demonstrate that the proposed framework achieves approximately 6 dBsm broadband RCS reduction over the frequency range from 1 to 4 GHz. By integrating these designs and analyses, Electromagnetic Sculptor provides an efficient and physically grounded solution for large-scale electromagnetic geometric optimization, with strong potential for accelerating practical electromagnetic system design.

Methods

This section describes the methods used in this study. A large language model was also used solely to improve the clarity and grammar of the manuscript; all scientific content and interpretations were developed, reviewed and verified by the authors.

Problem Statement

The electromagnetic shape optimization problem is a typical optimization problem aimed at manipulating the electromagnetic scattering energy of a target object through shape design. To simplify the problem complexity, our study focuses only on the PEC objects with plane wave excitation under far-field conditions. Assume that the object surface to be optimized can be represented as a 2D manifold, denoted by M . According to the electric field integral equation, the electric field intensity at any point in the external space of the object under the excitation of \mathbf{E}^{inc} is expressed as:

$$\mathbf{E}(\mathbf{r}) = \mathbf{E}^{\text{inc}}(\mathbf{r}) - \frac{j\omega\mu}{4\pi} \iint_M \frac{e^{-jk|\mathbf{r}-\mathbf{r}'|}}{|\mathbf{r}-\mathbf{r}'|} \mathbf{J}(\mathbf{r}') dS, \quad (1)$$

where ω refers to the angular frequency, μ refers to the magnetic permeability, k refers to the wave number, \mathbf{r} refers to the field point, \mathbf{r}' refers to the source point, \mathbf{J} refers to the induced current, and dS refers to the differential area element on the manifold M . The integral term in the righthand side of (1) represents the scattering field, which can be approximated under the far-field condition:

$$\lim_{r \rightarrow \infty} \mathbf{E}^{\text{sc}}(\mathbf{r}) \approx -\frac{j\omega\mu e^{-jkr}}{4\pi r} \iint_M \mathbf{J}(\mathbf{r}') e^{jk\hat{\mathbf{r}}' \cdot \hat{\mathbf{r}}} dS. \quad (2)$$

In electromagnetic geometry optimization problems, the objective is typically to optimize the scatterer geometry so that the scattered energy approximates to the desired target energy as closely as possible. Assuming that the test function at the receiving end is $\boldsymbol{\xi}(\mathbf{r}) = \hat{\mathbf{a}}\delta(\mathbf{r})$ (where $\delta(\mathbf{r})$ is the Dirac delta function), this test function essentially reflects a specific polarization selection. Under this definition of the test function $\boldsymbol{\xi}(\mathbf{r})$, the scattered energy can be expressed as

$\sigma^{\text{sc}} = [\hat{\boldsymbol{\xi}}(\mathbf{r}) \cdot \mathbf{E}^{\text{sc}}(\mathbf{r})]^2$. By defining the desired scattered energy as σ' , we can formulate the electromagnetic geometry optimization problem as:

$$M' = \arg \min_M \frac{1}{\Delta\Omega_{\text{rec}} \Delta\Omega_{\text{trans}} B} \iiint (\sigma^{\text{sc}} - \sigma')^2 d\hat{\mathbf{k}}^{\text{sc}} d\hat{\mathbf{k}}^{\text{inc}} d\omega, \quad (3)$$

where $\Delta\Omega_{\text{rec}}$ refers to the reception range, $\Delta\Omega_{\text{trans}}$ refers to the transmission range, B refers to the operating bandwidth, $\hat{\mathbf{k}}^{\text{sc}}$ refers to the scattering direction and $\hat{\mathbf{k}}^{\text{inc}}$ refers to the incidence direction. Here, we adopt a normalization factor to eliminate the influence of the integral domain size. Notably, when the desired scattering energy is set to zero, equation (3) describes an RCS optimization problem, which is the primary scenario we will examine in subsequent experiments.

Optimization Framework

Clearly, solving the geometry optimization problem is quite challenging, and it is nearly impossible to find the optimal manifold M . Typically, a quasi-optimal solution is sought as an alternative. Before solving the problem, we discretize the manifold M into a triangular mesh (V, E) for computational convenience, where V represents the set of vertices and E represents the set of edges.

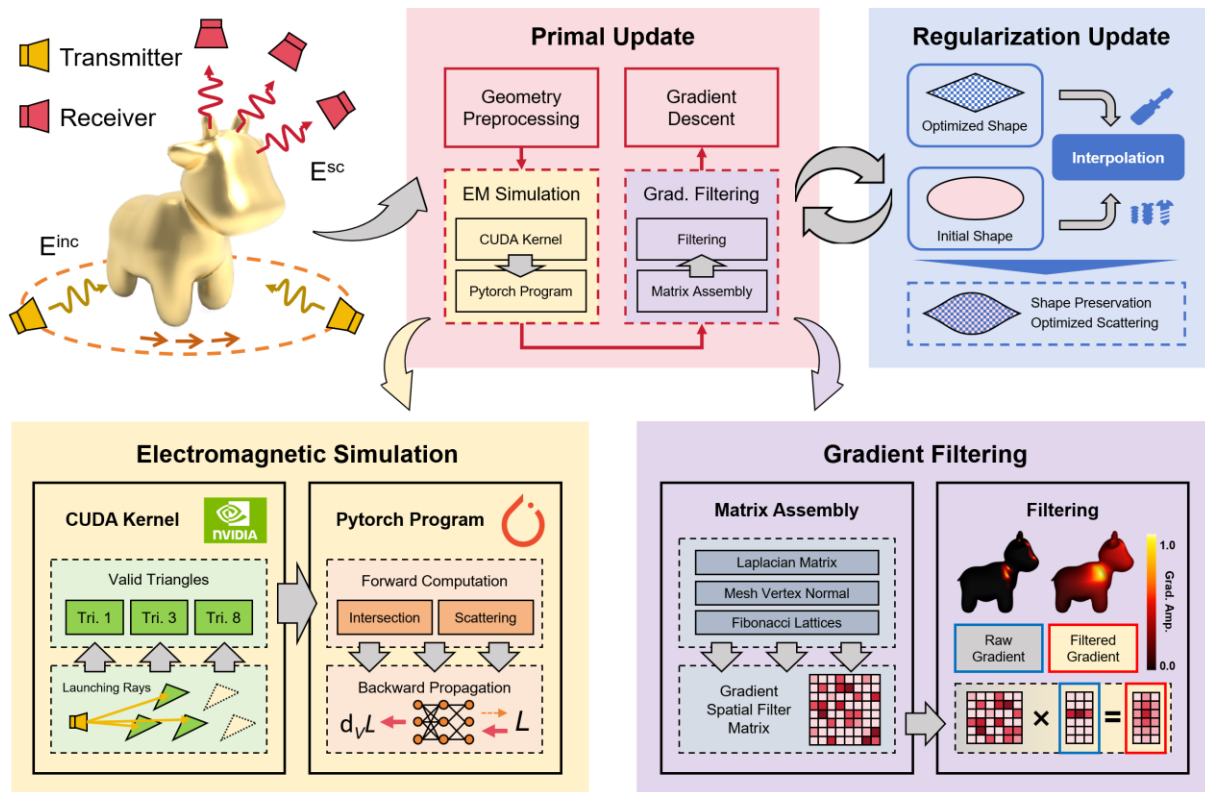


Fig. 1 The Electromagnetic Sculptor framework. This computational framework alternates between two core operations: a primal update conducting electromagnetic simulation (forward modeling and backward gradient computation) followed by filtered gradient descent optimization, and a regularization update enforcing geometric constraints through shape-preserving deformation.

We present a schematic diagram of the optimization framework in Fig. 1. Initially, we use a differentiable SBR simulator to compute the scattering energy of the object to be optimized. Next, we apply the backpropagation mechanism to calculate the derivatives of the loss function with respect to the vertices, and use a spatial filtering-based gradient descent optimizer to modify the vertex positions, yielding a new shape. After that, we perform a regularization step on the object's shape, interpolating the newly optimized shape with the initial shape to achieve the desired one. We repeat these processes until we get a satisfactory optimization result.

Differentiable SBR Method

Here, we construct a differentiable electromagnetic simulator based on the SBR method. Though the SBR method is widely used in electromagnetic analysis, it faces a fundamental limitation

stemming from the zero-volume nature of ray geometry: the probability of a ray path precisely intersecting any discrete receiver point is theoretically zero (unless the receiver happens to lie exactly on the propagation path). This results in null-field conditions that prevent the gradient backpropagation in the computational graph. While forward ray tracing followed by receiver-point position refinement may help, this solution inherently suffers from two critical drawbacks: they cannot fundamentally resolve the discrete sampling nature of ray methods, while simultaneously introduce substantial additional computational overhead. We introduce a PO approximation at the final scattering interaction⁵². This approach establishes a continuous far-field scattering model that enables electromagnetic field to query along arbitrary spatial directions, effectively overcoming the path discontinuity issues while endowing the SBR methodology with full differentiability. According to the theory of PO approximation, the final scattered field can be written as:

$$\mathbf{E}_{\text{SBR}}^{\text{sc}} = -\frac{j\omega\mu e^{-jkr}}{4\pi r} \iint_{S_{\text{ill}}} [2\hat{\mathbf{n}} \times \mathbf{H}^{\text{inc}}(\mathbf{r}'; V, E) \times \hat{\mathbf{k}}^{\text{sc}}] \times \hat{\mathbf{k}}^{\text{sc}} e^{j\mathbf{k}^{\text{sc}} \cdot \mathbf{r}'} dS, \quad (4)$$

where S_{ill} represents the mesh region illuminated by electromagnetic waves, \mathbf{H}^{inc} represents the direct and indirect incident magnetic field. Equation (4) provides the mathematical formulation of the SBR method, which indicates that the spatial scattered field is generated by the combined effects of direct and indirect electromagnetic excitation. The indirect excitation arises from electromagnetic scattering within the object and is closely related to the shape of object.

For practical implementation in electromagnetic simulations, Equation (4) must firstly be discretized into a computationally solvable form. Through mathematical derivation (see Supplementary Notes 1), we obtain its approximate expression under the far-field conditions:

$$\mathbf{E}_{\text{SBR}}^{\text{sc}} = -\frac{j\omega\mu e^{-jkr}}{4\pi r} \sum_{l=1}^N \{ [(2\hat{\mathbf{n}}_l \times \hat{\mathbf{h}}_l) \times \hat{\mathbf{k}}^{\text{sc}}] \times \hat{\mathbf{k}}^{\text{sc}} \} u_l v_l \text{sinc} \left[\frac{(\mathbf{k}_l^{\text{inc}} - \mathbf{k}^{\text{sc}}) \cdot \mathbf{u}_l}{2} \right] \text{sinc} \left[\frac{(\mathbf{k}_l^{\text{inc}} - \mathbf{k}^{\text{sc}}) \cdot \mathbf{v}_l}{2} \right] e^{-j(kd_l - \mathbf{k}^{\text{sc}} \cdot \mathbf{p}_l)}, \quad (5)$$

where N represents the total number of hit rays, $(\cdot)_l$ represents the index of a hit ray, $\hat{\mathbf{h}}$ represents the polarization direction, d represents the accumulated propagation distance, \mathbf{p} represents the hit point, and \mathbf{u}, \mathbf{v} are used to form the silhouette of a ray tube. From Equation (5), we can find that $\mathbf{E}_{\text{SBR}}^{\text{sc}}$ is determined by parameters $\hat{\mathbf{n}}, \hat{\mathbf{h}}, \hat{\mathbf{k}}^{\text{inc}}, \mathbf{u}, \mathbf{v}, d, \mathbf{p}$. In fact, these parameters can

be computed from the vertex and edge set (V, E) , which means $\mathbf{E}_{\text{SBR}}^{\text{sc}}$ can be viewed as a function of mesh (V, E) . Therefore, we establish a complete mapping from the vertex set V to the scattered field $\mathbf{E}_{\text{SBR}}^{\text{sc}}$, and further to the geometric optimization loss function L , enabling the derivation of the gradient expression $\nabla_V L$ via the chain rule:

$$\nabla_V L = \frac{\partial L}{\partial \mathbf{E}_{\text{SBR}}^{\text{sc}}} \sum_{l=1}^N \frac{\partial \mathbf{E}_{\text{SBR}}^{\text{sc}}}{\partial \hat{\mathbf{n}}_l} \frac{\partial \hat{\mathbf{n}}_l}{\partial V} + \frac{\partial \mathbf{E}_{\text{SBR}}^{\text{sc}}}{\partial \hat{\mathbf{h}}_l} \frac{\partial \hat{\mathbf{h}}_l}{\partial V} + \frac{\partial \mathbf{E}_{\text{SBR}}^{\text{sc}}}{\partial \mathbf{k}_l^i} \frac{\partial \mathbf{k}_l^i}{\partial V} + \frac{\partial \mathbf{E}_{\text{SBR}}^{\text{sc}}}{\partial \mathbf{u}_l} \frac{\partial \mathbf{u}_l}{\partial V} + \frac{\partial \mathbf{E}_{\text{SBR}}^{\text{sc}}}{\partial \mathbf{v}_l} \frac{\partial \mathbf{v}_l}{\partial V} + \frac{\partial \mathbf{E}_{\text{SBR}}^{\text{sc}}}{\partial d_l} \frac{\partial d_l}{\partial V} + \frac{\partial \mathbf{E}_{\text{SBR}}^{\text{sc}}}{\partial \mathbf{p}_l} \frac{\partial \mathbf{p}_l}{\partial V}. \quad (6)$$

The bouncing process causes a dependency between the ray's current and previous states, which makes the expansion of Equation (6) too complex to be computed directly. To address this issue, we develop a differentiable SBR simulation framework leveraging GPU parallel computing and PyTorch's automatic differentiation tools. The Electromagnetic Simulation module in Fig. 1 illustrates the proposed differentiable SBR framework, which consists of two main components: a CUDA kernel function and a PyTorch programs. Firstly, we perform GPU parallel ray tracing to obtain topological information about the regions hit by the rays, including the face indices and the vertex indices corresponding to each triangle. Next, the topological information is passed to the PyTorch program for scattering field calculations. During execution, PyTorch maintains a computation graph for the scattering field, which stores the dependencies between the vertices and other variables, ultimately enabling fast and accurate gradient computation through backpropagation.

While the preceding discussions are primarily focused on PEC structures, our method can be effectively extended to more general scenarios such as dielectric objects and dielectric-coated PEC structures. For detailed derivations and implementation specifics, we refer readers to the Supplementary Notes 2.

Spatial Filtering-Based Geometry Optimizer

Although the gradient descent has demonstrated broad utility in optimization tasks, its naive application to polygonal mesh deformation is fundamentally limited: severe shape distortion was found in prior work^{5,49,50}, which was rooted in both geometric and electromagnetic constraints. Geometrically, the method's greedy vertex-wise updates lack mechanisms to preserve the global

shape coherence, inevitably producing self-intersecting or topologically invalid configurations that corrupt the mesh integrity. While restricting updates to a subset of vertices might temporarily alleviate the distortion, such constraints nullify the core advantage of the gradient-based optimization: its ability to leverage full geometric freedom. On the other hand, the inherent spatial volatility of electromagnetic fields generates conflicting sensitivity gradients across the surface, where the unmodified gradient descent amplifies these disparities into physical implausibilities like creases and folds. These artifacts not only destabilize the optimization but degrade the fidelity of subsequent field simulations, creating a pathological feedback loop that traps the process in shallow local minima.

These challenges fundamentally stem from uncoordinated vertex updates leading to overly aggressive geometric deformations. To address this, we propose a spatial filtering-based geometric optimizer whose core innovation lies in intelligent gradient field diffusion to achieve coordinated vertex deformation. As illustrated in the Gradient Filtering module in Fig. 1, we apply an anisotropic spatial filter to vertex gradient field, enabling adaptive propagation of local gradient information to neighboring regions. This approach maintains the geometric features while ensures the overall morphological stability. Unlike the conventional Laplacian filters⁴⁹ that cause excessive smoothing, we employ the anisotropic filtering template⁵³, which achieves optimal balance between the high-frequency oscillation suppression and the geometric detail preservation through its direction-dependent weight allocation mechanism. The filtered gradient $F(\nabla_V L)$ is shown as:

$$F(\nabla_V L) = \frac{\sum_{j=1}^{N_s} e^{\frac{\mathbf{N}_v \hat{\mathbf{s}}_j}{\sigma_d}} (\mathbf{I} + \lambda_g \mathbf{L})^{-1} e^{\frac{\mathbf{N}_v \hat{\mathbf{s}}_j}{\sigma_\phi}} \nabla_V L}{\sum_{j=1}^{N_s} e^{\frac{\mathbf{N}_v \hat{\mathbf{s}}_j}{\sigma_d}} (\mathbf{I} + \lambda_g \mathbf{L})^{-1} e^{\frac{\mathbf{N}_v \hat{\mathbf{s}}_j}{\sigma_\phi}} \mathbf{1}}, \quad (7)$$

where \mathbf{N}_v represents the vertex normal matrix, $\hat{\mathbf{s}}_j$ represents the i -th direction among N_s sampling directions, \mathbf{L} represents Laplacian matrix, $\mathbf{1}$ represents an all-ones matrix, σ_d , σ_ϕ , λ_g are hyper-parameters.

The filtering operation described in Equation (7) involves a directional sampling process, where a set of uniformly distributed directions are sampled from a unit sphere to compute the

spatial weights for each vertex. This ensures that smoothing primarily occurs in the regions sharing similar normal vectors. Following the prior work⁵³, Sobol sequences were initially employed to achieve the uniform sampling on the unit sphere. However, our empirical observations reveal that this stochastic sampling approach leads to instability in the optimizer performance, producing inconsistent and erratic results. We note that while the Sobol sequences tend to generate equidistant points on the spherical surface, their inherent randomness introduces undesirable variance during optimization. To address this, we propose to replace the random sampling with the Fibonacci spherical lattice points, which provide deterministic and stable spatial distribution. The construction method for these Fibonacci lattice points is given by:

$$\begin{aligned} \theta_i &= \arccos\left(1 - \frac{2i}{N_F}\right), \quad \phi_i = 2\pi \cdot i \cdot 0.618, \\ \hat{\mathbf{s}}_i^F &= (\sin \theta_i \cos \phi_i, \sin \theta_i \sin \phi_i, \cos \theta_i)^T \end{aligned} \quad (8)$$

where $\hat{\mathbf{s}}_i^F$ refers to the i -th sample of N_F Fibonacci lattice points.

Ultimately, integrating the spatial filtering operation into UniformAdam^{49,53} yields a stable and effective geometric optimizer:

$$\begin{aligned} \mathbf{g} &\leftarrow \mathbf{F}(\nabla_V L) \\ \mathbf{m}_1 &\leftarrow \beta_1 \mathbf{m}_1 + (1 - \beta_1) \mathbf{g} \\ \mathbf{m}_2 &\leftarrow \beta_2 \mathbf{m}_2 + (1 - \beta_2) \mathbf{g}^2, \\ V &\leftarrow V - \frac{\eta}{\sqrt{\|\mathbf{m}_2\|_\infty}} \mathbf{F}(\mathbf{m}_1) \end{aligned} \quad (9)$$

where η represents the learning rate, $\mathbf{m}_1, \mathbf{m}_2$ are first- and second-order momentum, and β_1, β_2 are two hyper-parameters.

Shape-Preserving Regularization

To validate the effectiveness of Electromagnetic Sculptor, we conduct an RCS optimization experiment. For practical model fabrication and testing purposes, we further propose a shape-preserving regularization constraint based on the assumption that the object's geometry should maintain structural stability throughout the optimization process, as the optimized shape is typically required to retain essential geometric fidelity to the original design. Drawing on the research from computer graphics, we introduce as-rigid-as-possible (ARAP) regularization

technique to limit the vertex movements⁵¹. As a result, the final optimization output retains both low-scattering characteristics and shape preservation properties.

The optimization problem with the added regularization term can be rewritten as:

$$V = \arg \min_V E_{EM}(V, E) + \lambda_s E_s(V, V_0, E), \quad (10)$$

where E_{EM} represents the average electromagnetic scattered energy, λ_s denotes a regularization parameter, E_s represents the shape-preserving regularization energy, and V_0 represents the initial vertex set. Then, we define the term E_s by

$$E_s(V, V_0, E) = \sum_{T_k \in T} \left(\sum_{i, j \in T_k} w_{ij} \left| \|\mathbf{R}_k \mathbf{e}_{0,ij} - \mathbf{e}_{ij}\|^2 \right| \right) + \lambda_n a_{0,k} \left| \|\mathbf{R}_k \hat{\mathbf{n}}_{0,k} - \hat{\mathbf{n}}_k\|^2 \right|, \quad (11)$$

where T represents the set of mesh elements, including rims and spokes⁵¹ and faces, w_{ij} represents the cotangent weight of edge \mathbf{e}_{ij} ⁵⁴, $\mathbf{R}_k \in \text{SO}(3)$ represents the undetermined rotation matrix of element T_k , λ_n represents a positive regularization parameter, $a_{0,k}$ represents the element area, and finally, $\hat{\mathbf{n}}_k$ represents the unit normal vector of element T_k . Notably, rims/spokes and faces exhibit distinct impacts on the optimization process. Specifically, in rim-and-spoke structures, the higher density of overlapping vertices induces more pronounced deformation smoothing effects. In contrast, face-dominated structures contain fewer overlapping vertices, thereby better preserving localized geometric details during optimization. We will first employ the former approach to achieve coarse adjustments, followed by the latter for fine-tuning.

Clearly, deriving a closed-form solution for the nonlinear optimization problem shown in (11) is very challenging. Therefore, we adopt an iterative strategy, which alternates between the scattering optimization and shape-preserving regularization. To implement the latter, it is necessary to minimize Equation (11), which, however, is also quite difficult. To address this problem, we employ another iterative approach with two steps: local and global optimization. We show both of these two iterative processes in Supplementary Notes 3.

Results

In this section, we present a comprehensive evaluation of the Electromagnetic Sculptor framework. We begin by outlining the experimental setup and demonstrating the framework's baseline performance on standard monostatic and bistatic RCS reduction tasks (Subsection Experimental Setup and Baseline Validation). We then extend the evaluation to more complex scenarios, including 2D angular, broadband, and multi-polarization optimization, to verify the framework's versatility (Subsection Advanced Multi-Objective Optimization). To provide deeper insights, we analyze the underlying physical mechanisms through surface currents and compare our approach against existing heuristic and gradient-based methods (Subsection Physical Mechanism and Comparative Analysis). Furthermore, we investigate the critical relationship between sampling density and optimization consistency, along with hyperparameter sensitivity, to establish practical usage boundaries (Subsection Sampling Consistency and Parameter Sensitivity). Finally, we validate the simulation results through fabrication and compact range measurements of a 3D-printed prototype (Subsection Experimental Measurement and Fabrication).

Experimental Setup and Baseline Validation

We begin with a brief overview of the test conditions and data used during the validation phase. On the software side, our low-scattering optimization framework leverages several external libraries, including PyTorch 2.1.2 and CUDA 12.1. To enhance the efficiency of matrix inversion, we also integrated Eigen 3.4.0. All simulations were conducted on a Windows 11 operating system. On the hardware side, the testing platform consisted of a laptop equipped with 32GB of RAM, an Intel i9-13900H CPU, and an NVIDIA RTX 4060 GPU with 8GB of VRAM. The test geometries used in this study include models created by the authors (a sphere and an airplane) and standard publicly available benchmark models (the Stanford Bunny⁵¹ and CMU Spot⁵⁹, hereafter referred to as the bunny and the calf). The vast majority of these test models are constrained in a 2m×2m×2m cubic space, with an exception: the plane model, with dimensions of 4m in length, 4m in width, and 1.25m in height. All these models contain thousands of vertices. We generally control the size of the basic mesh elements to be less than 0.06m, and use Rhino software to

subdivide the elements that do not meet this requirement.

We conduct the tests to evaluate the RCS optimization performance under both monostatic and bistatic scenarios. Two models are selected for testing: a sphere and an airplane. Using the centroid of the target model as the origin, we define two sampling schemes. For the monostatic mode, the pitch angle is fixed at 90° , and we sample at 4° intervals in the azimuth plane, yielding 90 points. For the bistatic mode, with the transmitter fixed at (pitch= 90° , azimuth= 0°) and the receiver at (azimuth= 0°), we sample the pitch angle from 45° to 135° at 3° intervals, resulting in 30 points. The excitation frequency is set to 1 GHz and HH polarization is employed throughout the optimization processes. Additionally, the total number of epochs is set to 60. Under these conditions, we found that the optimization basically achieves convergence. For the first 30 epochs, the basic mesh elements for shape preservation are chosen as rims and spokes, while for the remaining 30 epochs, faces are used. Other parameters include an optimizer learning rate η of 0.1, with λ_n set to $1e3$, N_F set to 32, λ_g set to 20, σ_d and σ_ϕ set to 0.5 and 0.4, respectively. The latter four parameters are consistent with those in Ref. ⁵³, while the impacts of the first two parameters will be discussed subsequently.

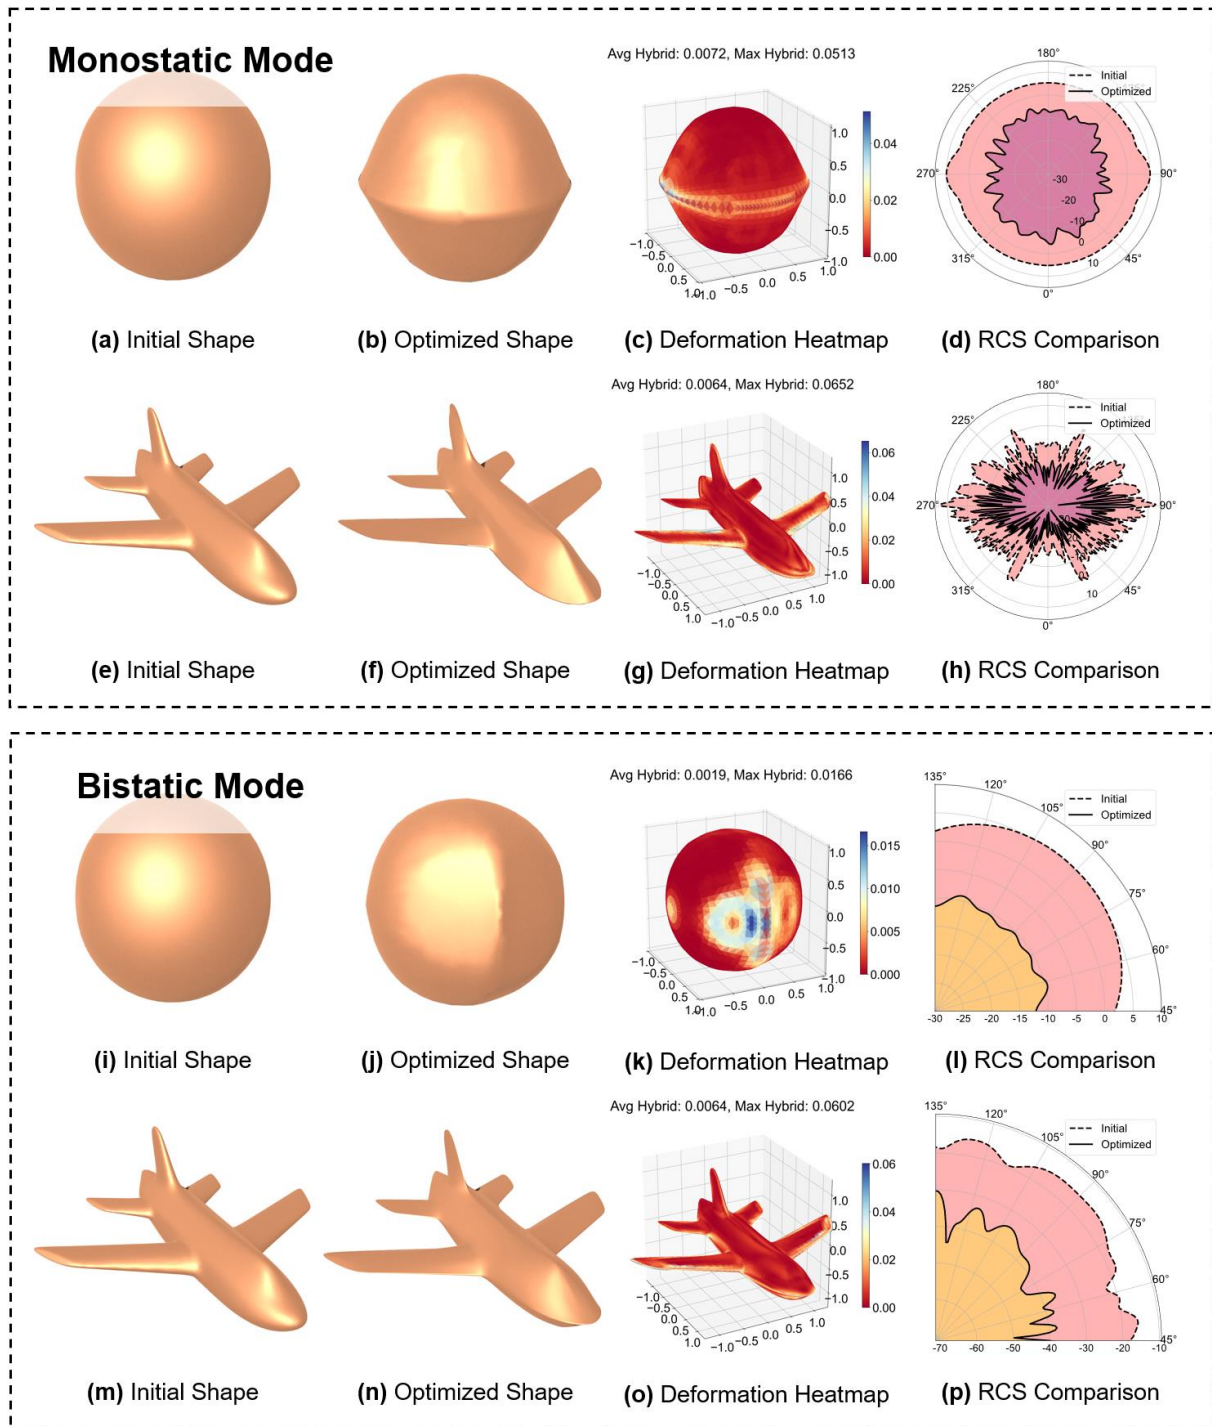


Fig. 2 Optimization results under monostatic and bistatic modes. (a, e, I, m) show the initial shape before optimization; (b, f, j, n) present the optimized low-scattering shape. We visualized the deformation quantitatively in (c, g, k, o), and demonstrated the radar cross section (RCS) comparison effects before and after optimization in (d, h, l, p).

Fig. 2(a-h) presents the optimization results for the monostatic RCS, while Fig. 2(i-p) displays

those for the bistatic RCS. Column-wise, Fig. 2 can be divided into four groups. The first two columns illustrate the model shapes before and after the optimization, respectively. It can be observed that the optimization sharpens the surfaces facing the illumination direction while flattens and smooths those in other directions. This effectively redirects the scattered energy toward less critical angles, which aligns with the expected characteristics of the low-scattering profiles. In the third column of Fig. 2, the deformation is quantitatively visualized using a hybrid metric, defined as the product of the Hausdorff distance⁵⁸ and the sine of normal vector deviation angle. This formulation captures both rigid transformations (e.g., rotation and translation) and non-rigid deformations (e.g., stretching). The distribution of this hybrid metric is consistent with the aforementioned description of surface sharpening and flattening. For instance, all of the protruding edges on the spherical surface in Fig. 2(c, k), the sharpened nose, flattened fuselage sides, wings, and empennage in Fig. 2(g, o) align with this pattern. The last column compares the RCS patterns before and after the optimization, with the sampling interval set to 1° for both. A clear RCS reduction is evident in the optimized results. Furthermore, Table I provides detailed data, including the number of mesh vertices, the average RCS before and after optimization, the computation time, and an efficiency metric. These results demonstrate the high effectiveness of our method, showing that satisfactory results can be achieved within just a few minutes, even for models with thousands of vertices and dozens of sampling points.

Advanced Multi-Objective Optimization

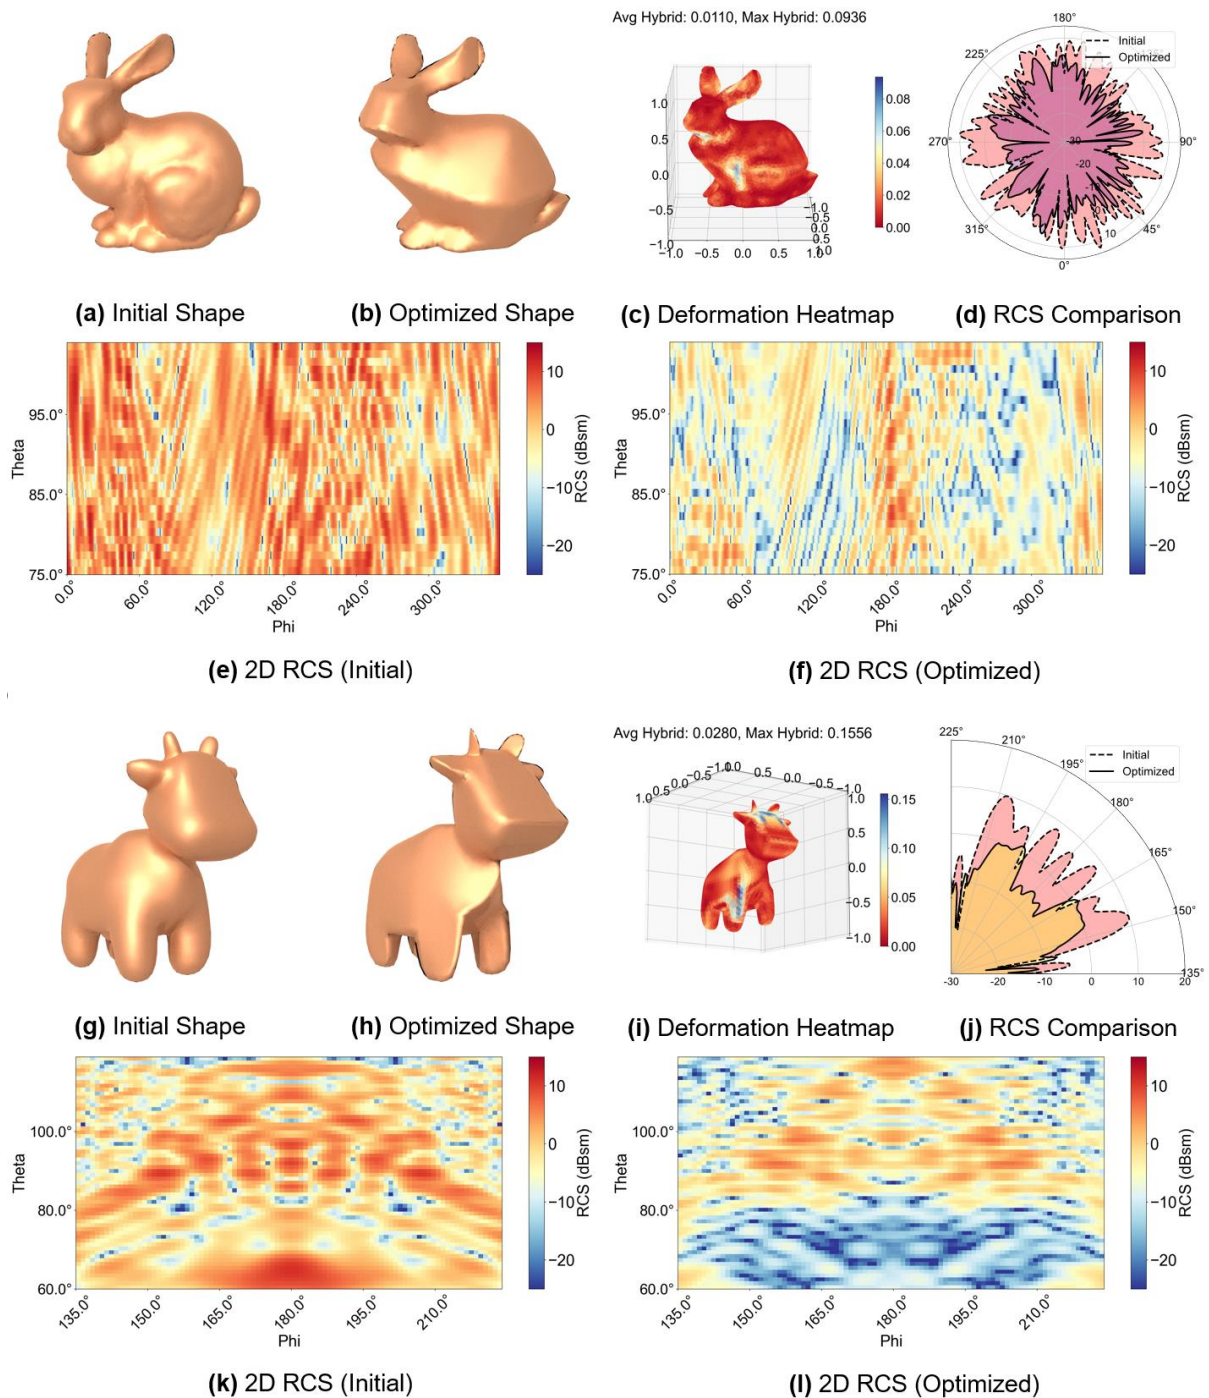


Fig. 3. 2D scattering optimization results. (a, g) show the initial shapes before optimization; (b, h) present the optimized low-scattering shapes; (c, i) display the corresponding deformation heatmaps; (d, j) illustrate the radar cross section (RCS) comparison before and after optimization on a representative slice; (e, k) and (f, l) respectively present the 2D RCS heatmaps before and after optimization.

We subsequently conducted 2D scattering optimization experiments on two models, the bunny and

the calf, both configured in the monostatic mode. For the bunny model, the optimization was performed over the full azimuth range, sampled at 6° intervals, and an elevation range from 75° to 105° , sampled at 5° intervals, resulting in 420 sampling points. The operating frequency was set to 1.5 GHz. For the calf model, the optimization region was defined with an azimuth range of 135° to 225° (sampled at 2° intervals) and an elevation range of 60° to 120° (sampled at 6° intervals), yielding 506 sampling points. The operating frequency was 2 GHz. All other configurations remain consistent with the previous setup. Fig. 3 presents the 2D RCS optimization results for the two models. Fig. 3(a, b, g, h) depicts the models before and after the optimization and Fig. 3(c, i) displays the hybrid deformation metric as heatmaps, quantifying the geometric differences. It can be observed that the bunny model undergoes clear changes in the mouth and lateral body regions, while the calf model exhibits deformation primarily concentrated in the head and leg areas. The per-face average of this hybrid metric is higher here than in Fig. 2(c, g, k, o) and the subsequent Fig. 3(c, k), indicating that the 2D optimization—aimed at satisfying angular requirements across multiple dimensions—induces more substantial deformation. The RCS slices at elevation angles of 80° and 90° are shown in Fig. 3(d, j), respectively, demonstrating average RCS reductions of 7.55 dB and 8.03 dB. The global RCS patterns are represented by the four 60×120 heatmaps in Fig. 3(e, f, k, l), which collectively exhibit effective RCS reduction across the observed domain. More detailed data are provided in Table II. Notably, the average RCS reduction reported here is somewhat smaller than the slice-specific values in Fig. 3(d, j). This discrepancy can be attributed to increased spatial anisotropy in the 2D optimization scenario. Further improvements could be achieved by increasing the sampling density or fine-tuning the parameters.

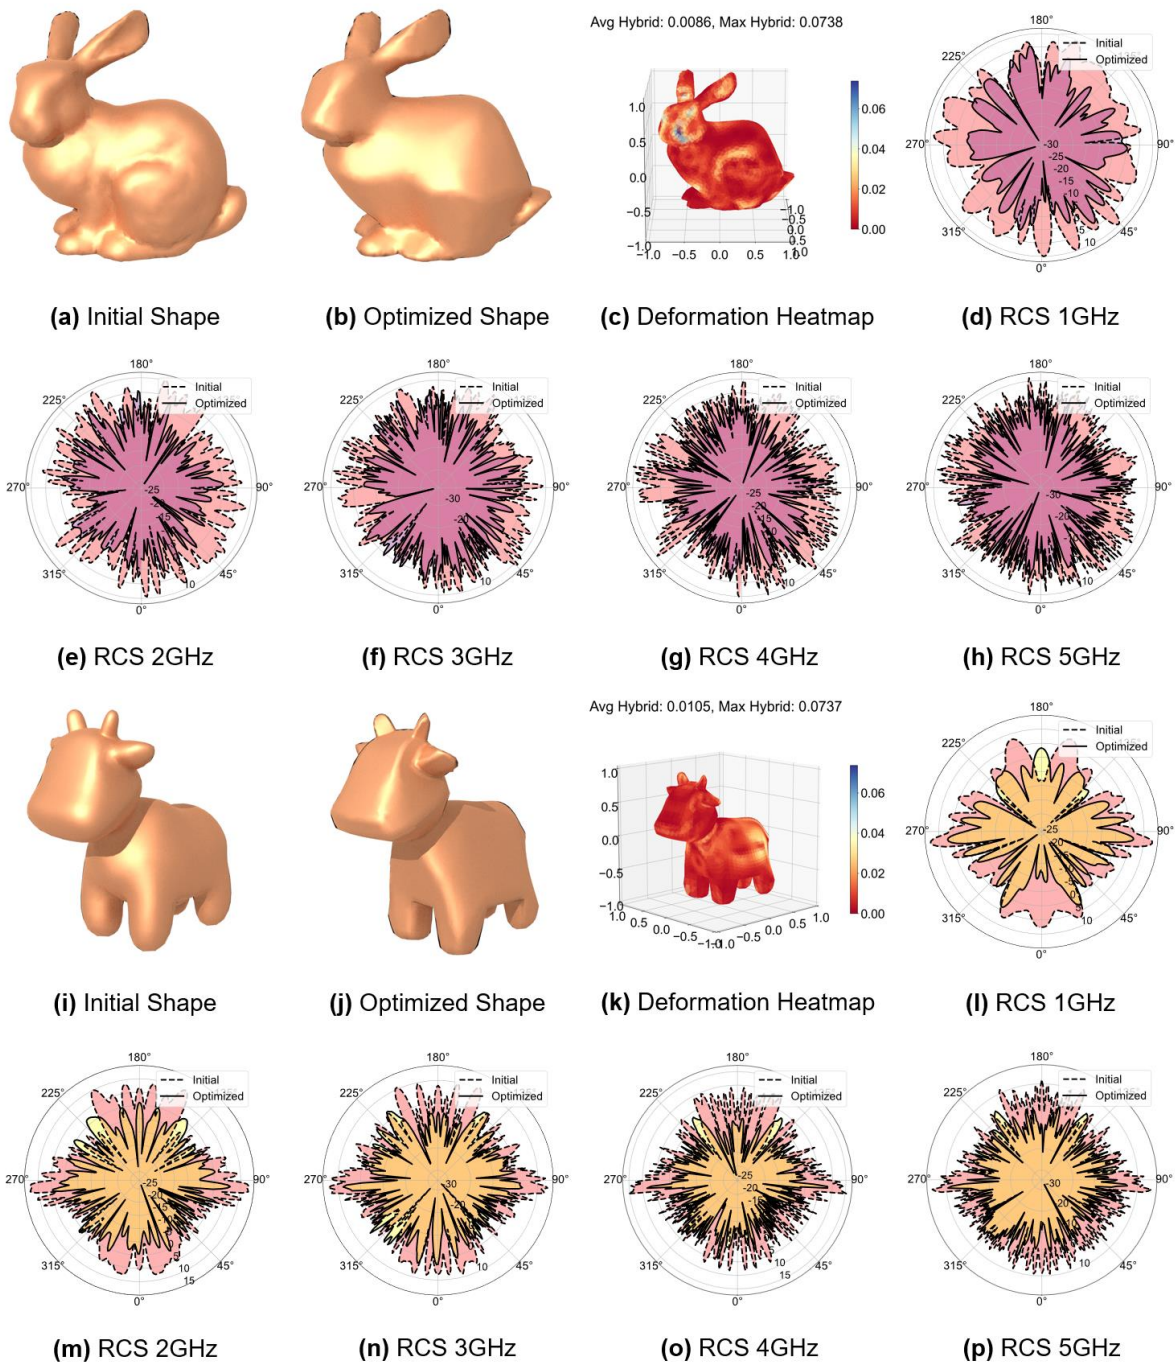


Fig. 4 Experimental results of broadband scattering optimization. (a, i) show the initial shapes before optimization; (b, j) present the optimized low-scattering shapes; (c, k) display the corresponding deformation heatmaps. (d-h) and (l-p) respectively illustrate the radar cross section (RCS) comparison results before and after optimization for the two models at the frequency range of 1–5 GHz.

To evaluate the broadband optimization performance of the framework, we select two models,

bunny and calf, and optimize the monostatic RCS as objective. We fix the excitation direction at 90° in pitch angle, set excitation directions every 4° in the azimuth angle range from 0 to 360° and employ HH polarization throughout the optimization processes. In the frequency domain, the excitation points are set every 1 GHz between 1 and 5 GHz, resulting in a total of 450 optimization sample points. Due to hardware platform limitations and the fact that our program had not been fine-tuned for optimization, performing 450 computations per epoch will lead to memory overflow. Hence, we adopt a mini-batch gradient descent strategy by dividing the one-epoch computations into 5 groups based on five frequencies. This effectively alleviates the memory pressure. In terms of other parameter settings, we keep parameters λ_n , N_F , λ_g , σ_d and σ_ϕ consistent with the previous configurations. For the bunny model, we set the optimizer learning rate η to 0.2, while for the calf model, the learning rate $\eta=0.1$. Fig. 4 shows the simulated results for the broadband monostatic RCS optimization task, including the shapes before and after optimization, deformation heatmaps and RCS plots vary from 1 to 5 GHz. Table III provides the average monostatic RCS values for the two models before and after optimization at different frequencies. The optimization results for such two models demonstrate that the optimization framework achieves approximately 6dB or more of average RCS reduction in the frequency range from 1 to 5 GHz. This indicates that the proposed framework has good optimization capabilities for low scattering in broadband.

As one of the most fundamental properties of electromagnetic field, polarization directly governs the resulting electromagnetic response. In our previous work, all evaluations were conducted under the HH polarization, which naturally raises the question of whether the optimized results remain effective under other polarization states. To investigate this, we performed additional experiments on the calf, bunny, and sphere models to evaluate their scattering performance across multiple polarization configurations. Specifically, the calf model was tested in monostatic mode with an excitation frequency of 1 GHz and the optimization region set to the azimuthal plane; the bunny model has the same experimental configuration as shown in Fig. 4; and the sphere model follows the bistatic experimental setup described in Fig. 2. As summarized

in Table SI in the Supplementary Notes 4, all models maintain consistently low RCS under both HH and VV co-polarization conditions. We believe that this may be related to the fact that the polarization mode has a relatively smaller influence on the scattered energy compared to the factors such as space and frequency. In contrast, such consistency is not observed under the cross-polarization. Even so, the corresponding scattering energy remains very weak and practically negligible. Further validation of VV-polarized optimization (also in Table SIII) confirms the same trend. Such results indicate that the resource allocation for polarization can be appropriately reduced during the optimization process, thus achieving consistent reduction of the scattered energy under co-polarization in a more cost-effective manner.

Physical Mechanism and Comparative Analysis

Previously, we visualized the optimization process by demonstrating the shape changes. We observe that as the optimization progresses, the scatterer's shape exhibits characteristics of surface flattening and smoothing. This aligns with our expectations, as such changes effectively alter the strong scattering direction of electromagnetic waves, suppressing the electromagnetic scattering in the observation direction and ultimately leading to a decrease in the average RCS (though some areas may show enhancement, the overall trend is reduction). Of course, this is a trend observed from a geometric perspective. Now, we aim to explain the changes in the scattering properties of objects during optimization from a more physical perspective. To achieve this, we firstly map the equivalent current density on the scatterer's surface, with the help of FEKO⁵⁶. Furthermore, we also incorporated with the hotspot maps, a highly effective tool for electromagnetic scattering analysis, calculated by FASTEM STUDIO (now designated as NESCS)⁵⁷. Composed of normalized electromagnetic energy carried by rays, these maps can reflect the contribution of each surface element to the RCS, thus addressing the limitation of equivalent current distribution maps (i.e., their irrelevance to the observation direction). We selected two models mentioned earlier (the calf, and a sphere) for monostatic and bistatic cases respectively. The calf was excited by a 1 GHz plane wave from two directions: (azimuth=0°, pitch=90°) and (azimuth=90°, pitch=90°), with HH

polarization mode. The sphere was excited from azimuth=0°, pitch=90° by the same frequency and polarization, and its scattering field was received at azimuth=0°, pitch=60°.

Fig. 5 presents the evolution of the transient surface equivalent current distribution and hotspot maps throughout the optimization process. A notable observation is the compression of major co-directional current regions on the scatterers, which are considered as the key contributors to strong electromagnetic scattering. Examples include the hip in Fig. 5(a2-a5), the abdomen and leg in Fig. 5(b2-b5), and the central area in Fig. 5(c2-c5). Conversely, the reverse current regions are expanded, promoting a better balance in electromagnetic scattering between co- and reverse-directional areas on the surface. The hotspot maps in Fig. 5(a7-a10, b7-b10, c7-c10) offer an intuitive visualization on how different regions contribute to the scattered field. A gradual attenuation in the normalized energy of the aforementioned regions can be observed, macroscopically contributing to the reduction in RCS. This trend is consistent with the RCS bar charts provided under specific excitation-observation angles (Fig. 5(a6, b6, c6)), which demonstrate a clear suppression of electromagnetic scattering post-optimization. However, it should be noted that hotspot maps do not account for interference effects between surface elements, thus limiting their accuracy in depicting scattering energy variations. For instance, the hotspot map suggests enhanced scattering from the head of the first calf, which may be misleading. This highlights the limitation of interpretations based solely on local energy considerations, which are often insufficient to capture interference effects.

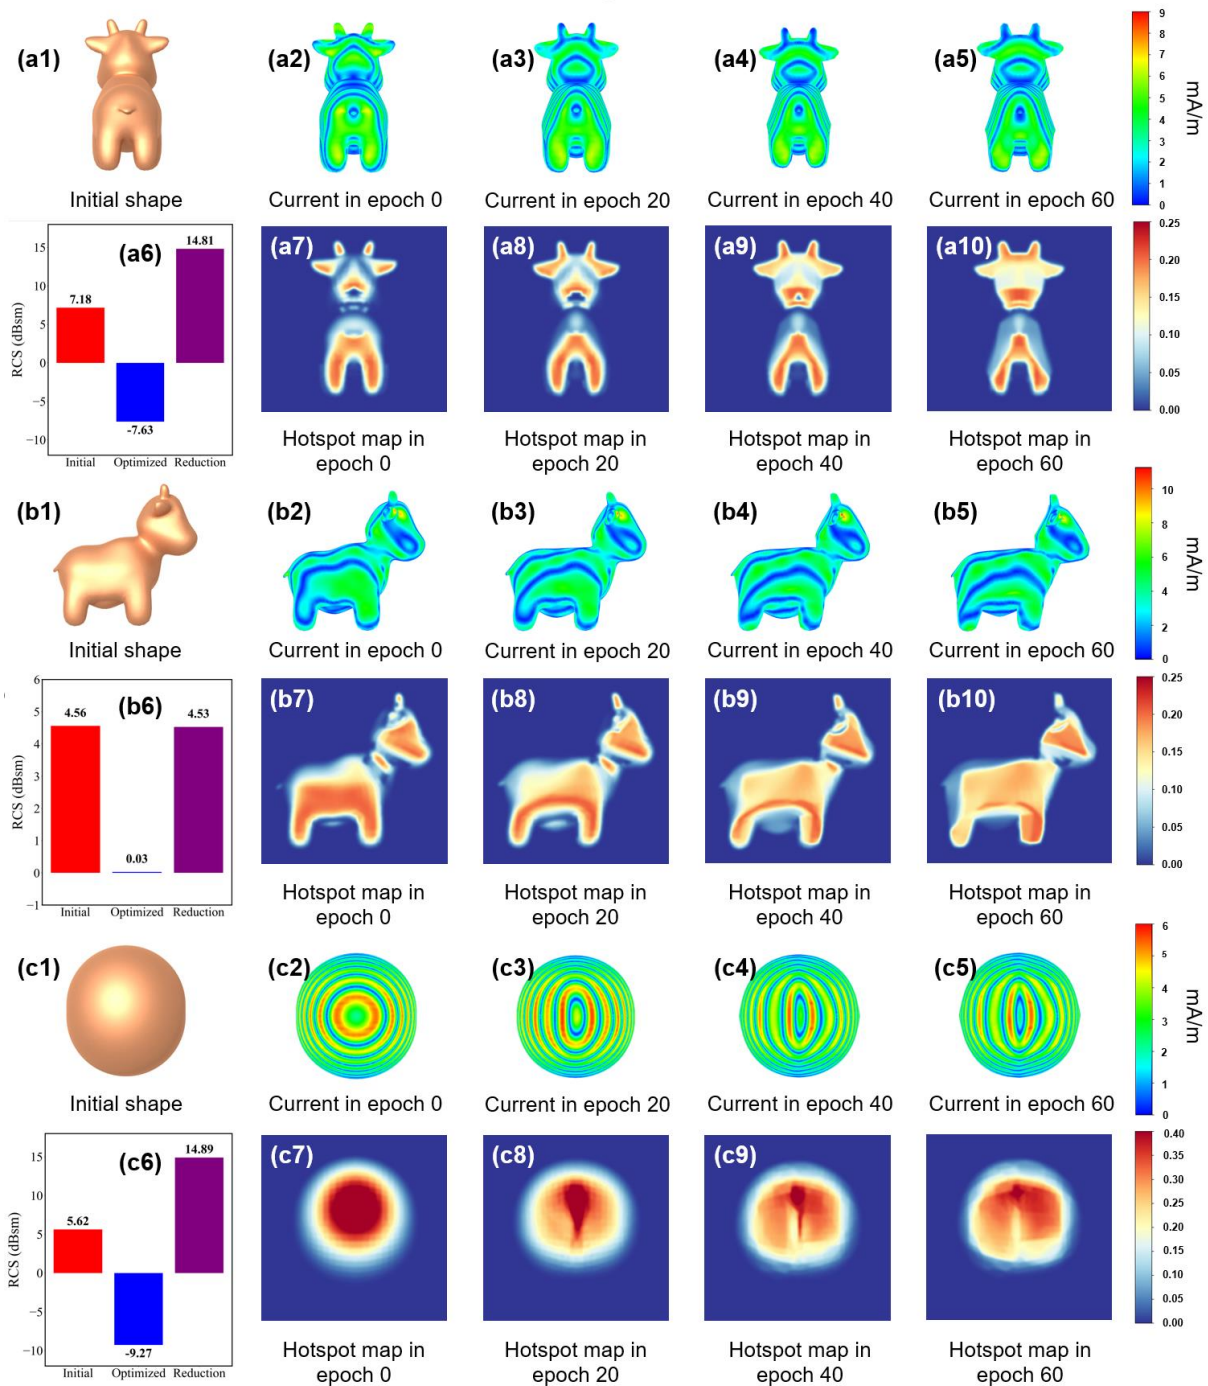


Fig. 5 Evolution of equivalent current distributions and hotspot maps during the optimization process. The figure is organized into three representative cases: a, b and c. For each case, (1) shows the rendered geometry under the given excitation angle, (2–5) illustrate the transient equivalent current distributions at different optimization epochs, (6) reports the corresponding radar cross section values under the specified excitation–observation configuration, and (7–10) present the associated hotspot maps at selected epochs.

To investigate the comparative advantages of our proposed method, we selected several alternative feasible solutions as the control groups and designed three sets of comparative experiments for the broadband monostatic RCS optimization problem illustrated in Fig. 4. Control Group 1 employs the Grey Wolf Optimizer (GWO)^{11,15}, a widely adopted heuristic approach with strong global search capability compared with classical genetic algorithms⁹ and particle swarm optimization¹⁰. Control Group 2 adopts a naive gradient descent strategy commonly used in electromagnetic optimization studies^{1,4,5,35-39}, with shape-preserving regularization applied. Building on Control Group 2, Control Group 3 further incorporates an optical energy-based optimization strategy⁵¹, conceptually related to the hotspot-based evaluation discussed above and implemented in the H-plane. We put the results in the Supplementary Notes 5. As shown in Fig. S1, the proposed method achieves lower RCS while better preserving geometric integrity than the baseline approaches. In contrast, Control Groups 1 and 2 exhibit pronounced geometric distortion that undermines optimization stability, while Control Group 3 maintains improved geometric smoothness but yields inferior RCS reduction over the 1–5 GHz range. These results confirm that strategies relying primarily on local energy metrics are insufficient, whereas the proposed framework more effectively exploits interference effects to achieve consistent broadband RCS reduction.

We have observed that all the optimization results above exhibit a trend toward surface flattening and smoothing of the model. This suggests that there may be a more fundamental physical intuition guiding the electromagnetic scattering optimization process. To explore this, we used graph-based low-pass filtering to simulate such flattening and smoothing operations, with detailed experimental procedures included in the Supplementary Notes 6. The results indicate that applying physical intuition (such as the low-pass filtering used here) can indeed reduce RCS in some cases, but this strategy is not universally effective. However, the optimization methods provide a more flexible and comprehensive solution; or, from another perspective, the optimization determines the regions that require or do not require smoothing.

Sampling Consistency and Parameter Sensitivity

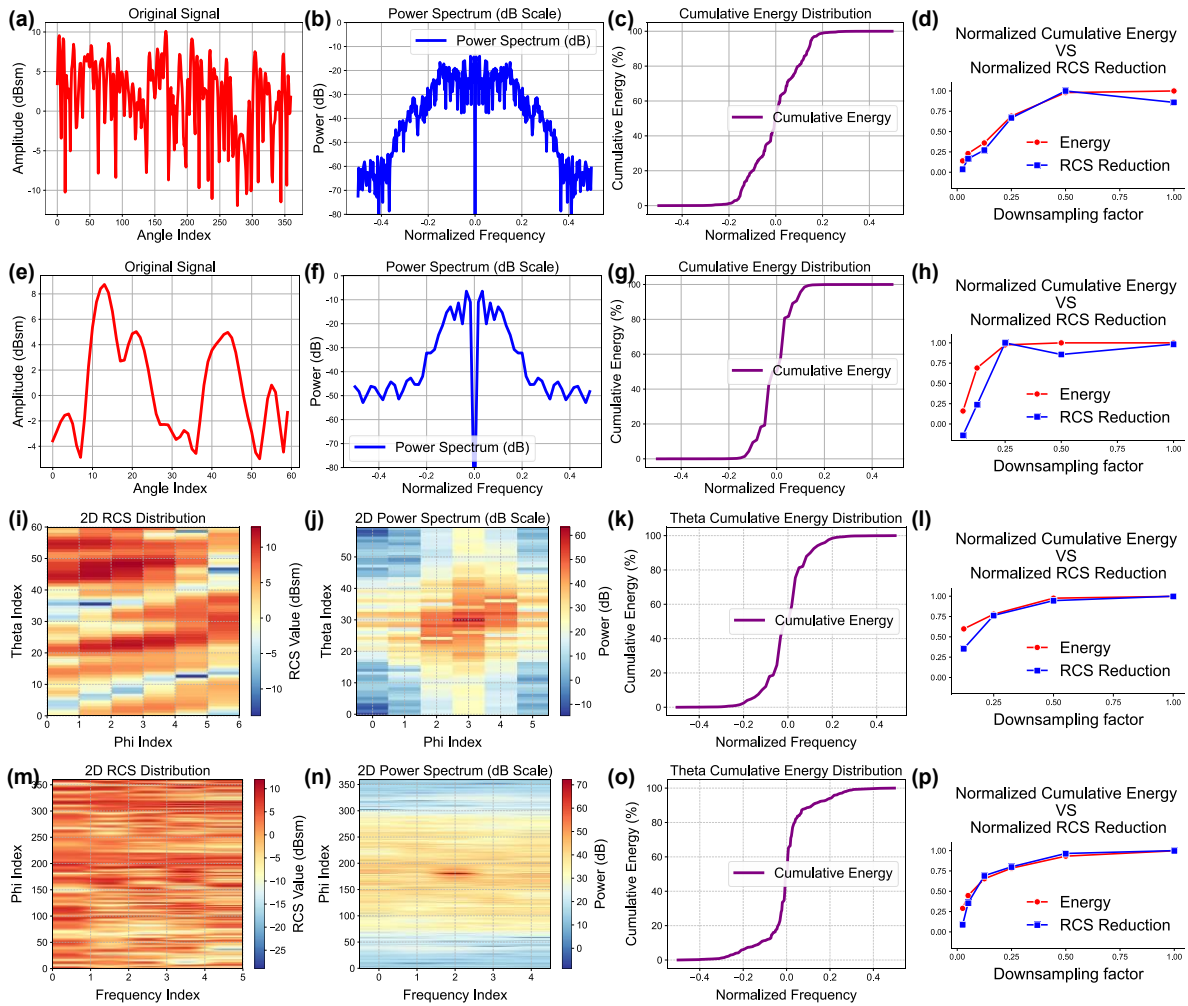


Fig. 6 Relationship between sampling and optimization consistency. Panels (a–d) correspond to the first experimental group, (e–h) to the second group, (i–l) to the third group, and (m–p) to the fourth group. For each group, (a, e, i, m) show the initial radar cross section (RCS) signals, (b, f, j, n) present the corresponding RCS power spectra, (c, g, k, o) display the cumulative energy distribution curves, and (d, h, l, p) illustrate the variation of cumulative energy and normalized RCS reduction (normalized by the maximum reduction) with respect to the downsampling rate.

Our previous experiments demonstrate the effectiveness of the proposed approach in manipulating electromagnetic waves to reduce scattering energy. We now address a fundamental question that determines the practical boundaries of the method: whether optimization in specific directions may lead to increased scattering intensity in other directions. This issue must be analyzed on a case-by-case basis. If such directions lie outside the optimized region, the proposed method cannot

guarantee consistent behavior, as the optimization process lacks information from external regions and therefore cannot achieve global coordination. This phenomenon is illustrated by a representative example shown in Fig. S3 of the Supplementary Notes 7. In contrast, when the directions correspond to unoptimized areas within the target region, strong optimization consistency can still be achieved. This observation is attributed to the inherent local correlation of electromagnetic scattering characteristics, whereby the scattering energy distribution tends to be locally stable. Consequently, ensuring that the sampling points sufficiently cover these stable regions is key to achieving consistent RCS reduction.

To investigate this effect, we conducted four simulation experiments using the bunny model, corresponding to the four experimental groups shown in Fig. 6(a–p). Experiments 1, 3, and 4 were performed in a monostatic configuration, while Experiment 2 employed a bistatic configuration. In Experiment 1, the frequency was fixed at 2 GHz with a constant elevation angle of 90° , and the azimuth angle spanned from 0° to 360° (Fig. 6(a–d)). Experiment 2 was conducted at 2 GHz with a fixed azimuth angle of 270° , while the elevation angle varied from 60° to 120° (Fig. 6(e–h)). Experiment 3 was also conducted at 2 GHz, where downsampling was applied only along the elevation angle, while the azimuth angle was restricted to a narrow window from -3° to 3° with seven sampling points (Fig. 6(i–l)). Experiment 4 swept the frequency from 1 to 3 GHz using five frequency points, while maintaining the same angular scanning protocol as Experiment 1 (Fig. 6(m–p)). All experiments employed HH polarization.

As shown in Fig. 6(a, e, i, m), we present the angular RCS signals of the initial shape. The corresponding power spectra are shown in Fig. 6(b, f, j, n), while the cumulative energy distribution curves are given in Fig. 6(c, g, k, o). The variation of cumulative energy and normalized RCS reduction with respect to the downsampling factor is summarized in Fig. 6(d, h, l, p). The power spectrum is obtained by computing the squared magnitude of the discrete Fourier transform of the angular RCS signal after removing the DC component, which corresponds to the mean RCS value. For both one-dimensional and two-dimensional cases, the Fourier transform is applied only along the angular dimension; the two-dimensional spectrum is constructed by

stacking angular spectra obtained at different secondary sampling indices.

A 1° angular sampling interval (corresponding to a downsampling factor of 1.0) is adopted as a practical baseline, as this sampling density already ensures that scattering energy leakage into unseen directions is small, enabling stable evaluation under the current configuration. When evaluating RCS reduction, the results are normalized by the maximum value observed over the downsampling factor range from 0 to 1. Although strong optimization performance is achieved at the sampled angles, insufficient sampling may cause increased scattering in unsampled directions, leading to a reduction in the overall RCS reduction.

Our results reveal a strong positive correlation between RCS reduction and cumulative sampling energy, as illustrated in Fig. 6(d, h, l, p). Notably, when the cumulative energy exceeds approximately 90%, the RCS reduction approaches saturation under the current optimizer configuration. These results indicate that consistent optimization performance can be achieved when the sampling density is compatible with the intrinsic spatial-frequency characteristics of the target's scattering response. In practice, this requires that the sampling range sufficiently covers the dominant energy components of the RCS power spectrum.

We now turn to the selection of parameters η and λ_n . Here, η controls the vertex update scale, while λ_n regulates the influence of the initial shape versus the low-scattering shape on the final result. Increasing η enhances RCS reduction but also amplifies geometric distortion and convergence oscillations when η exceeds 0.2. Accordingly, η is kept below 0.2 in our experiments. In contrast, λ_n exhibits a milder influence: sufficiently large values ($\lambda_n \geq 1e3$) lead to stable convergence and consistent deformation, whereas overly small λ_n suppress effective low-scattering optimization. Detailed parameter sweeps and visual results are provided in the Supplementary Notes 8.

Experimental Measurement and Fabrication

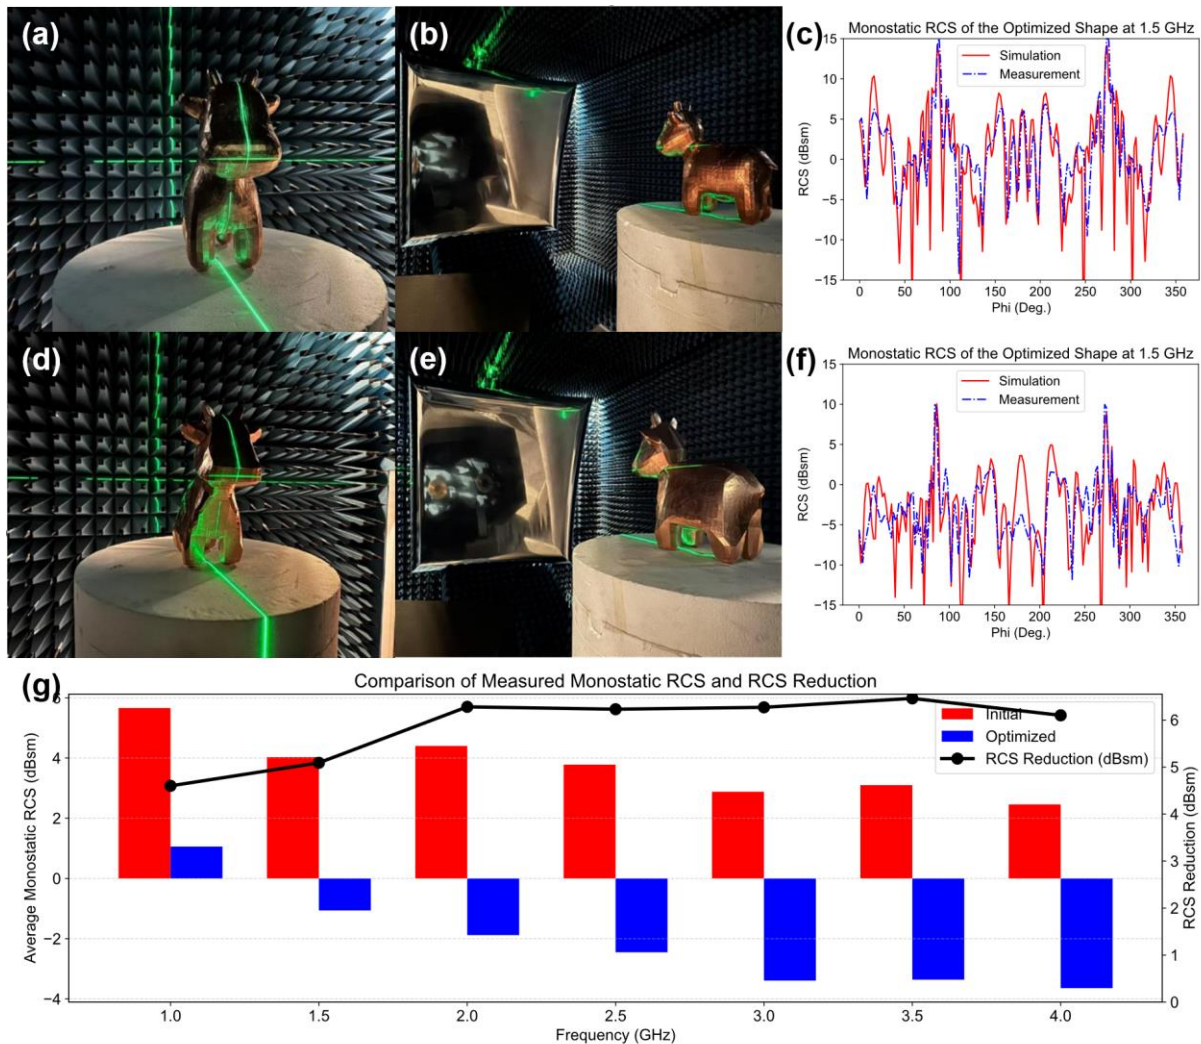


Fig. 7 Measured radar cross section (RCS) optimization results for the calf model. (a, b) Measurement setup of the initial model. (d, e) Measurement setup of the optimized model. (c) Simulation-measurement comparison of the initial model at the frequency of 1.5 GHz. (f) Simulation-measurement comparison of the optimized model at the frequency of 1.5GHz. (g) Frequency-domain RCS and RCS reduction (angle-averaged) in the frequency band from 1 to 4GHz.

Finally, to validate the effectiveness of the proposed framework for broadband (1-4GHz) and wide-angle (0-360°) monostatic RCS optimization, we conduct experiments using the calf model as test object. The baseline and optimized model obtained from our simulation studies are fabricated using 3D printing technology. All measurements are carried out in a compact antenna test range in the State Key Laboratory of Millimeter Waves, Southeast University. Due to spatial constraints (50cm diameter turntable), quarter-scale models are fabricated to accommodate the test environment. To

achieve proper PEC boundary conditions, the resin models are fully coated with copper foil (to be precise, copper foil tape) to prevent electromagnetic field penetration. In accordance with the scaling theorem⁵⁵, the actual measurement frequency range is proportionally increased to 4-16GHz. During the measurement, HH polarization is employed.

As shown in Fig. 7, the experimental process involved two separate measurement campaigns: (a) and (b) present the test configuration for the baseline model, while (d) and (e) show the setup for the optimized model. Measurements are conducted with 2° azimuthal increment and 2GHz frequency step. Fig. 7(c, f) presents the comparison between the simulated and measured RCS results. Overall, they show good consistency, demonstrating that the high-frequency approximation method adopted in this study can effectively capture the key electromagnetic characteristics of the scatterer while improving computational efficiency. However, we have also noted the discrepancies therein, which may be associated with factors including both modeling and measurement issues. A detailed discussion of the differences will be provided in the Supplementary Notes 10. Fig. 7(g) presents the angle-averaged RCS measurements (in dBsm) across broad frequency band. The results show consistent RCS reduction values of 4.60, 5.08, 6.29, 6.24, 6.35, 6.47, and 6.09 dBsm at 0.5GHz intervals from 1 to 4GHz, with an average reduction approximating 6dBsm. Compared to Table III, the mean deviation of 0.85dBsm (<1dBsm) falls in acceptable experimental error margins considering the systematic uncertainties in model fabrication and measurement, thereby validating both measurement consistency and effectiveness of this framework.

Discussion

It should be noted that all cases presented above are based on the condition of $\sigma' = 0$ in Eq. (3). This naturally raises the question of whether our method can be applied to more general electromagnetic design problems. The answer is affirmative in principle: our approach can be effectively extended to handle non-zero σ' scenarios, and is even capable of adapting to other types of electromagnetic data beyond the formulation in Eq. (3). Representative examples are

provided in the Supplementary Notes 11.

While these experiments demonstrate the promising potential of our method for the electromagnetic scattering optimization, several limitations should be acknowledged. For instance, the measurement results indicate that there remains room for improvement in the current forward modeling configuration. As previously noted, the optimization process tends to sharpen the edges of the scatterer geometry, which enhances diffraction effects. However, for the sake of simplicity, a diffraction model was not incorporated into our formulation. As a result, the accuracy of our forward modeling gradually declines over the course of optimization—as illustrated in the Supplementary Notes 11—leading to less-than-expected RCS reduction in practice. Furthermore, we observe that the optimization outcome is closely tied to the sampling density (as analyzed in Fig. 6). This implies that for higher-frequency problems or scenarios involving electrically larger structures, substantial computational resources are required. Such demands inevitably increase computational cost and may introduce instability in gradient-based optimization. Additionally, the shape preservation constraint itself imposes a limitation, confining the optimized results to remain close to the initial geometry and thereby excluding other potentially better configurations. Lastly, although the current method can be applied to arbitrary meshes, it does not ensure that the optimized geometry exhibits regularity or conforms to the CAD design standards—a critical consideration for practical manufacturing. These issues outline several natural directions for future work to complement and enhance the existing framework. Potential directions include incorporating a differentiable diffraction model, reducing sampling density, and introducing CAD-oriented constraints, among others.

Conclusion

We proposed the Electromagnetic Sculptor, an efficient geometry optimization framework for addressing the electromagnetic field manipulation. In contrast to the traditional empirical and heuristic optimization methods, our optimization framework is based on gradient descent, and hence leads to a clear improvement in optimization efficiency. This allows for a dramatic increase

in the number of optimization parameters and further extends the scope of optimization from small parameter models to more general mesh structures, thereby greatly expanding the application scenarios of automated electromagnetic design.

To build a feasible electromagnetic geometry optimization framework, we implemented two key modules. Firstly, we developed a differentiable SBR electromagnetic simulator based on parallel computing and automatic differentiation techniques. This effectively bridges the gap between electromagnetic simulation and optimization, allowing parameter gradients to be backpropagated along the forward computation path for efficient and reliable large-scale gradient calculation. Secondly, we incorporated a gradient-based optimizer equipped with a Fibonacci lattice-based spatial gradient filter. This ensures that the vertex gradients propagate through the connections between edges to local adjacent points, leading to more stable gradient descent and overcoming the sensitivity of vertex gradients to mesh shape. To evaluate the effectiveness of the Electromagnetic Sculptor framework, we conducted a RCS optimization experiment. For this application, we introduced a shape preservation module based on the computational graphics theory. This module minimizes the changes in the mesh model's shape while optimizes the scattering energy, imposing constraints on the gradient descent process to ensure that the optimization results remain physically reasonable. Finally, through a series of simulation and measurement results, we showed that the proposed low-scattering optimization framework performs well. It effectively supports low-scattering optimization designs under various requirements. Additionally, we further establish that consistent electromagnetic optimization requires the sampling density to be compatible with the intrinsic variation of the target's scattering characteristics. We believe that this framework will play an active role in the future development of flexible electromagnetic designs, broader adoption of differentiable electromagnetics, and multi-physics simulation technologies.

We note that the current framework still has several limitations that warrant further investigation. Firstly, the lack of a diffraction model reduces forward modeling accuracy as geometric sharpening intensifies during optimization, limiting achievable RCS performance.

Secondly, the method's reliance on dense sampling leads to high computational costs and instability in the gradient-based optimization, particularly for high-frequency or large-scale problems. Thirdly, while the shape preservation constrains the design space, it may also preclude the potentially superior configurations. Finally, the optimized geometries, while physically reasonable, may not directly satisfy standard CAD manufacturability requirements. Therefore, extending this approach still remains an open research direction.

Data Availability

The RCS measurement data used in this research has been uploaded to the Supplementary Data. Additionally, the RCS measurement data and mesh files used can be accessed via <https://github.com/yankaiqiao/EM-Sculptor/>.

Code Availability

The code used in this study is available from the corresponding authors upon reasonable request.

References

1. Dainese, P., Marra, L., Cassara, D. et al. Shape optimization for high efficiency metasurfaces: theory and implementation. *Light Sci. Appl.* 13, 300 (2024).
2. Liu, Z., Zhu, D., Rodrigues, S. P. et al. Generative Model for the Inverse Design of Metasurfaces. *Nano Lett.* 18, 6570-6576 (2018).
3. Fan, Y., Yang, C., Sun, Y. et al. "Hybrid Channel Tracking for THz Massive MIMO Communication Systems in Dynamic Environments," *IEEE J. Sel. Areas Commun.* 43, 988-1003 (2025).
4. Li, M. Bai, J., Li, L. et al. A gradient-based aero-stealth optimization design method for flying wing aircraft. *Aerosp. Sci. Technol.* 92, 156-169 (2019).
5. Li, M., Chen, J., Feng., X. et al. An efficient adjoint method for the aero-stealth shape optimization design. *Aerosp. Sci. Technol.* 118, 107017 (2021).
6. Chen, Y., Zhang, H., Ma, J. et al. Semantic–Electromagnetic Inversion With Pretrained Multimodal Generative Model. *Adv. Sci.* 11, 2406793 (2024).
7. Yang, L., Jiang, Y., Li, Y., et al. DPO-SDF: Differentiable Physical Optics Method With Signed Distance Field for 3-D Reconstruction of Perfect Electric Conductors. *IEEE Trans. Antennas Propag.* 72, 1759-1771 (2024).
8. Fan, Y., Yang, C., Yuan, R. et al. "Dynamic Delay-Doppler-Angle Domain Channel Tracking for THz Massive MIMO-OTFS Communication Systems," *IEEE Trans. Wireless Commun.* 24, 5163-5178 (2025).
9. Forrest, S. Genetic Algorithms: Principles of Natural Selection Applied to Computation. *Science.* 261,872-

- 878 (1993).
10. Banks, A., Vincent, J. & Anyakoha, C. A review of particle swarm optimization. Part I: background and development. *Nat. Comput.* 6, 467–484 (2007).
 11. Faris, H., Aljarah, I., Al-Betar, M. A. et al. Grey wolf optimizer: a review of recent variants and applications. *Neural Comput. & Applic.* 30, 413–435 (2018).
 12. Qing, A., Lee, C. K. & Jen, L. Electromagnetic inverse scattering of two-dimensional perfectly conducting objects by real-coded genetic algorithm. *IEEE Trans. Geosci. Remote Sens.* 39, 665-676 (2001).
 13. Salucci, M., Poli, L., Rocca, P., et al. Learned Global Optimization for Inverse Scattering Problems: Matching Global Search With Computational Efficiency. *IEEE Trans. Antennas Propag.* 70, 6240-6255 (2022).
 14. He, Z., Li, Y. S., Gu, P. F. et al. A Space-Mapping-Based Optimal EM Design of RCS Reduction for Electrically Large Targets. *IEEE Trans. Antennas Propag.* 69, 6702-6711 (2021).
 15. Wang, H., Zou, Q., & Lin, H. A Quasi-Optimal Shape Design Method for Electromagnetic Scatterers Based on NURBS Surfaces and Filter-Enhanced GWO. *IEEE Trans. Antennas Propag.* 71, 4236-4245 (2023).
 16. Elsaywy, M. M. R., Lanteri, S., Duvigneau, R. et al. Numerical Optimization Methods for Metasurfaces. *Laser & Photonics Rev.* 14, 1900445 (2020).
 17. Ji, W., Chang, J., Xu, HX. et al. Recent advances in metasurface design and quantum optics applications with machine learning, physics-informed neural networks, and topology optimization methods. *Light Sci. Appl.* 12, 169 (2023).
 18. Jin, J., Feng, F., Zhang, J. et al. Efficient EM Topology Optimization Incorporating Advanced Matrix Padé Via Lanczos and Genetic Algorithm for Microwave Design. *IEEE Trans. Microwave Theory Tech.* 69, 3645-3666 (2021).
 19. Cheng, M., Zhao, X., Dhimish, M. et al. A Review of Data-Driven Surrogate Models for Design Optimization of Electric Motors. *IEEE Trans. Transp. Electrification.* 10, 8413-8431 (2024).
 20. Yang, Y., Yu, W. M., & Cui, T. J. Efficient RCS Modeling With an Adaptive Design-Based Gaussian Process Method. *IEEE Antennas Wirel. Propag. Lett.* 23, 2006-2010 (2024).
 21. Elsaywy, M. M. R., Lanteri, S., Duvigneau, R. et al. Global optimization of metasurface designs using statistical learning methods. *Sci. Rep.* 9, 17918 (2019).
 22. Rezaee Jordehi A, Jasni J. Parameter selection in particle swarm optimisation: a survey. *J. Exp. Theor. Artif. Intell.* 25, 527-542 (2013).
 23. Jing, G., Wang, P., Wu, H. et al. Neural network-based surrogate model for inverse design of metasurfaces. *Photon. Res.* 10, 1462-1471 (2022).
 24. Zhang, Q., Liu, C., Wan, X. et al. Machine-Learning Designs of Anisotropic Digital Coding Metasurfaces. *Adv. Theory Simul.* 2, 1800132 (2019).
 25. Wang, Q., Pang, Z., Gao, D. et al. Machine Learning-Assisted Quasi-Bisection Method for Pixelated Patch Antenna Bandwidth Optimization. *IEEE Antennas Wirel. Propag. Lett.* 23, 4807-4811 (2024).
 26. Wu, Q., Chen, W., Yu, C. et al. Machine-Learning-Assisted Optimization for Antenna Geometry Design. *IEEE Trans. Antennas Propag.* 72, 2083-2095 (2024).
 27. Tanriover, I., Lee, D., Chen, W. Deep Generative Modeling and Inverse Design of Manufacturable Free-Form Dielectric Metasurfaces. *ACS Photonics.* 10, 875-883 (2023).
 28. Li, L., Wang, L. G., Teixeira, F. L. et al. DeepNIS: Deep Neural Network for Nonlinear Electromagnetic

- Inverse Scattering. *IEEE Trans. Antennas Propag.* 67, 1819-1825 (2019).
29. Zhang, H., Chen, Y., Wang, Z. et al. Semantic regularization of electromagnetic inverse problems. *Nat. Commun.* 15, 3869 (2024).
 30. Raissi, M., Perdikaris, P., Karniadakis, G. E. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *J. Comput. Phys.* 378, 686-707 (2019).
 31. Shan T. et al. Physics-Informed Supervised Residual Learning for Electromagnetic Modeling. *IEEE Trans. Antennas Propag.* 71, 3393-3407 (2023).
 32. Seo, D., Nam, D. W., Park, J. et al. Structural Optimization of a One-Dimensional Freeform Metagrating Deflector via Deep Reinforcement Learning. *ACS Photonics.* 9, 452-458 (2022).
 33. Jiang, H., Khoo, Y. & Yang, H. Reinforced Inverse Scattering. *SIAM J. Sci. Comput.* 46, 884-902 (2024).
 34. Yang, T., Cheng, D. & Wang, Y. Designing freeform imaging systems based on reinforcement learning. *Opt. Express* 28, 30309-30323 (2020).
 35. López, D. P. Programmable Integrated Silicon Photonics Waveguide Meshes: Optimized Designs and Control Algorithms. *IEEE J. Select. Topics Quantum Electron.* 26, 1-12 (2020).
 36. Lalau-Keraly, C. M., Bhargava, S., Miller, O. D. et al. Adjoint shape optimization applied to electromagnetic design. *Opt. Express* 21, 21693-21701 (2013).
 37. T. Qiao et al. Pixel Antenna Optimization Using the Adjoint Method and the Method of Moving Asymptote. *IEEE Trans. Antennas Propag.* 71, 2873-2878 (2023).
 38. Bahmani, S., Evlyukhin, A. B., Hassan, E. et al. Topology optimization of optical nanoantennas with desired multipoles. *Opt. Express* 33, 19418-19441 (2025).
 39. Sun, A., Xing, S., Deng, X. et al. Edge-guided inverse design of digital metamaterial-based mode multiplexers for high-capacity multi-dimensional optical interconnect. *Nat. Commun.* 16, 2372 (2025).
 40. Gongal-Reddy, V. M. R., Feng, F., Zhang, C. et al. Parallel Decomposition Approach to Gradient-Based EM Optimization. *IEEE Trans. Microwave Theory Tech.* 64, 3380-3399 (2016).
 41. Abadi M, Barham P, Chen J, et al. TensorFlow: a system for large-scale machine learning. *Proceedings of the 12th USENIX conference on Operating Systems Design and Implementation (OSDI 16)*. 265-283 (2016).
 42. Paszke, A., Gross, S., Massa, F. et al. PyTorch: An imperative style, high-performance deep learning library. *Adv. Neural Inf. Process. Syst.* 32, 8024-8035 (2019).
 43. Yu, N., Capasso, F. Flat optics with designer metasurfaces. *Nature Mater.* 13, 139-150 (2014).
 44. Hu Y, Jin Y, Wu X, et al. A theory-guided deep neural network for time domain electromagnetic simulation and inversion using a differentiable programming platform. *IEEE Trans Antennas Propag.* 70, 767-772 (2021).
 45. Müller, J., Philipp, D. & Günther, M. Differentiable cosimulator for electrically reconfigurable structures. *IEEE Access* 13, 191343-191356 (2025).
 46. Ruder S. An overview of gradient descent optimization algorithms. arXiv preprint arXiv:1609.04747 (2016).
 47. Staib, M., Reddi, S., Kale, S. et al. Escaping Saddle Points with Adaptive Gradient Methods. *Proceedings of the 36th International Conference on Machine Learning (PMLR)*. 97, 5956-5965 (2019).
 48. Encinas-Lago, G., Albanese, A., Sciancalepore, V. et al. A Cost-Effective RISs Deployment to Abate the Coverage Problem in B5G Networks. *IEEE Trans. Wireless Commun.* 23, 15276-15290 (2024).

49. Nicolet, B., Jacobson, A., & Jakob, W. Large steps in inverse rendering of geometry. *ACM Trans. Graphics*. 40, 1-13 (2021).
50. Wang, N., Zhang, Y., Li, Z. Pixel2Mesh: Generating 3D Mesh Models from Single RGB Images. *Proceedings of the European Conference on Computer Vision (ECCV)*. 52-67 (2018).
51. Tojo, K., Shamir, A., Bickel, B. et al. Stealth Shaper: Reflectivity Optimization as Surface Stylization. *ACM SIGGRAPH 2023 Conference Proceedings (SIGGRAPH '23)*. 20, 1–10 (2023).
52. Jin, J. M. Theory and computation of electromagnetic fields. *John Wiley & Sons* (2015).
53. Chang, W., Yang, X., Belhe, Y. et al. Spatiotemporal Bilateral Gradient Filtering for Inverse Rendering. *ACM SIGGRAPH Asia 2024 Conference*. 70, 1-11 (2024).
54. Sorkine, O., & Alexa, M. As-rigid-as-possible surface modeling. *Symposium on Geometry processing*. 4, 109-116 (2007).
55. Tulgar, O., Durgut, K., & Ergin, A. A. A simple and efficient SBR implementation for RCS calculation using target scaling. *IEEE Trans. Antennas Propag.* 70, 3680-3687 (2021).
56. Aguilar, A, Tonder, J., Jakobus, U. et al. Overview of recent advances in the electromagnetic field solver FEKO. 2015 9th European Conference on Antennas and Propagation (EuCAP). 1-5 (2015).
57. Mao, N. et al. The Research of Standard Model's RCS Calculation and its Analytical Efficiency Based on FASTEM-STUDIO. 2022 2nd International Signal Processing, Communications and Engineering Management Conference (ISPCEM). 194-197 (2022).
58. Belogay E, Cabrelli C, Molter U, et al. Calculating the Hausdorff distance between curves. *Inf. Process. Lett.* 64(1) (1997).
59. Crane, K., Pinkall, U. & Schröder, P. Robust fairing via conformal curvature flow. *ACM Trans. Graph.* 32, 1–10 (2013).

Acknowledgements

This work was supported by the National Key Research and Development Program of China under Grant No. 2023YFB3811501, National Natural Science Foundation of China under Grant No. 62301146 and No. 62288101, Natural Science Foundation of Jiangsu Province of China under Grant No. BK20230816, China Postdoctoral Science Foundation under Grant No. 2023M730554 and No. BX20220065 and Postgraduate Research and Practice Innovation Program of Jiangsu Province under Grant No. KYCX25_0463.

Author contributions

Kaiqiao Yang and Che Liu contributed equally to this work. Tie Jun Cui and Wenming Yu supervised the research. All authors contributed to the data analysis and revising of the manuscript.

Competing interests

The authors declare no competing interests.

Supplementary information

Supplementary information is available for this paper.

Tables

Table I The detailed data table for the simulation tests shown in Fig. 2.

Model Name	Vertex Count	Average Monostatic RCS (dBsm) ¹			Running Time (s)	Efficiency (s per sp per kv) ²
		Initial Model	Optimized Model	Reduction		
Sphere	2100	6.00	-8.51	14.51	220.46	1.17
Plane	12674	0.24	-12.21	12.45	575.94	0.50
Sphere	2100	4.38	-10.52	14.90	169.03	2.68
Plane	12674	-16.84	-37.13	20.29	393.27	0.69

¹ These values are calculated based on 1° sampling interval.

² The label "sp" represents sampling points. The label "kv" refers to the thousand vertices.

Table II The detailed data table for the simulation tests shown in Fig. 3.

Model Name	Vertex Count	Average Monostatic RCS (dBsm) ¹			Running Time (s)	Efficiency (s per sp per kv) ²
		Initial Model	Optimized Model	Reduction		
Bunny	2959	3.73	-2.81	6.54	481.93	1.84
Calf	2930	1.64	-3.53	5.17	3568.74	2.41

¹ These values are calculated based on 2D RCS.

² The label "sp" represents sampling points. The label "kv" refers to the thousand vertices.

Table III The detailed radar cross section (RCS) reduction data for the simulation tests shown in Fig. 4.

Model Name	Average Monostatic RCS Reduction(dBsm) ¹				
	1GHz	2GHz	3GHz	4GHz	5GHz
Bunny	6.99	6.96	6.09	6.80	6.55
Calf	5.22	6.98	7.27	7.37	8.04

¹ The values of average monostatic RCS are calculated based on 360 azimuth angle sampling points.

Editor's Summary:

Kaiqiao Yang and colleagues propose a differentiable geometric optimization framework for manipulating electromagnetic fields. The approach enables automated shape design by coupling ray-based electromagnetic modelling with gradient-based optimization.

Peer review information:

Communications Engineering thanks the anonymous reviewers for their contribution to the peer review of this work. Primary Handling Editors: [Yu-Cheng Chen] and [Wenjie Wang]. A peer review file is available.