

GLOBAL OPTIMIZATION OF ANTENNAS USING CNN-BASED INVERSE SURROGATES AND EXPEDITED FINE-TUNING

Slawomir Koziel⁽¹⁾, Anna Pietrenko-Dabrowska⁽²⁾

⁽¹⁾Reykjavik University

⁽²⁾Gdansk University of Technology

Keywords: Antenna design, global optimization, inverse modeling, CNN, machine learning

Abstract

Formal optimization methods are widely used in modern antenna engineering. Rigorous numerical algorithms are essential for handling multiple decision variables (e.g., geometry parameters), design objectives, and constraints. For an increasing number of applications—including antenna array development, metasurface design, system reconfiguration over broad operating conditions, or the synthesis of compact structures—local parameter tuning is no longer sufficient. In such cases, global optimization becomes preferable, or even necessary. However, conventional global optimization techniques, most notably population-based bio-inspired algorithms, are well known for their high computational cost. At the same time, machine-learning-based approaches often struggle to construct data-driven surrogate models of sufficient predictive accuracy.

This paper introduces a novel global optimization framework that incorporates inverse data-driven metamodels implemented using deep convolutional neural networks (CNNs). In the proposed surrogate, the inputs are the frequency-domain characteristics of the antenna under consideration, whereas the outputs correspond to the associated vectors of decision variables. To enable the construction of models capable of representing antennas implemented on substrates with different permittivity and thicknesses, a dedicated embedding mechanism is developed. Furthermore, residual connections within the CNN architecture allow the model to capture interactions between the real and imaginary components of antenna responses and the geometry parameters. The training data are generated within a dimensionally reduced space that captures the most influential directions in the design domain, which enables accurate model training using a relatively small number of samples. These directions are identified through a fast global sensitivity analysis conducted using a limited set of random samples.

The optimization process consists of two main stages. The first stage is global optimization, in which the inverse surrogate generates an initial design directly from the performance specifications (e.g., the desired operating bandwidth), mapped into a target frequency response. This response is constructed through an intermediate

step in which the intended operating conditions are transformed into a rational function characterized by appropriate values of the center frequency and quality factors. In the second stage, the initial design is refined using an accelerated gradient-based optimizer. When approaching convergence, the optimizer employs a rank-one Broyden update formula to reduce the number of full-wave simulations required.

The proposed methodology has been validated using two planar antennas optimized with respect to reflection performance and gain. For each antenna, multiple design scenarios were considered, corresponding to different center frequencies, bandwidths, and substrate materials. The results demonstrate the robustness, reusability, and versatility of the inverse surrogate, as well as its effectiveness in optimizing antenna designs across a wide range of frequencies, geometries, and substrate configurations. Extensive benchmarking against a broad spectrum of optimization approaches—including bio-inspired algorithms, machine-learning methods, and surrogate-assisted population-based techniques—confirms the advantages and practical engineering value of the proposed framework. One of its most important benefits is computational efficiency: the average optimization cost corresponds to fewer than thirty full-wave simulations, compared to at least 500 simulations required by the best-performing benchmark algorithms. An additional advantage is that, once the inverse surrogate has been constructed, the optimization process becomes deterministic and produces identical results across multiple runs, unlike the benchmark methods.