
MERGING AUTOMATIC DIFFERENTIATION AND THE ADJOINT METHOD FOR PHOTONIC INVERSE DESIGN

A PREPRINT

Alexander Luce
University Erlangen-Nürnberg &
ams OSRAM Group
Regensburg

Rasoul Alaei
ams OSRAM Group
Regensburg

Fabian Knorr
ams OSRAM Group
Regensburg

Florian Marquardt
University Erlangen-Nürnberg &
Max Planck Institute
for the Science of Light
Erlangen

02/10/2023

ABSTRACT

Optimizing shapes and topology of physical devices is crucial for both scientific and technological advancements, given its wide-ranging implications across numerous industries and research areas. Innovations in shape and topology optimization have been seen across a wide range of fields, notably structural mechanics, fluid mechanics, and photonics. Gradient-based inverse design techniques have been particularly successful for photonic and optical problems, resulting in integrated, miniaturized hardware that has set new standards in device performance. To calculate the gradients, there are typically two approaches: implementing specialized solvers using automatic differentiation or deriving analytical solutions for gradient calculation and adjoint sources by hand. In this work, we propose a middle ground and present a hybrid approach that leverages and enables the benefits of automatic differentiation and machine learning frameworks for handling gradient derivation while using existing, proven solvers for numerical solutions. Utilizing the adjoint method, we turn existing numerical solvers differentiable and seamlessly integrate them into an automatic differentiation framework. Further, this enables users to integrate the optimization environment with machine learning applications which could lead to better photonic design workflows. We illustrate the approach through two distinct examples: optimizing the Purcell factor of a magnetic dipole in the vicinity of an optical nanocavity and enhancing the light extraction efficiency of a μ LED.

Keywords Shape optimization · Adjoint method · Automatic differentiation · Gradient-based optimization · Light extraction efficiency · Nanophotonic devices · Finite-difference time-domain (FDTD) · Outcoupling structures

1 Introduction

Optimization is a crucial aspect in the development of structural devices that dictate the physical properties of waves and fields in order to yield higher performance compared to those created using traditional approaches. The application domain for optimization for light is vast and rapidly evolving [1], encompassing numerous techniques that modify parameters or geometries based on specific update algorithms. Generally, physical systems subject to optimization do not exhibit a convex loss landscape, resulting in inherent limitations when seeking global optima. Consequently, most optimization algorithms can be categorized as either global or local optimization. Global optimization techniques for photonic optimization include evolutionary algorithms [2] and Bayesian optimization techniques [3, 4]. Although global optima are typically preferred over local optima, these algorithms face significant limitations regarding their

applicability. For instance, Bayesian optimization becomes increasingly expensive for large problems with many parameters and data points [5]. Similarly, evolutionary algorithms are often sample-inefficient and unsuitable for computationally expensive evaluations. Emerging machine-learning approaches offer new possibilities for global optimization and have demonstrated promising results in photonic optimization applications [6, 7, 8, 9, 10]. However, they cannot overcome the fundamental issue of the curse of dimensionality [11]. Conversely, identifying local optima in high-dimensional problems is much more manageable than obtaining a global minimum [12], as evidenced by the vast number of neural network parameters being optimized during training for deep learning and machine learning [13]. By utilizing gradient-based optimization and adaptive stepsize methods such as ADAM [14] or line-search [15], it is feasible to optimize numerous parameters simultaneously and achieve a local optimum. Gradient-based optimization is often applied in numerical optimization algorithms for device parameter, shape, or topology optimization, which is commonly referred to as inverse design [16, 17, 18, 19, 20]. To employ gradient-based optimization, it is essential to compute the gradients of the desired function with respect to a loss or target value. This computation can be challenging, as the loss often depends on the solution of the system of equations governing the underlying physical problem. The adjoint method [21, 22] allows for obtaining analytical gradients by deriving adjoint equations, which can then be manually integrated into numerical solvers and update equations for the geometry parameters [23, 18, 17, 24, 25]. However, this approach can be tedious and problem-specific, requiring new derivations for different physical settings or optimization targets. Automatic differentiation (AD) offers an alternative by automating the complex and elaborate process of deriving gradients [26]. With AD, one only needs to implement the forward function, provide a suitable parameterization, define functions for solving the governing partial differential equation (PDE), postprocess the PDE solution, and define the target/loss function. Although implementing functions for the postprocessing and the loss is typically straightforward, obtaining gradients of a physical solution for a PDE can be challenging with AD, especially when established numerical solvers do not support it. Consequently, an end-to-end approach using AD requires implementing a new solver directly within the AD framework to derive the PDE solution and perform backpropagation [27, 28, 29, 30]. This task can be impractical since migrating an existing, validated physical model to an AD solver and framework poses significant work overhead.

In this work, we propose a "hybrid approach" that merges the benefits of automatic differentiation (AD) with the applicability of the adjoint equation to established solvers. By utilizing analytically derived results from previous works on the physical adjoint problem, we directly integrate established numerical solvers into an AD framework. Importantly, the internal workings of the numerical solver must not be accessible to the user. Hence, we consider the adjoint computation to be an atomic step in the computational graph of the AD framework. This combination creates an end-to-end AD-enabled process incorporating established numerical solvers. The solvers are then seamlessly integrated into the computational graph of the AD framework, effectively rendering them auto-differentiable for the optimization. This approach allows users to leverage the functionality and efficiency of modern AD frameworks while selecting the optimal numerical solver for their specific problem, regardless of AD compatibility¹.

We demonstrate the application of AD integration by performing shape optimization on two distinct photonic problems of scientific and engineering interest. In the first example, we aim to enhance the Purcell factor of a photonic nanocavity by deforming the cavity geometry. In the second example, we apply shape optimization to the outcoupling structure of a μ LED. In both cases, we implement the analytically derived equations for the shape gradient into PyTorch, while utilizing PyTorch-provided functionality for postprocessing and loss calculations [33].

2 Combining the adjoint method with autodifferentiation

A general optimization problem for photonic applications is the overarching goal to achieve the highest possible value of a physical property. The problem is typically described by a partial differential equation (PDE) A that governs the dynamics. Here, we consider only linear PDEs for brevity, but in general, the dynamics could also encompass non-linear equations or be described by an eigenvalue problem [21]. With a set of geometrical design parameters p , the solution u of the physical system is given by the equation $A_p u = b_p$ where b_p denotes source terms. The system under consideration is embedded in a simulation region \mathcal{D} on which the solution is computed. A figure of merit or loss $J(u)$ must also be defined which evaluates the solution of the physical system given by the parameters p . Typically, this loss functional is given by an integral over the computational domain $J = \int_{\mathcal{D}} u(s) ds$ or a sum over discrete physical properties of the solution. The optimization problem is formulated by $\min_p J(u_p)$ such that $A_p u = b_p$. Typical examples are increasing the quality factor of a cavity, focusing the emission of light into a particular solid angle, or increasing the Purcell effect for an emitter [34, 2]. The problem of computing the optical characteristics for a given task is usually tackled by employing various numerical solvers such as rigorous coupled wave analysis (RCWA), finite difference time domain (FDTD), finite-difference frequency-domain (FDFD), or finite element method (FEM).

¹It is crucial for the solver to provide an interface that enables loading external geometries or parameters, and sources, as well as exporting numerical solutions, which typically is a supported functionality [31, 32].

Depending on the type of problem, the approximations and discretizations performed by the software are particularly useful for a specific problem. As a result, many specialized types of solvers and software solutions exist today. However, the situation is usually not as straightforward as choosing a set of parameters that can be applied to the PDE and then be evaluated directly. For more complex problems, geometry generation and postprocessing of the raw solution of the PDE require additional work on top of finding the solution. The evaluation of the design of an optimization process can be separated into the following steps: the geometry definition (Figure 1a), the numerical simulation of the physical problem (Figure 1b), the postprocessing of the results (Figure 1c), and the evaluation by a loss functional (Figure 1d).

Optimizing a large set of parameters p by gradient-based optimization requires the gradient of the loss functional with respect to the design parameters $\delta J(u)$. Obtaining gradients of a loss functional J with respect to the input parameters p is difficult at first glance since it involves computing the variation of the loss functional with respect to all input parameters. This would lead to a computational complexity that scales linearly with the number of input parameters. For large problems, this poses a heavy computational burden. Fortunately, the computational complexity can be improved to a constant dependency on the input parameters by either employing the adjoint method or using (backward mode) automatic differentiation to compute the parameter gradients. Both approaches are elucidated and combined in the following.

2.1 Automatic differentiation

Backward mode automatic differentiation or backpropagation is the idea of applying the chain rule algorithmically to a numerically executed computation. The gradient is separated into atomic computations for which the derivatives are known analytically. For a simple computation $y = f(g(x))$, the chain rule gives the derivative of y with respect to x as $\frac{\partial y}{\partial x} = \frac{\partial f}{\partial g} \frac{\partial g}{\partial x}$. The functions f and g can now be differentiated individually and the analytical solution of the respective derivative reused whenever the functions appear again. For arbitrarily big computations, it suffices to separate the computation into individual function applications and keeping track of the order in which all functions have been applied. This concept is known as the computational graph. Consider a complicated computational graph $J = f_G(p)$ that represents solving the linear PDE of the aforementioned physical system $A_p u = b_p$ and computing the loss of the system $J(u)$. Computing $\frac{\partial J(u)}{\partial p}$ is then reduced to tracing back the individual steps of the computational graph from the loss J to the parameters p and chaining the respective derivatives of the intermediate function applications. To distinguish which part of the computation is executed, any mathematical function of the AD framework provides a forward and a backward method. The normal function evaluation is applied when the forward method is invoked since the computational graph is traversed in the forward direction from the input to the loss. The derivative is computed by the backward method which receives the gradient information in backward order starting from the loss and propagating to the input values. Automatic differentiation is straightforward to use and many scientific AD frameworks are readily available such as Jax, PyTorch, Hips/autograd, and more [35, 33, 36, 37, 38]. The drawback of using AD frameworks for optimization is that the numerical solver for the PDE must support automatic differentiation and create a computational graph during the numerical solution for u . Although this is a rapidly evolving field and a number of solvers employing automatic differentiation have been developed [28, 27, 30], many popular choices and industry-standard solvers do not support AD [31, 32].

2.2 Adjoint Method

On the other hand, adjoint methods approach the gradient derivation from the manual side. Here, we briefly recall the basis of the adjoint method [21, 23, 18]. Consider again a system governed by the linear PDE $A_p u = b_p$ with loss functional J . The functional derivative of $J(u)$ can be expressed by $\frac{dJ}{dp} = \frac{\partial J}{\partial p} + \frac{\partial J}{\partial u} \frac{\partial u}{\partial p}$. The term $\frac{\partial u}{\partial p}$ exhibits the undesired effect of linear scaling of the computational cost with the number of input parameters p if it is evaluated via finite differences. However, this term can be expressed by taking the derivative of the PDE wrt. to p $\frac{\partial A u}{\partial p} - \frac{\partial b}{\partial p} = \frac{\partial A}{\partial p} u + A \frac{\partial u}{\partial p} - \frac{\partial b}{\partial p} = 0$ which is then rearranged to

$$\frac{\partial u}{\partial p} = A^{-1} \left(\frac{\partial b}{\partial p} - \frac{\partial A}{\partial p} u \right). \quad (1)$$

Putting everything together, we are given an equation that splits into three parts - the derivative of the loss with respect to the parameter, the s.c. *adjoint solution* and the *forward solution*:

$$\frac{dJ(u)}{dp} = \frac{\partial J}{\partial p} + \underbrace{\left(\frac{\partial J}{\partial u} A^{-1} \right)}_{\text{adjoint solution}} \underbrace{\left(\frac{\partial b}{\partial p} - \frac{\partial A}{\partial p} u \right)}_{\text{forward solution}}. \quad (2)$$

Here, the name adjoint solution derives from the reformulation of the left side of Equation 2 to $A^\dagger v = \frac{\partial J}{\partial u}$ where v is the solution of this adjoint PDE. For many linear PDE, the adjoint system equations A^\dagger are straightforward to derive. The boundary terms are given by the sensitivity of the loss functional with respect to the solution of the original PDE. The right side of Equation 2 is then comprised of the solution u and the sensitivity of the system matrix A and boundary conditions of the original PDE. The adjoint method is appealing due to its generality. It is applicable to physical optimization problems in many settings, also outside of photonics. In particular, it is possible to compute the forward and adjoint solution with many different types of numerical solvers as long as the adjoint system equations can be used by the solver. However, it poses an analytical overhead before the optimization since the derivation of the required equations is done by hand. Especially for complicated problems where the system equations have a complicated dependency on the parameters in $\frac{\partial A}{\partial p}$ and $\frac{\partial b}{\partial p}$ or if the loss functional involves lengthy and tedious postprocessing $J = f \circ g \circ h \dots$ the adjoint method becomes impractical.

2.3 Integrating the adjoint method into automatic differentiation

Interestingly, the advantages and disadvantages of AD and the adjoint method seem to complement each other. While the adjoint method is difficult to use but generally compatible with most numerical solvers, AD is easy to use but only applicable by using appropriate solvers. Here, we show how to combine both methods and leverage the advantages of each other to cancel their disadvantages.

The key idea is to incorporate the adjoint method directly into the computational graph of the AD framework. This integration provides the benefit of the AD framework’s flexibility without the necessity of rewriting efficient numerical solvers. In order to integrate the adjoint method into the backward calculation of an AD framework, we need to identify the appropriate terms in the derivation. Fortunately, the forward and adjoint solution derivation is similar to the distinction between the forward and backward methods for AD. The forward method should receive the input parameters p that determine the system equations A_p and source terms b_p and return the solution of the PDE u back to the computational graph. During gradient computation, we crucially depend on the input from the backward methods to receive the *adjoint source* $\frac{\partial J}{\partial u}$ with which the adjoint solution v can be computed. In the backward method, we should therefore receive gradient information from the loss function and any other postprocessing steps that were computed from the forward solution u . Then the adjoint solution and the forward solution are multiplied (Equation 2) and the solution is returned back to the computational graph for further processing back to the root of the parameters. The remaining terms $\frac{\partial A}{\partial p}$ and $\frac{\partial b}{\partial p}$ depend on the discretization of the solver and the applied optimization scheme and must be treated accordingly if they are not part of the AD framework.

There are two popular schemes that are mostly used for the geometric optimization of photonics - topology and shape optimization. In topology optimization, the entire distribution of material is considered to change point-wise throughout the optimization domain. In shape optimization, the boundary $\partial\Omega$ of a shape Ω is continuously deformed but the shape remains connected during the deformation. Importantly, the approaches require different treatments of the gradients. We will focus on shape optimization in the following but similar steps can be applied for topology optimization [39]. The optimization target must be reformulated slightly since in shape optimization the target is to optimize over the possible geometrical shapes instead of the parameters. A shape Ω is a connected region within the computational domain $\mathcal{D} \subset \mathbb{R}^n$ with fixed optical material properties. The PDE can then be solved on \mathcal{D} which yields solution $u_\Omega \in \mathbb{R}^k$. A loss functional J is then applied to evaluate the solution. Shape optimization derives how to deform the boundary of the shape Ω to improve the loss $J(u_\Omega) = \int_{\mathcal{D}} ds u_\Omega$. By taking the variation of the loss functional to first order, we obtain [23, 20, 40]

$$\delta J_\Omega(\delta\Omega) = \int_{\partial\Omega} ds \delta\Omega \hat{n} (c_1 - c_2) u_\Omega v_\Omega = \int_{\partial\Omega} ds \delta\Omega \hat{n} V_\Omega(s). \tag{3}$$

$\delta\Omega$ denotes the variation of the shape which is equivalent to a test function in functional analysis. At iteration i , the variation deforms the shape $\Omega_{i+1} = (\mathbb{1} + \delta\Omega)(\Omega_i)$. \hat{n} denotes the normal vector on the boundary, and $V_\Omega(s)$ denotes the s.c. sensitivity field (also known as gradient field [17] or velocity field [41]). Since our goal is to minimize the loss functional, we see that this can be achieved by setting the geometry deformation to $\delta\Omega = -\hat{n} V_\Omega(s)$. Then, the loss functional is guaranteed to decrease to first order in every iteration. The sensitivity field acts on the shape as a vector field that drags the boundary along the direction of the vector field [42, 41].

The forward and adjoint solutions u_Ω and v_Ω are directly inserted in the shape gradient. c_1 and c_2 are parameters of the computational domain inside and outside of the shape Ω . For photonic optimization, the parameters are given by the relative electric permittivity ε_i for the parallel component of the electric fields in u and v and $1/\varepsilon_i$ for the normal component [17].

In many interesting scenarios, the functional $J(u_\Omega)$ has a complicated dependence on the solution from the post-processing and the shape has parameter dependencies. Generally, this postprocessing dependence is described by a

function $f : \mathbb{R}^k \times \mathcal{D} \mapsto \mathcal{F}$, where \mathcal{F} is a vector space, usually chosen to be \mathbb{R}^m . The postprocessing, which can be arbitrarily complicated, acts on the solution u_Ω before integrating the result of the postprocessing $J(f(u_{\Omega(p)})) = \int_{\mathcal{F}} dy f(u_{\Omega(p)})$ to obtain the loss J . An example of such a postprocessing function is the farfield transformation that is used in subsection 3.2 that projects the boundary values of the solution u_Ω from the 2D computational domain \mathcal{D} on a linear farfield [43]. The chain rule for functionals [44, 45] for chaining functionals with ordinary differentiable functions allows us to separate the postprocessing and the geometry definition from the shape gradients

$$\delta J_\Omega(\delta\Omega) = \int_{\partial\Omega} ds \delta\Omega \hat{n} V_{\Omega(p)}(s) \frac{\partial\Omega(p)}{\partial p}. \quad (4)$$

The functional derivative can then be written as

$$\frac{\delta J_\Omega}{\delta p} = \hat{n} (c_1 - c_2) u_\Omega \underbrace{\left(\frac{\partial J}{\partial f} \frac{\partial f}{\partial u_{\Omega(p)}} A^{-1} \right)}_{v_\Omega} \frac{\partial\Omega(p)}{\partial p}. \quad (5)$$

Here, we see again how to employ automatic differentiation on the adjoint source computation for the adjoint solution v_Ω and then continue with the backpropagation of the shape. For illustrative purposes, the process is also depicted in Figure 1.

2.4 Software considerations

On a practical level, integrating the adjoint method into an AD framework and thus making AD compatible with conventional solvers boils down to implementing a *differentiable simulation* function that receives a numerical representation of the geometry boundary $\partial\Omega$. The differentiable simulation is shown in pseudocode in subsection 5.2. Boundary support points that give rise to the shape Ω are an ideal representation since they easily integrate with numerical AD frameworks. Here, we focus only on 2D shapes but this can be extended to 3D by modifying the equations appropriately. The differentiable simulation method starts by initializing the simulation by evaluating the geometry support points. Then, the solution of the simulation is computed and returned to the AD framework. The AD framework then undertakes the postprocessing of the solution, along with the evaluation of the loss functional. Taking the derivative of the loss functional up to the adjoint simulation is handled by the AD framework using backpropagation. Crucially, the AD framework returns the *adjoint source* $\frac{\partial J}{\partial u}$, see Figure 1e-f), which is usually derived manually in former applications [17, 25, 18, 23].

The adjoint source is then passed to the *backward* method of the *differentiable simulation* function as the *gradient input*, shown in Figure 1g). It initializes the adjoint simulation and computes the adjoint solution v , which is particularly easy for Maxwell’s equations with linear materials due to their time-reversal properties [18]. For the Maxwell equations, the adjoint system is given by $A^\dagger v_\Omega = T A T v_\Omega = \frac{\partial J}{\partial u}$ with T , the time reversal operator. Then, computing the adjoint solution can be done by solving $A T v_\Omega = T \frac{\partial J}{\partial u}$.

Taking the derivative of a function often involves retaining the results from the forward function evaluation. In the case of the adjoint method, most importantly the solution u_Ω but it is useful to also save the boundary reference and other additional parameters for the solver. Since AD frameworks are often required to store the solution of the forward pass, the framework provides the functionality to store intermediate results which are necessary for the gradient calculation during backpropagation. Together with the forward solution, the backward method computes the sensitivity field for the given geometry with the shape calculus shown in Equation 5 [46, 23, 18].

The sensitivity field will act as a gradient by deforming the geometry since $\delta\Omega$ will decrease the loss functional to first order. The movement of the boundary is projected on the normal since the tangential movement of the boundary has no influence on the loss functional which is shown in Equation 5. We need to take into account that the geometry is represented by boundary support points where movement of the support points drags the connecting edge along. The edge displacement must be carefully taken into account which is detailed in subsection 5.1. The support point sensitivities are interpreted as the support point gradients and returned by the backward method. The AD framework continues with the backpropagation to the geometry parameters p . In this way, it is particularly straightforward to create parameterized geometries which, for example, serve to introduce geometry constraints that ensure favorable properties such as manufacturability. To illustrate the core functionality of the *differentiable simulation* function we present it in pseudo-code in subsection 5.2.

Finally, the shape can be updated via an optimizer depicted in Figure 1h). The optimizer uses the computed gradients in order to update the shape parameters with a stepsize estimated by the selected algorithm. Many different approaches exist and exhibit advantages and disadvantages. The simplest optimizer is gradient descent where the parameters are updated based on a stepsize η selected at the start following the rule $p_{i+1} = p_i \pm \eta \nabla p$. However, many more refined




Figure 1: Overview of a typical geometry optimization problem. **a)** The optimization starts with a parametrized initial shape $\Omega(p)$ within a computational domain \mathcal{D} . **b)** Then, the physical problem is solved using a numerical solver with the source distribution appropriate for the problem at hand. **c)** Next, the physical solution is evaluated in the domain or on its boundary and transformed during postprocessing, for example by projecting a recorded near field to the far field or selecting different waveguide modes. **d)** Finally, the result from postprocessing is evaluated and integrated by a loss functional J . **e-f)** Differentiating the loss functional up to the solution from the numerical solver is easy due to automatic differentiation (AD). **g)** However, it becomes more difficult to obtain gradients for the simulation parameters from the simulation solution since many numerical simulations do not provide AD or it is not efficient to use. Taking the derivative with respect to the shape parameters is therefore difficult. By combining AD with the adjoint method, backpropagation computes the gradients up to the solution from the solver at which point the adjoint method is employed to continue with the gradient computation and passes the gradients further to the shape parameters. Computing the gradient of a given shape is then reduced to backpropagating through the computational graph of the AD framework, which is equivalent to backpropagation through the simulation itself by means of the adjoint method. **h)** Finally, the gradient of the shape is evaluated by AD, and the shape parameters are updated by an optimization algorithm. Then the next iteration begins with the updated shape.

techniques exist such as quasi-newton methods [15] that estimate the hessian iteratively with gradient information or moment estimation methods such as ADAM [14] which approximate first and/or second-order moments of the stepsize.

3 Application examples

To showcase the integration of a conventional simulation into an AD framework, we apply it to two different problems of current interest. In the first example, we enhance the spontaneous emission rate for an optical nanocavity while in the second example, we increase the farfield intensity distribution within a given solid angle by optimizing the outcoupling structure of a 2D μ LED.

3.1 Purcell effect optimization

In the example shown in Figure 2, we aim to increase the spontaneous emission rate for an optical nanocavity. Increasing the spontaneous emission rate is a research problem both for inverse design [47, 48] and classical approaches [49]. More precisely, our goal is to increase the Purcell factor, which is proportional to the time-averaged poynnting vector on the boundary of the domain $\max \frac{P}{P_0} = \max \int_{\partial\mathcal{D}} \hat{\mathbf{n}} \text{Re}[\mathbf{E}_\Omega \times \mathbf{H}_\Omega^\dagger] / 2 \text{ ds}$ where P_0 is the dipole emission in free space,

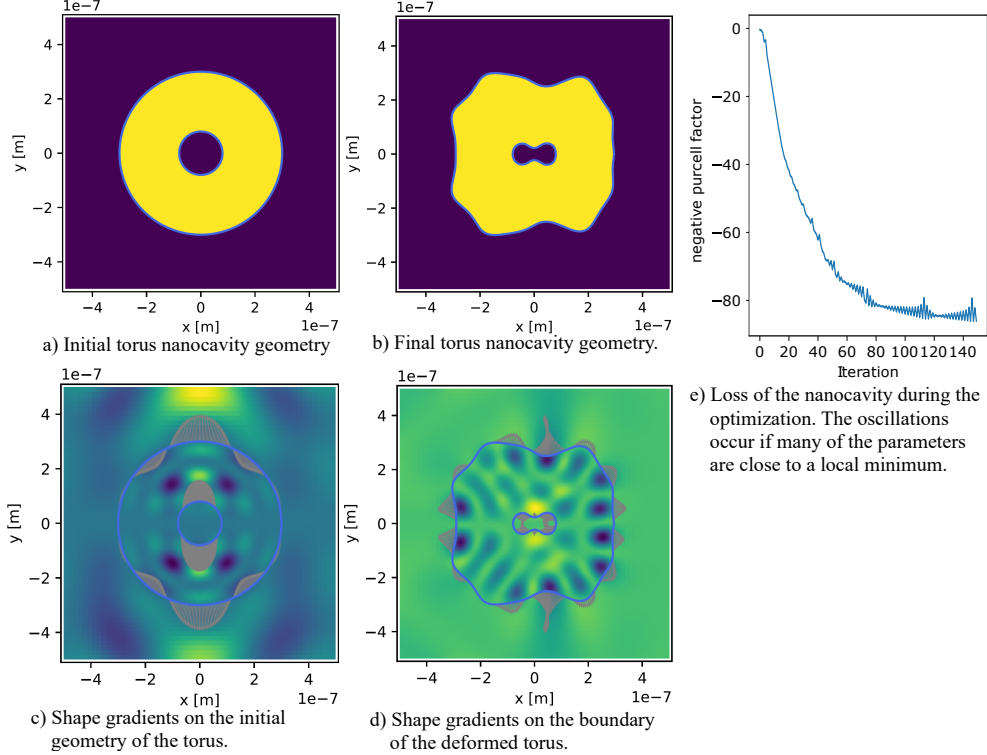


Figure 2: Overview of the gradient descent optimization results for the enhancement of the Purcell factor in a torus nanocavity. **a)** depicts the initial torus geometry. The dipole emitter sits in the center of the torus at $\vec{r} = (0, 0)^T$. The refractive index of the torus has a value of $n = 3.5 + 0i$. The blue contour represents the deformable boundary which is subject to optimization. The optimized torus geometry after 150 iterations is shown in **b)**. Geometry deformations during optimization led to a loss reduction, as demonstrated in **e)**. The loss exhibits a steady decline throughout the optimization process with slight oscillations which emerge if large parts of the boundary are close to a local minimum. For this optimization, we chose the negative Purcell factor for the loss with gradient descent. After the optimization, the Purcell factor was increased by about $\times 220$. The deformations are computed via the adjoint method, see section 2. The sensitivity field is shown in **d)** together with the shape gradients that are obtained by evaluating the sensitivity field on the boundary of the deformable geometry.

which is used for normalization while \hat{n} denotes the boundary normals and \mathbf{E}_Ω and \mathbf{H}_Ω denote electric and magnetic fields on the boundary Ω . We employ the FDTD solver from Lumerical [31] to obtain the electric and magnetic field solutions for the selected wavelength of 600nm. For simplicity and to reduce the simulation time, we consider a 2D problem that is infinitely extended in the z -direction. In the center of the simulation domain, we place a magnetic dipole emitter with the magnetic current oriented in the x -direction and surround it with a torus structure with a real refractive index of $n = 3.5$. Both the inner and the outer torus boundary, shown in Figure 2-, are subject to optimization but are represented by a tensor of an AD framework. In the presented case, we employ PyTorch as AD framework. For the update scheme, we employ simple gradient descent but chose a physically inspired step size on a length scale for which we expected a change of the Purcell factor. We also decrease the step size proportionally to the Purcell factor. The optimization shows converging behavior after 150 iterations, in which the Purcell factor increased to $\frac{P}{P_0} \approx 83$ starting from $\frac{P}{P_0} \approx 0.4$, which is in the range of expected improvement. Since we use a simple gradient update scheme, the presented solution is highly dependent on the initial geometry and is potentially far away from a global maximum for the Purcell factor. The optimization also results in a dumbbell shape around the dipole which has been observed in other works, too [49, 50, 47, 51, 48].

3.2 Shape optimization for spatially distributed dipole emission of a μ LED

In the example shown in Figure 3, our primary objective is to enhance the farfield emission of a μ LED (nanoscale light-emitting diode) within a specific solid angle, denoted as Γ , also known as LEE (light extraction efficiency). The field of μ LED development is rapidly advancing, attracting substantial research interest and holding considerable industrial relevance [52, 53, 54, 55]. The targeted optimization can be mathematically represented by the function

$\max_{\Omega} \text{LEE}_{\Gamma} = \max_{\Omega} P_{\Gamma} / P_0 = \max_{\Omega} \frac{1}{P_0} \int_{\Gamma} \int_{\Lambda} \text{Re}[E_{\Omega, \text{Farfield}} \times H_{\Omega, \text{Farfield}}^{\dagger}] / 2 d\Gamma d\lambda$. LEDs exhibit incoherent dipole emissions that originate from all over the quantum well region. To approximate this emission behavior, we distribute dipole emitters over the quantum well and compute the average emitted intensity within the solid angle Γ and normalize with respect to the injected power P_0 into the simulation. By solving the emission problem for each individual dipole, we can account for the incoherent nature of the μLED 's emission. To simplify the problem and reduce computational time, we adopt a two-dimensional FDTD model. The wavelength range of interest for this example is from 600 to 650 nanometers ($\Lambda = [600, 650]$ nm). The initial μLED model is depicted in Figure 3a). The μLED features a gold substrate passivated with a thin layer of silicon dioxide and a semiconductor material composed of gallium nitride, which is connected to the gold substrate via a thin layer of indium tin oxide. The quantum well region is made up of indium gallium phosphide. To enhance the light extraction efficiency within the solid angle Γ , we deform the upper boundary of the μLED 's top side, referred to as the outcoupling structure, which consists of gallium nitride. The emission then radiates into the air.

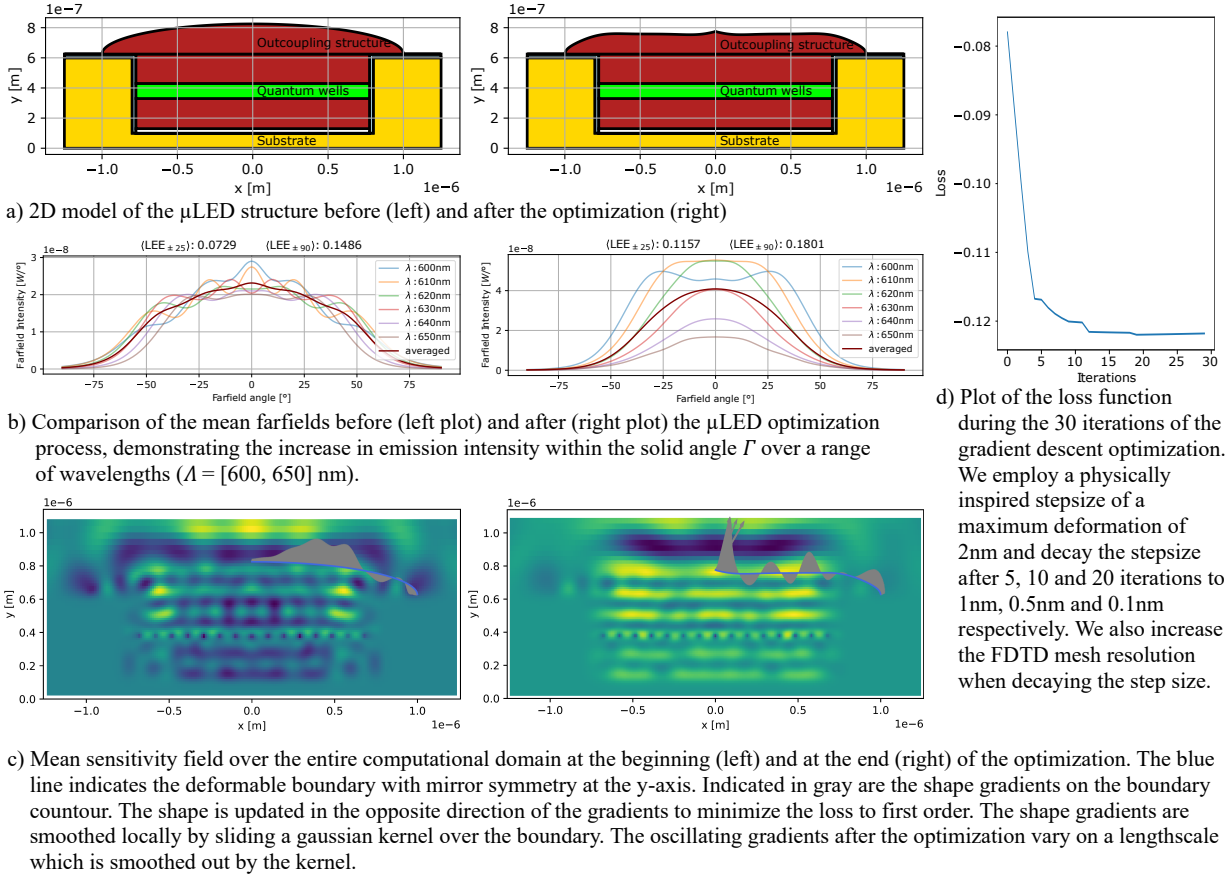


Figure 3: Illustration of the μLED optimization process: **a)** presents the initial model, where dipole emitters are distributed within the quantum well region (green). The averaged farfield emission of the initial μLED and LEE into solid angle $\Gamma = 25^\circ$ is shown on the left plot of **b)**. For reference, **b)** also displays the farfield for a range of individual wavelengths and the total LEE. Using the adjoint method with shape calculus section 2 we obtain the averaged sensitivity field over the entire domain, which is shown in **c)**. By following the direction of the steepest descent and updating the boundary geometry according to the shape gradients for 30 iterations, we deform the boundary and decrease the loss over the optimization duration, see **d)**. To avoid obtaining a shape with very small features, we smooth the gradients locally over the deformable boundary. At the end of the optimization, the outcoupling structure of the μLED is shown in **a)** on the right. The corresponding farfield for the optimized μLED shown on the right in **b)** experiences a reshaping with an improvement of the overall LEE directed into Γ of $\Delta\text{LEE}_{\pm 25} = 0.0428$ and a total LEE improvement of $\Delta\text{LEE}_{\pm 90} = 0.0315$.

To enable the use of automatic differentiation, we represent the outcoupling structure geometry with a PyTorch tensor. The differentiable simulation receives a reference to this boundary representation and creates an STL file from the tensor that is imported to the solver after each iteration. Furthermore, the projection of the fields on the boundary of the μLED are projected into the farfield by a PyTorch implementation of the equivalence principle in order to use AD

for the postprocessing [43]. Employing a standard gradient descent optimization technique, we run the optimization for 30 iterations. During the optimization, the outcoupling structure is continuously deformed following the mean sensitivity field computed by the adjoint method. The shape gradients and mean sensitivity field for the initial and final structure are shown in Figure 3 a) in the left and right figure respectively. The boundary approximation of the outcoupling structure is much finer than the mesh size of the solver, thus we interpolate the field solution to obtain smooth field values on the boundary. For optimization stability, we decay the stepsize after 5, 10, and 20 iterations from 2nm to 1nm, 0.5nm, and 0.1nm. By increasing the mesh resolution, after 10 and 20 iterations, we make sure to resolve the geometry changes sufficiently when the stepsize becomes small. To avoid creating small features on the outcoupling structure boundary, we smooth the shape gradients locally by sliding a Gaussian kernel over the boundary. Therefore, the standard deviation of the kernel controls the size of the features. This explains the oscillating gradients on the boundary after the optimization (see Figure 3 c) plot on the right), which are smoothed out after applying the Gaussian kernel. The final optimized μ LED structure is presented in Figure 3 a), while the loss throughout the optimization process is depicted in Figure 3 d). As the main objective is proportional to the farfield intensity within the solid angle Γ , we also provide the mean farfield at both the beginning and end of the optimization in Figure 3 b). Additionally, we show the farfield intensity for a range of wavelengths. After the optimization, the averaged farfield is focused in the target solid angle Γ , and the LEE_{Γ} is increased by 0.0428, an improvement of 63.01% compared to the initial farfield.

4 Conclusion

In this work, we have presented an general approach to include existing numerical solvers into an automatic differentiation framework to simplify and make the optimization of photonic structures faster and more convenient, particularly with existing models and solvers. To this end, the adjoint method is the key to implementing the forward and backward methods for automatic differentiation which allows us to make conventional solvers end-to-end auto-differentiable. Due to our focus on continuous geometries, we presented the optimization in the context of shape optimization and computed shape gradients on geometry boundaries which enable gradient-based optimization algorithms to improve the optical characteristics of the geometry with respect to the loss function. Generally, the approach is also suitable for topology optimization.

We demonstrate the approach with two different physical problems, optimizing an optical nanocavity to increase the spontaneous emission rate and optimizing the outcoupling structure of a μ LED to increase the light extraction efficiency into a solid angle in the farfield. For both cases, we show a significant reduction of the loss while employing Lumerical FDTD to solve the Maxwell equations.

Two key advantages of employing automatic differentiation (AD) for numerical optimization are parallelized multiphysics optimization and compatibility with machine learning. By integrating various physics solvers into AD, conducting joint optimization (e.g., thermal and optical optimization) becomes considerably more straightforward. Integrating the adjoint method into AD frameworks could also facilitate the development of novel AI applications for designing and optimizing optical devices.

Acknowledgments

We thank the supporters of this work, particularly Heribert Wankerl (ams-OSRAM Group) and Maike Stern (OTH Regensburg) for constructive feedback of this manuscript. Furthermore, we thank Daniel Grünbaum (ams-OSRAM Group) and Philipp Schwarz (ams-OSRAM Group & University Regensburg) for fruitful and helpful discussions. Finally, we thank Harald Laux (ams-OSRAM Group) for his organizational support.

Publication Funding Acknowledgments

This work was funded by the ams-OSRAM Group.

References

- [1] K. Y. Bliokh, E. Karimi, M. J. Padgett, M. A. Alonso, M. R. Dennis, A. Dudley, A. Forbes, S. Zahedpour, S. W. Hancock, H. M. Milchberg, S. Rotter, F. Nori, S. K. Özdemir, N. Bender, H. Cao, P. B. Corkum, C. Hernández-García, H. Ren, Y. Kivshar, M. G. Silveirinha, N. Engheta, A. Rauschenbeutel, P. Schneeweiss, J. Volz, D. Leykam, D. A. Smirnova, K. Rong, B. Wang, E. Hasman, M. F. Picardi, A. V. Zayats, F. J. Rodríguez-Fortuño, C. Yang, J. Ren, A. B. Khanikaev, A. Alù, E. Brasselet, M. Shats, J. Verbeeck, P. Schattschneider, D. Sarenac, D. G. Cory, D. Pushin, M. Birk, A. Gorlach, I. Kaminer, F. Cardano, L. Marrucci, M. Krenn, and F. Marquardt. Roadmap on structured waves, 2023.

- [2] Eric Johlin, Sander Mann, Sachin Kasture, and A. Koenderink. Broadband highly directive 3d nanophotonic lenses. *Nature Communications*, 9, 11 2018.
- [3] Heribert Wankerl, Christopher Wiesmann, Laura Kreiner, Rainer Butendeich, Alexander Luce, Sandra Sobczyk, Maike Lorena Stern, and Elmar Wolfgang Lang. Directional emission of white light via selective amplification of photon recycling and bayesian optimization of multi-layer thin films. *Scientific Reports*, 12(1):5226, Mar 2022.
- [4] Xavier García Santiago. *Numerical methods for shape optimization of photonic nanostructures*. PhD thesis, Karlsruher Institut für Technologie (KIT), 2021. 43.32.02; LK 01.
- [5] Jost Tobias Springenberg, Aaron Klein, Stefan Falkner, and Frank Hutter. Bayesian optimization with robust bayesian neural networks. In D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc., 2016.
- [6] Jiang An Qing, Yoshie Osamu, and Liangyao Chen. Multilayer optical thin film design with deep q learning. *Scientific Reports*, 10, 07 2020.
- [7] Heribert Wankerl, Maike L. Stern, Ali Mahdavi, Christoph Eichler, and Elmar W. Lang. Parameterized reinforcement learning for optical system optimization, 2020.
- [8] Sunae So, Trevon Badloe, Jaebum Noh, Jorge Bravo-Abad, and Junsuk Rho. Deep learning enabled inverse design in nanophotonics. *Nanophotonics*, 9(5):1041–1057, 2020.
- [9] John Peurifoy, Yichen Shen, Li Jing, Yi Yang, Fidel Cano-Renteria, Brendan Delacy, Max Tegmark, John D. Joannopoulos, and Marin Soljacic. Nanophotonic particle simulation and inverse design using artificial neural networks, 2017.
- [10] Christopher Yeung, David Ho, Benjamin Pham, Katherine T. Fountaine, Zihan Zhang, Kara Levy, and Aaswath P. Raman. Enhancing adjoint optimization-based photonic inverse design with explainable machine learning. *ACS Photonics*, 9(5):1577–1585, 2022.
- [11] Eamonn Keogh and Abdullah Mueen. *Curse of Dimensionality*, pages 314–315. Springer US, Boston, MA, 2017.
- [12] Yurii Nesterov. *Lectures on Convex Optimization*. Springer Publishing Company, Incorporated, 2nd edition, 2018.
- [13] Abeba Birhane, Atoosa Kasirzadeh, David Leslie, and Sandra Wachter. Science in the age of large language models. *Nature Reviews Physics*, 5(5):277–280, 2023.
- [14] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization, 2017.
- [15] Roger Fletcher. *Practical Methods of Optimization*. John Wiley & Sons, New York, NY, USA, second edition, 1987.
- [16] Sean Molesky, Zin Lin, Alexander Y. Piggott, Weiliang Jin, Jelena Vucković, and Alejandro W. Rodriguez. Inverse design in nanophotonics. *Nature Photonics*, 12(11):659–670, Oct 2018.
- [17] Christopher M. Lalau-Keraly, Samarth Bhargava, Owen D. Miller, and Eli Yablonovitch. Adjoint shape optimization applied to electromagnetic design. *Opt. Express*, 21(18):21693–21701, Sep 2013.
- [18] Paul Hansen and Lambertus Hesselink. Accurate adjoint design sensitivities for nano metal optics. *Opt. Express*, 23(18):23899–23923, Sep 2015.
- [19] Yannick Augenstein and Carsten Rockstuhl. Inverse design of nanophotonic devices with structural integrity. *ACS Photonics*, 7(8):2190–2196, 2020.
- [20] J.S. Jensen and O. Sigmund. Topology optimization for nano-photonics. *Laser & Photonics Reviews*, 5(2):308–321, 2011.
- [21] Steven G. Johnson. Notes on adjoint methods for 18.335, April 2021.
- [22] Pontriagin L. S. *The Mathematical Theory of Optimal Processes*. CRC Press, second edition, 1987.
- [23] Nicolas Lebbe. *Contribution in topological optimization and application to nanophotonics*. Theses, Université Grenoble Alpes, November 2019.
- [24] Stephane Durand, Ivan Cimrák, and Peter Sergeant. Adjoint variable method for time-harmonic maxwell equations. *COMPEL International Journal of Computations and Mathematics in Electrical*, 28:1202–1215, 09 2009.
- [25] Christopher M. Lalau-Keraly, Samarth Bhargava, Owen D. Miller, and Eli Yablonovitch. Adjoint shape optimization applied to electromagnetic design. *Opt. Express*, 21(18):21693–21701, Sep 2013.
- [26] Atılım Günes Baydin, Barak A. Pearlmutter, Alexey Andreyevich Radul, and Jeffrey Mark Siskind. Automatic differentiation in machine learning: A survey. *J. Mach. Learn. Res.*, 18(1):5595–5637, jan 2017.
- [27] Logan Su, Dries Vercautse, Jinhie Skarda, Neil V. Sapro, Jan A. Petykiewicz, and Jelena Vuckovic. Nanophotonic inverse design with spins: Software architecture and practical considerations, 2019.

- [28] Benjamin Vial. Nannos. <https://nannos.gitlab.io/license.html>, 2023.
- [29] Momchil Minkov, Ian AD Williamson, Lucio C Andreani, Dario Gerace, Beicheng Lou, Alex Y Song, Tyler W Hughes, and Shanhui Fan. Inverse design of photonic crystals through automatic differentiation. *ACS Photonics*, 7(7):1729–1741, 2020.
- [30] Tyler W Hughes, Ian AD Williamson, Momchil Minkov, and Shanhui Fan. Forward-mode differentiation of maxwell’s equations. *ACS Photonics*, 6(11):3010–3016, 2019.
- [31] Ansys Canada Ltd. Lumerical.
- [32] Comsol. COMSOL Multiphysics 6.
- [33] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. In *NIPS-W*, 2017.
- [34] Kerry J. Vahala. Optical microcavities. *Nature*, 424(6950):839–846, August 2003.
- [35] James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao Zhang. JAX: composable transformations of Python+NumPy programs. 2018.
- [36] Dougal Maclaurin, David Duvenaud, and Ryan P Adams. Autograd: Effortless gradients in numpy. In *ICML 2015 AutoML Workshop*, volume 238, page 5, 2015.
- [37] Frames White and Miles Lubin. juliadiff. <https://juliadiff.org/>, 2023.
- [38] Allan Leal. autodiff. <https://autodiff.github.io/>, 2023.
- [39] Rasmus E. Christiansen and Ole Sigmund. Inverse design in photonics by topology optimization: tutorial. *J. Opt. Soc. Am. B*, 38(2):496–509, Feb 2021.
- [40] Rasmus E. Christiansen and Ole Sigmund. Inverse design in photonics by topology optimization: tutorial. *J. Opt. Soc. Am. B*, 38(2):496–509, Feb 2021.
- [41] Yaguang Wang, Zhan Kang, and Pai Liu. Velocity field level-set method for topological shape optimization using freely distributed design variables. *International Journal for Numerical Methods in Engineering*, 120(13):1411–1427, 2019.
- [42] Michel Delfour and J.-P Zolésio. Velocity method and lagrangian formulation for the computation of the shape hessian. *Siam Journal on Control and Optimization - SIAM J CONTR OPTIMIZAT*, 29, 11 1991.
- [43] John Schneider. Understanding the finite-difference time-domain method. 2011.
- [44] Walter Greiner and Joachim Reinhardt. *Field Quantization*. Springer-Verlag, Berlin–Heidelberg–New York, 1996. With a foreword by D. A. Bromley.
- [45] Eberhard Engel and Reiner M. Dreizler. *Density Functional Theory*. Springer Berlin Heidelberg, 2011.
- [46] Michael. C. Delfour and Jean-Paul Zolésio. *Shapes and Geometries*. Society for Industrial and Applied Mathematics, second edition, 2011.
- [47] Jesse Lu, Stephen Boyd, and Jelena Vučković. Inverse design of a three-dimensional nanophotonic resonator. *Opt. Express*, 19(11):10563–10570, May 2011.
- [48] Jesse Lu and Jelena Vučković. Inverse design of nanophotonic structures using complementary convex optimization. *Opt. Express*, 18(4):3793–3804, Feb 2010.
- [49] Sandro Mignuzzi, Stefano Vezzoli, Simon A. R. Horsley, William L. Barnes, Stefan A. Maier, and Riccardo Sapienza. Nanoscale design of the local density of optical states. *Nano Letters*, 19(3):1613–1617, 2019. PMID: 30786717.
- [50] Fengwen Wang and Ole Sigmund. Optimization of photonic crystal cavities. In *2017 International Conference on Numerical Simulation of Optoelectronic Devices (NUSOD)*, pages 39–40, 2017.
- [51] Omer Yesilyurt, Zhaxylyk A. Kudyshev, Alexandra Boltasseva, Vladimir M. Shalaev, and Alexander V. Kildishev. Efficient topology-optimized couplers for on-chip single-photon sources. *ACS Photonics*, 8(10):3061–3068, 2021.
- [52] Panpan Li, Hongjian Li, Yifan Yao, Norleakvisoth Lim, Matthew Wong, Mike Iza, Michael J. Gordon, James S. Speck, Shuji Nakamura, and Steven P. DenBaars. Significant quantum efficiency enhancement of ingan red micro-light-emitting diodes with a peak external quantum efficiency of up to 6%. *ACS Photonics*, 0(0):null, 2023.
- [53] Panpan Li, Aurelien David, Hongjian Li, Haojun Zhang, Cheyenne Lynsky, Yunxuan Yang, Mike Iza, James S. Speck, Shuji Nakamura, and Steven P. DenBaars. High-temperature electroluminescence properties of InGaN red 40x40 μm^2 micro-light-emitting diodes with a peak external quantum efficiency of 3.2%. *Applied Physics Letters*, 119(23), 12 2021. 231101.

- [54] Tetsuya Taki and Martin Strassburg. Visible leds: More than efficient light. *ECS Journal of Solid State Science and Technology*, 9(1):015017, nov 2019.
- [55] Haejun Chung. Computational upper-limit of directional light emission in nano-led via inverse design. *Opt. Express*, 30(6):9008–9020, Mar 2022.

5 Appendix

5.1 Support Points behavior

The representation of the shape with support points is quite important for the gradient computation. The shape of the discretized geometries are deformed by dragging the support points following the gradients. First of all, the gradient in shape optimization is only meaningful in direction of the normal of the boundary. The gradient direction is therefore a weighted combination of the normal vectors of the two adjacent edges to the support points. Furthermore, depending on the density of the support points, one can not simply evaluate the sensitivity field on the support points position x_{n_i} directly since the sensitivity field can vary between support points. Moving an individual support point leads to moving the edge between the previous support point $x_{n_{i-1}}$ and the next support point $x_{n_{i+1}}$ see Figure 4. The influence of a small displacement of an individual support point n_i on the variation of the loss functional can be computed by [17]

$$\delta n_i = \int_{x_{n_{i-1}}}^{x_{n_i}} dx \frac{x - x_{n_{i-1}}}{x_{n_i}} \hat{n} V_{\Omega}(x) - \int_{x_{n_i}}^{x_{n_{i+1}}} dx \frac{x - x_{n_i}}{x_{n_{i+1}}} \hat{n} V_{\Omega}(x). \quad (6)$$




Figure 4: Edge movement for a displacement of an individual support point. The edge moves much further close to the support point n_i than close to the adjacent support points. For a small displacement of support point n_i , the variation of the support point is given by Equation 6.

5.2 Pseudo-code for the differentiable simulation function

The forward and backward methods seamlessly integrate into the computational graph of the chosen autograd framework where using forward and backward is often the naming convention. The inputs to the forward are necessary to specify the boundary conditions for the numerical solution. For Lumerical, this would mean activating the forward sources for the forward calculation and activating the adjoint sources for the backward calculation. After the simulation, the numerical solution from the simulation is collected and passed to the next function in the computational graph. After calling backward on the loss value and backpropagating up to the *differentiable simulation*, the adjoint simulation is initialized with the adjoint sources and the given shape. Note that the time reversal operation which is performed on the gradient input and the adjoint solution. The forward and adjoint solutions are then used to compute the sensitivity field [23]. c_1 and c_2 are constants, which depend on the permeability of the materials inside and outside of the boundary. Finally, the sensitivity field is evaluated on the boundary and projected on the boundary normal. These values then represent the gradients on the shape which are returned from to the computational graph.

Algorithm 1 Pseudocode for a *differentiable simulation* autograd function for automatic differentiation to derive shape gradients via the adjoint method.

```

1: class DIFFERENTIABLESIMULATION(AUTOGRAD)
2:
3:   function FORWARD(self, simulation, forward_source, c1, c2, shape)
4:     self.store  $\leftarrow$  simulation
5:     self.store  $\leftarrow$  c1, c2
6:     self.store  $\leftarrow$  shape
7:     simulation  $\leftarrow$  forward_Source
8:     simulation  $\leftarrow$  shape
9:     u  $\leftarrow$  simulation.run()
10:    self.store  $\leftarrow$  u
11:    return u
12:
13:   function BACKWARD(self, grad_input)
14:     simulation  $\leftarrow$  self.store
15:     shape  $\leftarrow$  self.store
16:     adjoint_source  $\leftarrow$  time_reverse(grad_input)
17:     simulation  $\leftarrow$  adjoint_source
18:     simulation  $\leftarrow$  shape
19:     v  $\leftarrow$  simulation.run()
20:     u  $\leftarrow$  self.store
21:     c1, c2  $\leftarrow$  self.store
22:     sensitivity_field  $\leftarrow$  u, v, c1, c2
23:     V_Omega  $\leftarrow$  sensitivity_field(shape)
24:     normals  $\leftarrow$  shape
25:     gradients  $\leftarrow$  normals  $\cdot$  V_Omega
26:     return gradients
27: end class

```
