

In-Situ Inverse Design of a Plasma Metamaterial Beam Steering Device

Katherine P. Bronstein, Noah A. Harris, Aleczaender J. Harder, Jennay L. Edmondson, and Jesse A. Rodríguez*



Cite This: <https://doi.org/10.1021/acsphotonics.6c00188>



Read Online

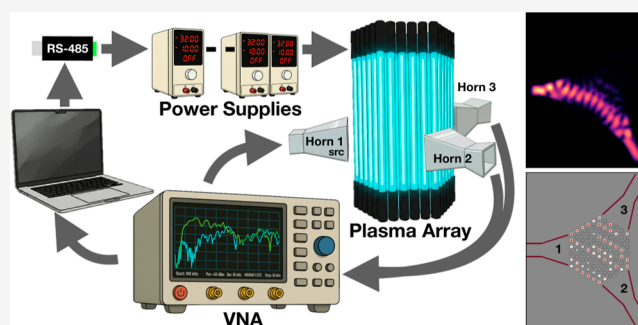
ACCESS |

Metrics & More

Article Recommendations

ABSTRACT: Inverse design is a commonly used methodology for creating devices that manipulate electromagnetic (EM) waves by algorithmically modifying the device parameters to achieve a desired functionality. Utilizing plasma, a dynamically tunable medium, allows the optimization of the design process to be conducted directly on the experimental hardware (in situ). A key advantage of this method is the creation of devices that are inherently switchable and dynamically reconfigurable. Bayesian optimization is used to tune the plasma density of 91 independent discharges that make up a plasma metamaterial (PMM) device to steer incoming EM waves to the desired exit waveguides. Measurements were conducted in an automated loop, where a vector network analyzer records the PMM transmission characteristics for each device setting. By relying only on measured scattering parameters, this gradient-free approach is robust to experimental drift and noise and does not require complex full-wave models. Significant performance improvements over traditional simulation-based (in-silico) inverse design are demonstrated, with in situ Bayesian optimization achieving up to 10,000× higher isolation between ports than the best *in-silico* design at the same target frequency. This work also presents guidelines for applying Bayesian optimization to noisy, highly dimensional physical systems.

KEYWORDS: plasma metamaterial, inverse design, beam steering, reconfigurable photonics, optical computing, Bayesian optimization, microwave metamaterials



INTRODUCTION

Inverse design has become a powerful way to discover electromagnetic devices that would be difficult to design by hand. In this approach, the device permittivity distribution is encoded by a set of tunable parameters, and those parameters are optimized so that the measured or simulated fields match the desired behavior. Plasma metamaterials (PMMs) are a useful platform for inverse design because each plasma element can be tuned electronically over a wide range of effective permittivities, making the same hardware reusable for many different device functions that can be switched dynamically.

In electromagnetics, inverse design is widely used to create devices that are effectively impossible to design by hand.^{1–7} These systems are constrained by Maxwell's equations and require numerical solvers to evaluate and optimize each design; heretofore referred to as *in-silico* inverse design. Many inverse-designed electromagnetic devices have been demonstrated, including metalenses,^{8,9} filters and wavelength demultiplexers,¹⁰ photonic crystals,^{11–15} conformal volumetric metamaterials,¹⁶ and even optical computing devices like boolean logic gates¹⁷ and matrix multiplication devices.¹⁸ These operations underlie tasks such as signal processing, commu-

nication, and machine-learning inference; therefore, faster implementations translate directly into higher-throughput specialized computing hardware.

Unlike prefabricated metamaterial structures,^{18–21} PMMs offer real-time control of wave propagation by tuning plasma density.^{22–24} The epsilon near zero (ENZ) regime is crucial to operating these PMM-based devices. By tuning plasma density near the critical frequency, the effective permittivity can approach zero, enabling unique field confinement, phase control, and impedance matching effects.^{25,26} In PMMs, this is achieved through the electrical control of the plasma frequency, making the ENZ behaviors dynamically accessible and providing an advantage over conventional dielectric metamaterials. This allows the PMMs to serve many different purposes without any refabrication. Earlier plasma photonic

Received: January 23, 2026

Revised: April 16, 2026

Accepted: April 27, 2026

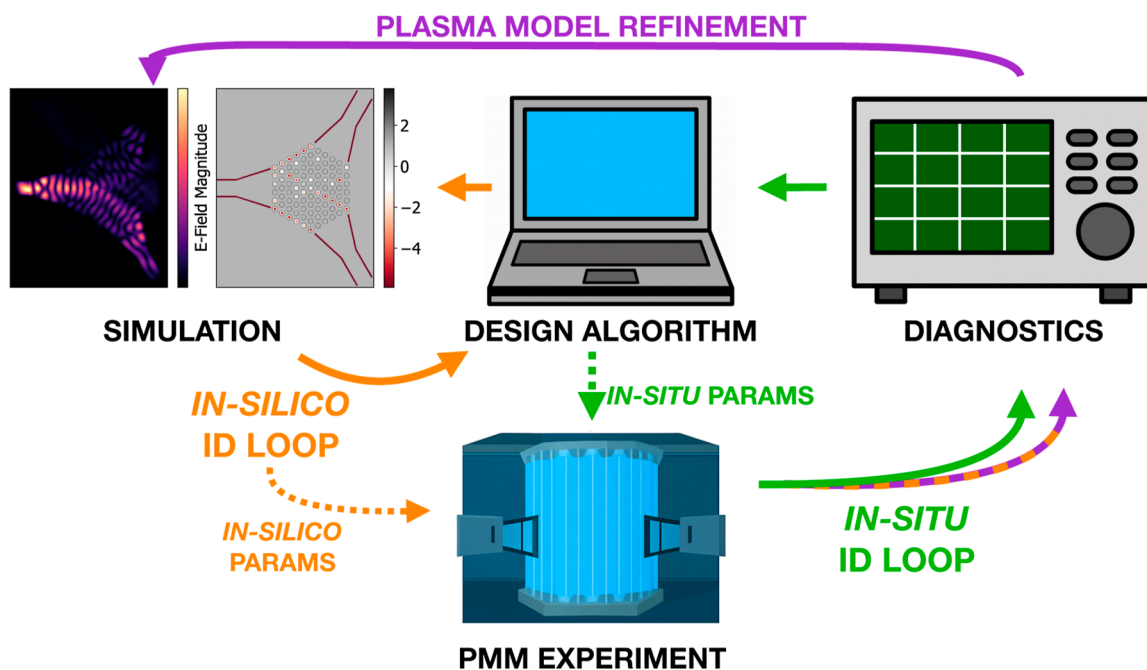


Figure 1. Illustration of the inverse-design workflow showing the complementary roles of in-silico and in situ optimization. This PMM experiment graphic was created with assistance from Gemini Pro based on an author-generated Keynote draft.

crystal experiments have demonstrated bandgap devices, waveguides, and lattice-resonance effects in related discharge geometries, establishing plasma as a viable platform for reconfigurable electromagnetic structures.^{27–35}

The tunability of the plasma elements that constitute a PMM arises from the Drude dielectric permittivity of low temperature plasmas. The permittivity (neglecting the ion’s response, an appropriate assumption for low temperature plasma) is of the following form

$$\epsilon = 1 - \frac{\omega_p^2}{\omega^2 + i\omega\gamma} \quad (1)$$

where γ is the collision/damping rate, $\omega_p^2 = \frac{n_e e^2}{\epsilon_0 m_e}$ is the plasma frequency squared, n_e is the electron density, e is the electron charge, m_e is the electron mass, and ϵ_0 is the free-space permittivity. We can see here that in the collisionless limit, the plasma permittivity can be tuned continuously in the interval $(-\infty, 1)$ by varying the plasma density n_e for a given operating frequency ω . This is done in practice by supplying more power to the discharge. By choosing an operating frequency that is close to the plasma frequency of the elements, we can easily make drastic changes to their electromagnetic response, i.e., to that of a metallic element or a near-zero-index medium. This dispersion relation not only makes the configuration space infinite as opposed to the binarized parametrization schemes most often used in the inverse design of electromagnetic devices, but more importantly, this allows a single array of plasma elements to perform several functions. Moreover, since the plasma elements can be quickly deactivated or tuned, you also can switch the device on and off or to a different function very quickly in ~ 10 ms or less.

A major opportunity created by PMMs is the ability to run inverse design directly on the experimental device and link it to a traditional in-silico design in a closed discovery loop. In this work, we compare in-silico inverse design that utilizes a

realistic numerical plasma model²² with *in situ* inverse design that optimizes the plasma frequency of the plasma elements using only measured scattering parameters. Related *in situ* training ideas have recently been demonstrated in photonic neural networks, where the photonic hardware itself is optimized using measured gradients.³⁶ The *in-silico* loop provides physics-aware gradient-based updates but depends on assumptions about density profiles, collision rates, and boundary conditions as well as the dimensionality of the simulation, whereas the *in situ* loop operates on the real hardware and automatically includes experimental imperfections such as collisional damping, electrode geometry, misalignment, and three-dimensional effects. Comparing the optimal configurations from these two loops reveals where the plasma model and experimental plasma frequency mapping fail to match reality, turning the device into a platform for both performance optimization and plasma-model validation.

This paired approach turns the PMM into both a reconfigurable device and a probe of its own underlying physics. When the *in situ* and *in-silico* optimizations disagree, differences in performance and the resulting plasma-density distribution throughout the device elements highlight specific gaps in our modeling of the plasma elements. Over many optimization runs, the accumulated experimental data can be used to refine the plasma model and to train surrogate models that accelerate future design cycles. Figure 1 summarizes this workflow: a design algorithm exchanges parameters with both the numerical solver and the physical PMM, while diagnostics close the loop by feeding measurements back into the model.

In this study, we present the first electromagnetic device that has a domain created entirely via in situ inverse design. We demonstrate how the in situ inverse design process, when coupled with the PMM device paradigm, leads to vastly better performance than the in-silico inverse design while also enabling devices that are inherently multifunctional and tunable.




Figure 2. Experimental setup: (left) schematic of the PMM control system, and (right) photograph of the physical implementation of the 91-element plasma array. The left-panel schematic was created with assistance from Gemini Pro based on an author-generated Keynote draft.

We begin by describing the PMM hardware, voltage-to-plasma frequency mapping, and inverse-design objectives and methodology. Next, we compare *in-silico* and *in situ* beam-steering results for standard, narrowband, and broadband objectives. Finally, we discuss how these results inform plasma-model refinement and future PMM-based electromagnetic wave control, including optical computing devices.

METHODS

The device utilized in this study is a 2D plasma metamaterial composed of 91 individually addressable discharge tubes arranged in a triangular lattice in the style of a photonic crystal. The elements form a hexagon with six bulbs along each side, so that the center-to-center spacing (20 mm) is fixed across the array, as seen in Figure 2. The spacing is small compared with the 4–10 GHz wavelengths of interest, so the array can behave like a single contiguous metamaterial. The array is mounted between a source microwave horn antenna (port 1) and two receiver horns (ports 2 and 3), and the entire assembly is placed inside a compact anechoic chamber lined with microwave absorber to suppress reflections during *in situ* optimization and S-parameter measurements, as seen on the right side of Figure 2.

Microwave transmission is measured with a Rohde & Schwarz ZNB40 vector network analyzer (VNA) configured to measure S_{21} and S_{31} transmission spectra between the input (port 1) and each receiver horn antenna (ports 2 and 3). The VNA is calibrated using an automatic calibration kit (R&S ZN-Z54) along with the SMA test cables used to connect the VNA to the horn antennas. During optimization of the PMM devices, the VNA is left in a fixed measurement state with a frequency sweep covering the band of interest (0–20 GHz) with 10,000 points and an averaging factor of 10 to reduce noise. For each array configuration selected by the optimizer, the control script triggers a fresh sweep, waits for completion, and then queries the complex S_{21} and S_{31} traces along with the corresponding frequency vector. In postprocessing, we convert these to magnitudes in dB and use them to compute scalar objective values.

Each plasma element is a custom UV germicidal discharge tube made of high-purity quartz (relative permittivity $\epsilon = 3.8$) that is filled with Argon and a small amount of mercury to provide UV radiation. We used UV tubes instead of commercial lighting bulbs to avoid fluorescent coatings that are opaque at microwave frequencies. The quartz envelope has an outer diameter of 15 mm with a thickness of 1 mm, yielding an inner diameter of 13 mm. In operation, the neutral gas temperature is approximately 315–330 K (measured via an infrared thermometer), and the discharges are designed for AC drive at 20–50 kHz. The discharge waveform is controlled via DC-to-AC ballast circuits with a maximum driving DC voltage of about 20 V. At

this maximum driving voltage, each ballast-discharge circuit draws about 30 W. The discharges have a rise-time of about 14 ms and a fall-time of about 45 ms. These times are defined according to the plasma's dielectric response; see App. A of ref 37.

To enable software-controlled operation, each discharge-ballast circuit is driven by a dedicated programmable Longwei LW-3010 DC power supply. The power supplies are connected in groups of ten to RS-485 communication buses and controlled using MODBUS serial commands from a custom python library. To compare the *in-silico* optimal parameters to those from the *in situ* optimization algorithm, we introduce a fixed mapping between the simulated element indices and the hardware addresses of the bulbs and power supplies. In the numerical simulations, the design variables form a vector ρ whose entries correspond to elements ordered in a triangular lattice. In the hardware, each discharge is linked to that mapping.

To connect the *in-silico* plasma frequency parameters to experimental control settings, we construct a quasiexperimental mapping from power-supply voltage to plasma frequency. To measure the plasma density, we use a 0-D kinetic solver (BOLSIG+^{38,39}) to relate the measured reduced electric field across the discharge E/n to the electron energy distribution and consequently the electron mobility and drift velocity, which can be used to calculate the electron density. For each operating point, we measure the RMS voltage and current across a single discharge. We then estimate the neutral gas density from the fill pressure (200–300 Pa, proprietary to the discharge producer and thus not precisely known) and a range of assumed gas temperatures (315–330 K) via the ideal gas law. Using the discharge length and an estimate of the cathode voltage drop,⁴⁰ we infer the effective electric field in the positive column. BOLSIG+ then provides the electron mobility μ_e , which we multiply by the measured electric field to obtain the average electron drift velocity, $u_d = \mu_e E$. Combining this drift velocity with the measured RMS current and the tube cross-sectional area gives an estimate of the electron density $n_e = I_{\text{meas,RMS}} / (eu_d A_{\text{tube,inner}})$, and thus, the plasma frequency. We repeat this procedure over a range of voltage-limited points, sweep over possible fill pressures, temperatures, and cathode voltage drops, and then fit a curve that maps our voltage to plasma frequency, following the approach detailed in Appendix A of ref 37.

The resultant fit for plasma frequency f_p in GHz follows a logarithmic function of the form

$$f_{p,\text{GHz}} = S[(10.5 + 2.5k)\log_5(V - 4.8) - 4.5] \quad (2)$$

where k tunes between low and high-density envelopes and the global scale factor S shifts the entire curve to account for systematic over- and underestimates of electron density. The mapping is plotted in Figure 3.

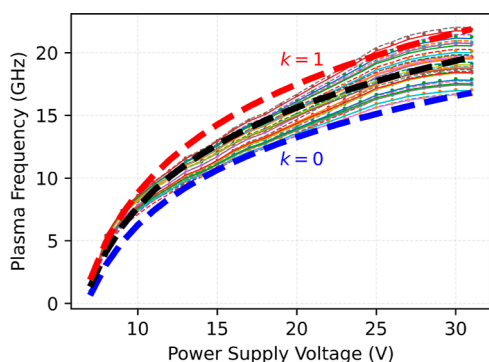


Figure 3. Quasi-experimental mapping for plasma frequency tuning, showing three fitted voltage-frequency mappings (thick dashed curves) that bound the low-, mid-, and high-density cases encoded by k in eq 2. Each of the thin colored curves plotted underneath the fitted function show the BOLSIG + -based estimate of plasma frequency for a specific combination of gas pressure and temperature (e.g., 220 Pa and 300 K) resulting from a sweep across possible values.

For the *in-silico* optimization, the device is modeled with an autograd-compliant two-dimensional finite difference frequency domain (FDFD) Maxwell solver, Ceviche,⁴¹ where the domain includes the input and output horns, perfectly matched layers at the boundaries, and a hexagonal array of discharges using the plasma element model detailed in.²² The source mode E -field is polarized along the length of the plasma discharges (E_z with z out of the side of the page in Figure 4). The design vector ρ encodes the average plasma frequency and thus permittivity of each of the elements, and the solver returns the fields for a given source frequency. The inverse design algorithm utilizes objectives based on overlap integrals between the simulated field and the desired waveguide mode, the same as those used in prior work.^{17,22,23} These differentiable objectives are used to

optimize ρ using the Adam optimizer⁴² with the learning rate α ranging from 0.01 to 0.05 and the default Adam hyperparameters $\beta_1 = 0.9$ and $\beta_2 = 0.999$. It is important to note that since the device is simulated in the frequency domain, the numerical results represent the steady state device behavior for a single frequency, and thus, the full-spectrum response of the device is not factored into the *in-silico* optimization routine.

To mitigate any reproducibility issues associated with transient thermal effects or elemental hysteresis, each optimization follows a standard sequence. Before optimization, the array is allowed to “warm up” by cycling the entire array on and off at the desired experimental duty cycle for approximately 10 min. Once the optimization commences, each objective evaluation follows the same procedure: First, the supplies are pulsed to a higher ignition voltage to strike the discharges. Once all bulbs are lit, the voltages are adjusted to their operating values specified by the parameter vector ρ and held for a short time (on the order of a second) to allow any transient effects associated with element tuning to subside. The array remains on while the vector network analyzer (VNA) acquires S -parameters, after which the discharges are switched off and allowed to cool. To avoid overheating the tubes and ballasts, we enforce a 25% duty cycle (roughly 15 s per minute) during extended runs.

To optimize the device *in situ*, we treat the PMM as a black-box function that maps a set of power-supply settings to a scalar performance metric. The experimental objective functions are defined as discrete sums over differences between the measured scattering parameters at ports 2 and 3. For the standard beam steering objective where port 2 is the desired exit port, we first compute, at each frequency f , the dB-scale difference: $d(f) = S_{21,\text{dB}}(f) - S_{31,\text{dB}}(f)$, so that a large positive d indicates more transmission into the correct exit port. The objective function then consists of summing $d(f)$ over a specified frequency band $B = [f_{\text{op}} - \Delta f/2, f_{\text{op}} + \Delta f/2]$ where f_{op} is the desired operating frequency. For the standard objective, Δf was 0.5 GHz.

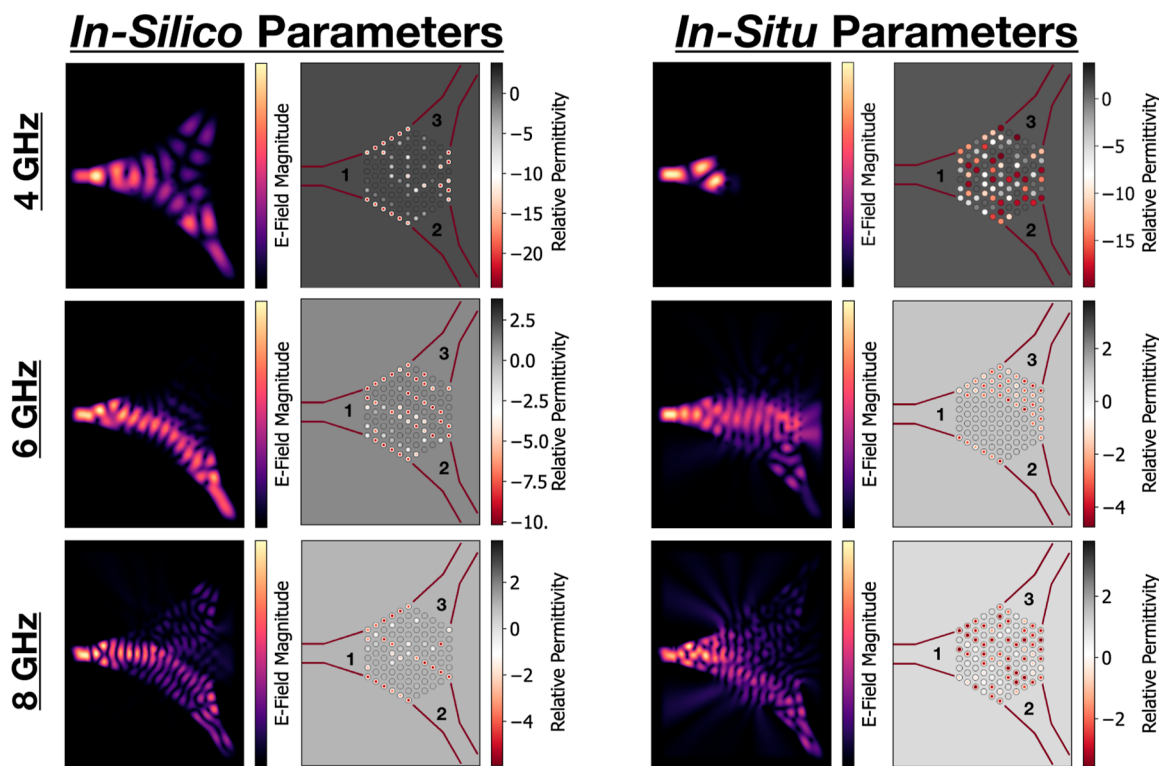


Figure 4. PMM FDFD simulation results at 4, 6, and 8 GHz: steering comparing *in-silico* and *in situ* parameters, showing stronger steering with the *in-silico* parameters. The left panels for each parameter set plot E-field intensity $|E_z|^2$ in the domain and the right panels show the relative permittivity in the domain. The discharges with lowest permittivity correspond to those with the highest plasma density.

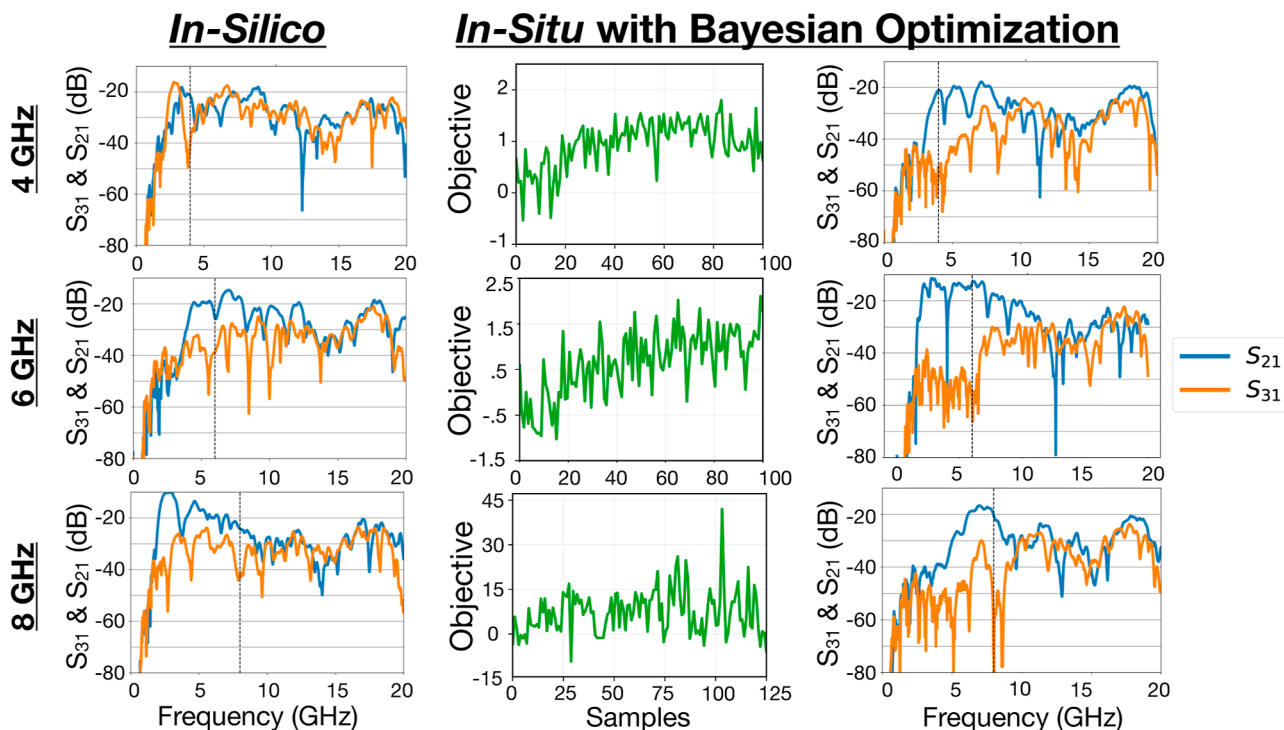


Figure 5. PMM experimental transmission spectra for the standard beam steering objective at 4, 6, and 8 GHz for the *in-silico* parameters (left) and the *in situ* parameters (right) along with the objective evolution over time (center). The operating frequency is denoted by a vertical dashed black line. The objective values are in arbitrary units. For each target frequency, *in situ* optimization improves isolation over the *in-silico* parameters by a factor of up to 10^3 at 6 GHz and 10^4 at 8 GHz on a power scale.

$$\mathcal{L}_{\text{standard}} = \sum_{f \in B} d(f) \quad (3)$$

For narrowband beam steering, we desire a narrow band of isolation in favor of the correct output port and minimal difference in transmission between the ports at other frequencies. In this case, we define an in-band window (still with $\Delta f = 0.5$ GHz) around a target frequency f_0 and treat everything outside that window as out-of-band (Right and Left). The narrowband objective is computed as

$$\mathcal{L}_{\text{narrow}} = w_{\text{in}} \sum_{f \in B} d(f) - w_{\text{out}} \sum_{f \notin B} d(f)^2$$

where $w_{\text{in}} = 2.0$ and $w_{\text{out}} = 0.5$ are objective hyperparameters used to differentially penalize or promote spectral differences within, below, and above the operating band. These parameters are adjusted by hand, guided by experimental results. The squared penalties ensure that any difference in the spectra out-of-band are discouraged.

The extra-broadband objective is designed to push performance over a wide contiguous frequency range. It rewards isolation across a broad frequency band while penalizing rapid fluctuations in the transmission difference. The set B is expanded to include the entire operating range of the microwave horns, 0–20 GHz. The objective is

$$\begin{aligned} \mathcal{L}_{\text{broadband}} &= w_{\text{sep}} \sum_{f \in B} d(f) \\ &\quad - w_{\text{flat}} \sum_i [d(f_{i+1}) - d(f_i)]^2, \end{aligned} \quad (4)$$

where $w_{\text{sep}} = 1.0$ and $w_{\text{flat}} = 0.05$. The second term penalizes roughness by summing the squared differences between adjacent frequency samples.

We implement Bayesian optimization⁴³ with a Gaussian-process surrogate (implemented via `gp_minimize` from `scikit-optimize`) to explore each “black-box” objective. The optimizer proposes a new set of parameters, ρ by maximizing an acquisition function (expected improvement) over the bounded design space. For

each proposed ρ , we apply the simulation to hardware mapping, convert the entries to plasma frequencies and voltage set points, execute the ignition and warm-up sequence, measure S_{21} and S_{31} , compute the chosen objective (standard, narrowband, broadband), and return it to the optimizer. The initial set of parameter samples used to condition the surrogate are either randomly sampled or a “warm-start” is conducted using samples from previous runs. The initial parameter vector before the first set of samples corresponds to a configuration with the incorrect exit port effectively “blocked” by high plasma density elements so that the most power was directed toward the desired exit at the start of the procedure. Each evaluation is logged along with diagnostic plots of the S-parameters and the objective value. To avoid plateaus, the parameters may be manually perturbed to encourage the exploration of other regions of the configuration space. We repeat this loop until the objective converges or a predefined evaluation number is reached and then the best performing iteration is selected, yielding *in situ* optimized configurations for the chosen objective.

RESULTS

We first applied *in situ* optimization to the standard beam steering objective at 4, 6, and 8 GHz. Separately, the device was optimized *in-silico* using Ceviche at the same three frequencies following the procedure laid out in refs 22,23 We present simulated fields for the optimal parameters from both inverse design approaches (*in-silico* vs *in situ*) in Figure 4. As expected, the *in-silico* parameters produce strong performance in the simulated device, whereas the *in situ* parameters result in a comparatively poor performance. This is in stark contrast to the experimental results on the PMM device, where the trend reverses.

In Figure 5, we see that the *in situ* optimized device produces higher measured isolation between ports 2 and 3 than does the *in-silico* design at the target frequencies. At 6 GHz, the *in-silico*

design yields about 15 dB of isolation, while the *in situ* design achieves about 45 dB isolation, corresponding to a $\sim 1,000\times$ improvement. At 8 GHz, the *in-silico* design likewise gives about 15 dB isolation, and the *in situ* about 60 dB precisely at 8 GHz, corresponding to about a $\sim 10,000\times$ improvement. Thus, for the same task, the *in-silico* parameters moderately outperform the numerical model, but the *in situ* parameters perform far better on the actual device.

Next, we applied *in situ* optimization with the narrowband objective to measure how much isolation the device can provide at a single frequency. For this objective, the initial parameters were set such that all of the discharges were inactive. At individual target frequencies between 4 and 8 GHz, including the 4 and 7 GHz cases presented in Figure 6, the

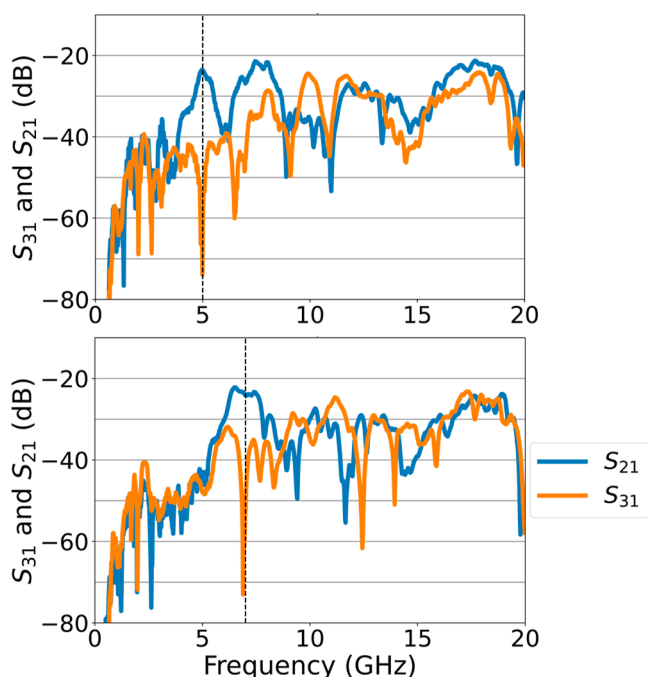


Figure 6. *In-situ* narrow band optimization results targeting 5 (top) and 7 (bottom) GHz. The optimizer finds patterns that yield about 50 dB isolation at the target frequency while maintaining high transmission into the desired port and minimal isolation at other frequencies.

device achieved approximately 50 dB of isolation between the desired and undesired ports, at least as high as the isolation obtained with the standard objective but over a much smaller bandwidth. Where the standard objective results in isolation from about 2 to 10 GHz of varying strength, the narrow band objective limits the strong isolation to a frequency band of approximately 0.3 GHz in width with the remainder of the spectrum limited below about 10 dB isolation in either direction.

These experiments were conducted over many weeks, and some discharges and ballast circuits were replaced as they degraded or failed. As a check on experimental drift and to determine if these changes had a strong effect, the previously optimized narrowband parameters for 4 and 8 GHz were reapplied after the experimental campaign concluded. The array still produced clear but degraded beam steering with reduced, broadened, or frequency-shifted isolation peaks, indicating that component replacement and long-term drift can partially erode the stored optimum and would likely

require reoptimization to fully recover the original narrowband performance. To test this, a follow-up optimization starting from the previously optimal parameters targeting 4 GHz shifted the isolation peak back to the target frequency within three samples (~ 30 dB separation) and recovered the full ~ 50 dB narrowband isolation after just nine samples (i.e., in less than 10 min), showing that deep narrowband performance can be restored quickly. These results indicate that the system remains sufficiently close to the desired operating regime for rapid warm-start reoptimization, suggesting that periodic recalibration would be sufficient to maintain performance in practice. Moreover, this rapid re-establishment of optimal performance even after several parts were replaced shows how this device paradigm is robust to major hardware damage, making it suitable for extreme environments. In future implementations, the system could query the power draw of each ballast at a fixed reference voltage during optimization and use changes in that baseline as an indicator of ballast or tube degradation, helping flag drift in the apparatus before it impacts device performance.

Finally, we tested the *in situ* inverse design procedure on the broadband objective. In this case, the objective rewards a positive dB-scale difference between ports 2 and 3 across as many frequencies as possible while smoothing the rapid variations in transmission. As seen in Figure 7, the optimizer




Figure 7. *In-situ* results for the broadband objective. *In-situ* optimization produces a configuration that maintains useful isolation between ports 2 and 3 across a wide frequency range, with clear separation from approximately 4 to 12 GHz.

finds a configuration that maintains nonzero isolation over a wide band, with clear separation between ports 2 and 3 from roughly 4 to 12 GHz ranging from 55 to 15 dB isolation, with fewer drops in transmission to the correct port across the isolation band, with an additional region of isolation at the higher frequency band from 13 to 15 GHz.

DISCUSSION

The results show a clear mismatch between how the plasma metamaterial behaves in simulation and how it behaves in the experiment, highlighting why *in situ* optimization is necessary. The simulation results in Figure 4 utilize a plasma element model that includes several nonideal factors, including nonuniform density profiles, collisionality/loss, robustness to random perturbation/error, and quartz envelopes. With these factors considered and the voltage–frequency mapping used in the model with the best case fit tuning factors k and S , the numerical model fails to reproduce the stronger performance observed experimentally with the *in situ* design. This indicates

that, although the gradient-based *in-silico* optimization is doing its job and finding a performant design given our plasma model, the model does not fully capture the real plasma and hardware behavior, including the actual density profile, loss, glass response, horn coupling, and, perhaps most importantly, three-dimensional effects. Although the discharges are much longer than the horn antenna aperture dimensions, the presence of the discharge electrodes or even axial plasma density variations may be playing a role in the shortcomings of the numerical model. At the same time, the *in situ* pattern still often provides reasonable beam steering performance in the simulations, which suggests that the numerical model captures at least the first-order physical behavior of the device.

Mutual coupling between plasma columns is implicitly captured in the *in situ* optimization because the objective is evaluated from measured device-level scattering parameters of the full 91-element array, so any interelement electromagnetic interaction or weak nonlinear effect that affects the transmitted fields is automatically included in the optimization loop. By contrast, the *in-silico* model can miss or mis-estimate these interactions because it relies on a reduced representation of the experiment, where the plasma elements are modeled in the frequency domain as identically structured (save for modified average plasma densities) static DC discharges. In that sense, the discrepancy between the *in situ* and *in-silico* optima is not evidence that coupling is absent but rather evidence that the experimental device includes higher-order interactions that are only fully captured when the optimization is performed directly on the hardware.

The experimental results in Figure 5 show how *in situ* optimization can sidestep errors in numerical models. This objective focused on target frequencies in the range of 4–8 GHz. This range was chosen because the nominal operating plasma frequency (driving at the suggested ballast voltage) of the discharges is estimated to be about 10 GHz, so the device is expected to have the most dynamic range in a frequency band slightly below that value following the Drude permittivity. Within this band, the *in situ*-optimized patterns provide much higher measured isolation than the transferred *in-silico* design: at 6 GHz, the isolation improves by roughly 3 orders of magnitude, and at 8 GHz, it improves by around 4 orders of magnitude, bringing transmission in port 3 all the way to the noise floor of the VNA. This is because all nonideal factors are automatically taken into account when optimizing *in situ*.

This does, however, present a unique potential method for plasma model refinement. By running both the *in-silico* and *in situ* inverse design in parallel and comparing results, we can comprehensively evaluate the accuracy of our plasma model. As opposed to a slow experiment-simulation feedback loop for model refinement that may neglect edge cases and device performance, the optimization algorithms will probe the full dynamic range of the real and simulated devices, ensuring that if the performance of the optimal configurations from both methods converge, the plasma model must represent the underlying physics extremely well. The large amount of experimental data can also be used to produce low-cost machine learning surrogate models that may speed up both the *in-silico* and *in situ* inverse design loops.

The narrow band *in situ* results in Figure 6 show what the device can do when it is only asked to perform at a single frequency. This resulted in a similarly high isolation to the standard objective but in very precise frequency bands. This highly distinct device behavior exhibits the dynamic range of

the PMM device. In a single device with ~ 10 ms switching times, operating frequency bands can be shifted, expanded, and contracted.

The broadband results in Figure 7 test how wide of a frequency band a single device setting can maintain for useful beam steering. The optimized configuration provides clear and consistent separation between ports but with a lower peak isolation than the narrow band designs. The key difference between this result and those from the standard objective is the influence of the smoothing term. While the standard objective achieves isolation in the correct direction over a similar frequency band in some cases, the strength of isolation is highly variable, sometimes even reversible, such as what is seen in the 6 GHz case. The broadband result also highlights the limited range due to the voltage-plasma frequency tuning range of the current hardware. Since this objective sought to increase isolation across the entire measured spectrum, the poor performance at higher frequencies is indicative of a shortcoming of the device. This was expected, as the maximum plasma frequency of the elements is estimated to be around 15 GHz. Following Drude permittivity, the plasma elements will become increasingly more transparent above that frequency, reducing their capacity to affect EM wave propagation. This suggests that future devices will need higher-voltage drives and more resilient discharges to access broader frequency spans. If operated in short pulses, much higher plasma frequencies can be reached without damaging the hardware.⁴⁴

These experiments also expose several limitations of the current *in situ* optimization loop. The control space is very large (91 plasma discharges), which is not ideal for Bayesian optimization. As the array size grows, this dimensionality challenge is expected to become more severe. Gaussian-process-based Bayesian optimization is attractive in the present system because it is sample-efficient and tolerant of experimental noise, but its performance generally degrades as the number of independently controlled plasma elements increases and the search space becomes harder to model globally. For larger PMM arrays, more structured parametrizations,^{45,46} dimensionality-reduction strategies,⁴⁷ warm-starting from previously learned configurations, or alternative black-box optimizers such as CMA-ES⁴⁸ or reinforcement-learning-based⁴⁹ methods may be required. The 91-element system studied here therefore represents both a practical demonstration of *in situ* optimization at relatively high dimensions and a useful benchmark for developing more scalable optimization workflows.

Because the acquisition function is designed to balance exploration and exploitation, the optimizer will often jump to quite different configurations to “explore” uncertain regions and then move back toward settings that previously looked good. As a result, the sequence of objective values is nonmonotonic and can be very noisy, even when the algorithm is slowly converging over time. In addition, long ignition and warm-up times, duty-cycle limits, and occasional misfires in bulb ignition restrict how many evaluations we can afford to run and introduce slow drifts in the underlying objective, which further challenge the optimizer. Even with these limitations, we nevertheless demonstrate a functioning *in situ* optimization loop, presenting a great opportunity for improved performance on more complex objectives in future work by utilizing better optimization algorithms and workflows.

CONCLUSION & FUTURE WORK

The experiments show that *in situ* optimization enables strong beam steering on the plasma metamaterial device when the plasma settings are tuned directly from measurements rather than from inverse design algorithms based on simulations. In numerical simulations, *in-silico* inverse design parameters appear to give the best steering and isolation. On the real device, those same configurations do not perform as well as those found by Bayesian optimization. This mismatch highlights the impact of modeling uncertainty and further justifies the *in situ* inverse design approach. We have shown that Bayesian optimization can tune a plasma metamaterial beam-steering device directly from experimental measurements. Crucially, this single device is shown to be rapidly reconfigurable to arbitrary operating frequencies with a very highly variable operating bandwidth, a functionality that is not possible with traditional electromagnetic devices. This charts a path toward a fully programmable and generalizable electromagnetic device based on the PMM paradigm.

Moving to fully *in situ* optimization also removes some of the limitations on device geometry and physics. 3D geometry is generally far too computationally costly to use for *in-silico* inverse design, but it incurs no extra cost in *in situ* inverse design. Thus, we could use device structures like that of refs 32,50. Even for the 2D configuration in this study, each iteration in the optimization procedure can take 5–10 min with a 64-core workstation, while the *in situ* iterations can be limited to mere seconds in practice regardless of the geometry. To combat the loss inherent to gaseous plasma, low-loss dielectric scaffolds could be incorporated into the inverse design workflow via a bilevel optimization scheme.⁵¹ We have also shown in prior work that our plasma sources have a pronounced gyrotropic response when magnetized,^{33,35} but materials with anisotropic permittivity tensors like magnetized plasma cannot be modeled with our simulation tool. By using a large Helmholtz configuration, entire PMM devices can be magnetized to take advantage of the very rich physics inherent to magnetized plasmas, enhancing the dynamic range of our device even further, introducing nonreciprocity and increasing the potential for other exotic device behaviors like dynamic cloaking.^{37,52} More sophisticated black box optimization techniques that handle many parameter systems more effectively such as deep learning-based reinforcement learning⁵³ or CMA-ES⁵⁴ may be used to fully leverage these more complex device phenomena.

ASSOCIATED CONTENT

Data Availability Statement

Data available upon reasonable request.

AUTHOR INFORMATION

Corresponding Author

Jesse A. Rodríguez – Department of Mechanical, Industrial and Manufacturing Engineering, Oregon State University, Corvallis, Oregon 97331, United States;
Email: jesse.rodriguez@oregonstate.edu

Authors

Katherine P. Bronstein – Department of Mechanical, Industrial and Manufacturing Engineering, Oregon State University, Corvallis, Oregon 97331, United States;
orcid.org/0009-0004-7608-8608

Noah A. Harris – Department of Mechanical, Industrial and Manufacturing Engineering, Oregon State University, Corvallis, Oregon 97331, United States

Aleczer J. Harder – Department of Mechanical, Industrial and Manufacturing Engineering, Oregon State University, Corvallis, Oregon 97331, United States

Jennay L. Edmondson – Department of Mechanical, Industrial and Manufacturing Engineering, Oregon State University, Corvallis, Oregon 97331, United States

Complete contact information is available at:

<https://pubs.acs.org/10.1021/acsp Photonics.6c00188>

Notes

The authors used ChatGPT (OpenAI) for limited coding assistance and Gemini Pro (Google) to assist in the creation of Figures 1 and 2 based on author-generated Keynote drafts. All scientific content and figure interpretations were reviewed and verified by the authors.

The authors declare no competing financial interest.

ACKNOWLEDGMENTS

The authors acknowledge the School of Mechanical, Industrial, and Manufacturing Engineering at Oregon State University as well as the Nesbitt Faculty Scholar in Energy Engineering Fund for providing funding to support this work.

REFERENCES

- (1) Su, L.; Vercruyse, D.; Skarda, J.; Sapra, N. V.; Petykiewicz, J. A.; Vučković, J. Nanophotonic inverse design with SPINS: Software architecture and practical considerations. *Applied Physics Reviews* **2020**, *7*, 011407.
- (2) Miller, O. D. *Photonic Design: From Fundamental Solar Cell Physics to Computational Inverse Design*, 2013; .
- (3) Liu, D.; Gabrielli, L. H.; Lipson, M.; Johnson, S. G. Transformation inverse design. *Opt. Express* **2013**, *21*, 14223–14243.
- (4) Hughes, T. W.; Minkov, M.; Williamson, I. A. D.; Fan, S. Adjoint Method and Inverse Design for Nonlinear Nanophotonic Devices. *ACS Photonics* **2018**, *5*, 4781–4787.
- (5) Molesky, S.; Lin, Z.; Piggott, A. G.; Jin, W.; Vučković, J.; Rodríguez, A. W. Inverse design in nanophotonics. *Nat. Photonics* **2018**, *12*, 659–670.
- (6) Christiansen, R. E.; Sigmund, O. Inverse design in photonics by topology optimization: tutorial. *J. Opt. Soc. Am. B* **2021**, *38*, 496–509.
- (7) Andrade, N. M.; Hooten, S.; Fortuna, S. A.; Han, K.; Yablonovitch, E.; Wu, M. C. Inverse design optimization for efficient coupling of an electrically injected optical antenna-LED to a single-mode waveguide. *Opt. Express* **2019**, *27*, 19802–19814.
- (8) Christiansen, R. E.; Lin, Z.; Roques-Carmes, C.; Salamin, Y.; Kooi, S. E.; Joannopoulos, J. D.; Soljačić, M.; Johnson, S. G. Fullwave Maxwell inverse design of axisymmetric, tunable, and multi-scale multi-wavelength metalenses. *Opt. Express* **2020**, *28*, 33854–33868.
- (9) Meem, M.; Banerji, S.; Pies, C.; Oberbiermann, T.; Majumder, A.; Sensale-Rodriguez, B.; Menon, R. Large-area, high-numerical-aperture multi-level diffractive lens via inverse design: erratum. *Optica* **2020**, *7*, 1323.
- (10) Piggott, A. G.; Lu, J.; Lagoudakis, K. G.; Petykiewicz, J.; Babinec, T. M.; Vučković, J. Inverse design and demonstration of a compact and broadband on-chip wavelength demultiplexer. *Nat. Photonics* **2015**, *9*, 374–377.
- (11) Borel, P. I.; Harpoth, A.; Frandsen, L. H.; Kristensen, M.; Shi, P.; Jensen, J. S.; Sigmund, O. Topology optimization and fabrication of photonic crystal structures. *Opt. Express* **2004**, *12*, 1996–2001.
- (12) Minkov, M.; Williamson, I. A. D.; Andreani, L. C.; Gerace, D.; Lou, B.; Song, A. Y.; Hughes, T. W.; Fan, S. Inverse Design of

Photonic Crystals through Automatic Differentiation. *ACS Photonics* **2020**, *7*, 1729–1741.

(13) Martin Burger Stanley, J.; Osher, E. Y. Inverse Problem Techniques for the Design of Photonic Crystals. *Transactions on Electronics* **2004**, *E87-C*, 258–265.

(14) Lin, Z.; Pick, A.; Lon čar, M.; Rodríguez, A. W. Enhanced Spontaneous Emission at Third-Order Dirac Exceptional Points in Inverse-Designed Photonic Crystals. *Phys. Rev. Lett.* **2016**, *117*, 107402.

(15) Jiang, J.; Chen, M.; Fan, J. A. Deep neural networks for the evaluation and design of photonic devices. *Nature Reviews Materials* **2021**, *6*, 679–700.

(16) Huang, Q.; Gan, L. T.; Fan, J. A. Conformal Volumetric Grayscale Metamaterials. *Adv. Mater.* **2023**, *35*, 2204688.

(17) Rodríguez, J. A.; Abdalla, A. I.; Wang, B.; Lou, B.; Fan, S.; Cappelli, M. A. Inverse Design of Plasma Metamaterial Devices for Optical Computing. *Physical Review Applied* **2021**, *16*, 014023.

(18) Nikkhah, V.; Pirmoradi, A.; Ashtiani, F.; Edwards, B.; Aflatouni, F.; Engheta, N. Inverse-designed low-index-contrast structures on a silicon photonics platform for vector–matrix multiplication. *Nat. Photonics* **2024**, *18*, 501–508.

(19) Pestourie, R.; Pérez-Arancibia, C.; Lin, Z.; Shin, W.; Capasso, F.; Johnson, S. G. Inverse design of large-area metasurfaces. *Optics Express* **2018**, *26*, 33732–33747.

(20) Chung, H.; Miller, O. D. Tunable Metasurface Inverse Design for 80% Switching Efficiencies and 144° Angular Deflection. *ACS Photonics* **2020**, *7*, 2236–2243.

(21) Lin, X.; Rivenson, Y.; Yardimci, N. T.; Veli, M.; Luo, Y.; Jarrahi, M.; Ozcan, A. All-optical machine learning using diffractive deep neural networks. *Science* **2018**, *361*, 1004–1008.

(22) Rodríguez, J. A.; Cappelli, M. A. Inverse design of plasma metamaterial devices with realistic elements. *J. Phys. D: Appl. Phys.* **2022**, *55*, 465203.

(23) Rodríguez, J. A.; Cappelli, M. A. Inverse design and experimental realization of plasma metamaterials. *Physical Review Applied* **2023**, *20*, 044017.

(24) Sakai, O.; Tachibana, K. Plasmas as metamaterials. *Plasma Sources Science and Technology* **2012**, *21*, 013001.

(25) Engheta, N. Pursuing near-zero response. *Science* **2013**, *340*, 286–287.

(26) Liberal, I.; Engheta, N. Near-zero refractive index photonics. *Nat. Photonics* **2017**, *11*, 149–158.

(27) Righetti, F.; Wang, B.; Cappelli, M. A. Enhanced Attenuation due to Lattice Resonances in a Two-Dimensional Plasma Photonic Crystal. *Physics of Plasmas* **2018**, *25*, 124502.

(28) Wang, B.; Cappelli, M. A. A Plasma Photonic Crystal Bandgap Device. *Appl. Phys. Lett.* **2016**, *108*, 161101.

(29) Wang, B.; Cappelli, M. A. Waveguiding and Bending Modes in a Plasma Photonic Crystal Bandgap Device. *AIP Advances* **2016**, *6*, 065015.

(30) Sakai, O.; Maeda, J.; Shimomura, T.; Urabe, K. Functional Composites of Plasmas and Metamaterials: Flexible Waveguides and Variable Attenuators with Controllable Phase Shift. *Physics of Plasmas* **2013**, *20*, 073506.

(31) Sakai, O.; Iwai, A.; Omura, Y.; Iio, S.; Naito, T. Wave Propagation in and Around Negative-Dielectric-Constant Discharge Plasma. *Physics of Plasmas* **2018**, *25*, 031901.

(32) Wang, B.; Rodríguez, J. A.; Cappelli, M. A. 3D Woodpile Structure Tunable Plasma Photonic Crystal. *Plasma Sources Science and Technology* **2019**, *28*, 02LT01.

(33) Houriez, L. S.; Bernety, H. M.; Rodríguez, J. A.; Wang, B.; Cappelli, M. A. Experimental Study of Electromagnetic Wave Scattering from a Gyrotropic Gaseous Plasma Column. *Appl. Phys. Lett.* **2022**, *120*, 223101.

(34) Mehrpour Bernety, H.; Houriez, L. S.; Rodríguez, J. A.; Wang, B.; Cappelli, M. A. A Characterization of Plasma Properties of a Heterogeneous Magnetized Low Pressure Discharge Column. *AIP Advances* **2022**, *12*, 115220.

(35) Mehrpour Bernety, H.; Houriez, L. S.; Rodríguez, J. A.; Wang, B.; Cappelli, M. A. A tunable microwave circulator based on a magnetized plasma as an active gyrotropic element. *Physics of Plasmas* **2022**, *29*, 112114.

(36) Pai, S.; Sun, Z.; Hughes, T. W.; Park, T.; Bartlett, B.; Williamson, I. A. D.; Minkov, M.; Milanizadeh, M.; Abebe, N.; Morichetti, F.; Melloni, A.; Fan, S.; Solgaard, O.; Miller, D. A. B. Experimentally realized in situ backpropagation for deep learning in photonic neural networks. *Science* **2023**, *380*, 398–404.

(37) Rodríguez, J. A. Electromagnetic Wave Manipulation with Plasma Metamaterials. Ph.D. Thesis; Stanford University, 2023; .

(38) Hagelaar, G. J. M.; Pitchford, L. C. Solving the Boltzmann Equation to Obtain Electron Transport Coefficients and Rate Coefficients for Fluid Models. *Plasma Sources Science and Technology* **2005**, *14*, 722–733.

(39) SIGLO Database *Electron-Neutral Collision cross-section Database Used with BOLSIG+*.

(40) Hilscher, A. Determination of the Cathode Fall Voltage in Fluorescent Lamps by Measurement of the Operating Voltage. *J. Phys. D: Appl. Phys.* **2002**, *35*, 1707–1712.

(41) Hughes, T. W.; Williamson, I. A. D.; Minkov, M.; Fan, S. Forward-Mode Differentiation of Maxwell's Equations. *ACS Photonics* **2019**, *6*, 3010–3016.

(42) Kingma, D. P.; Ba, J. *Adam: A Method for Stochastic Optimization*, 2014; .

(43) Frazier, P. I. *A Tutorial on Bayesian Optimization*, 2018; .

(44) Chiang, W.-H.; Mariotti, D.; Sankaran, R. M.; Eden, J. G.; Ostrikov, K. Microplasmas for advanced materials and devices. *Adv. Mater.* **2020**, *32*, 1905508.

(45) Guo, J.; Yu, L.; Xiang, H.; Zhao, Y.; Liu, C.; Dai, D. Realization of advanced passive silicon photonic devices with subwavelength grating structures developed by efficient inverse design. *Advanced Photonics Nexus* **2023**, *2*, 026005.

(46) Colburn, S.; Majumdar, A. Inverse design and flexible parameterization of meta-optics using algorithmic differentiation. *Communications Physics* **2021**, *4*, 65.

(47) Constantine, P. G. *Active Subspaces: Emerging Ideas for Dimension Reduction in Parameter Studies*; SIAM, 2015; .

(48) Hansen, N.; Müller, S. D.; Koumoutsakos, P. Reducing the Time Complexity of the Derandomized Evolution Strategy with Covariance Matrix Adaptation (CMA-ES). *Evolutionary Computation* **2003**, *11*, 1–18.

(49) Degrave, J.; et al. Magnetic control of tokamak plasmas through deep reinforcement learning. *Nature* **2022**, *602*, 414–419.

(50) Wang, B.; Rodríguez, J. A.; Miller, O.; Cappelli, M. A. Reconfigurable plasma-dielectric hybrid photonic crystal as a platform for electromagnetic wave manipulation and computing. *Physics of Plasmas* **2021**, *28*, 043502.

(51) Lou, B.; Rodríguez, J. A.; Wang, B.; Cappelli, M.; Fan, S. Inverse design of optical switch based on bilevel optimization inspired by meta-learning. *ACS photonics* **2023**, *10*, 1806–1812.

(52) Rodríguez, J.; Wang, B.; Cappelli, M. Dual-polarization Dirac cones in a simple 2D square lattice photonic crystal. *Opt. Lett.* **2020**, *45*, 2486–2489.

(53) Haarnoja, T.; Zhou, A.; Abbeel, P.; Levine, S. In *International Conference on Machine Learning*, 2018; pp 1861–1870.

(54) Nomura, M.; Shibata, M. CMAES: A Simple yet Practical Python Library for CMA-ES. *arXiv preprint arXiv:2402.01373* 2024, (accessed 2025–02–14, 2025) .