Optimizing Photonic Structures with Large Language Model Driven Algorithm Discovery

Haoran Yin
h.yin@liacs.leidenuniv.nl
LIACS, Leiden University
Leiden, Netherlands

Anna V. Kononova
a.kononova@liacs.leidenuniv.nl
LIACS, Leiden University
Leiden, Netherlands

Thomas Bäck
t.h.w.baeck@liacs.leidenuniv.nl
LIACS, Leiden University
Leiden, Netherlands

Niki van Stein
n.van.stein@liacs.leidenuniv.nl
LIACS, Leiden University
Leiden, Netherlands

Abstract

We study how large language models (LLMs) can be used in combination with evolutionary computation techniques to automatically discover optimization algorithms for the design of photonic structures. Building on the Large Language Model Evolutionary Algorithm (LLaMEA) framework, we introduce structured prompt engineering tailored to multilayer photonic problems such as Bragg mirror, ellipsometry inverse analysis and solar cell antireflection coatings. We systematically explore multiple configurations of Evolution Strategies, including (1+1), (1+5), (2+10) and others, to balance exploration and exploitation. Our experiments show that LLM-generated algorithms, generated using small-scale problem instances, can match or surpass established methods like Quasi-oppositional Differential Evolution on large-scale realistic real-world problem instances. Notably, LLaMEA's self-debugging mutation loop, augmented by automatically extracted problem-specific insights, achieves strong anytime performance and reliable convergence across diverse problem scales. This work demonstrates the feasibility of domain-focused LLM prompts and evolutionary approaches in solving optical design tasks, paving the way for rapid, automated photonic inverse design. Such an approach is broadly applicable and can transform design processes across a wide range of scientific and engineering domains.

CCS Concepts

• **Theory of computation → Design and analysis of algorithms**; **Optimization with randomized search heuristics**; *Continuous optimization*; *Evolutionary algorithms*; • **Computing methodologies → Heuristic function construction**; • **Applied computing** → *Physics*.

Keywords

Large Language Models, Automated Algorithm Design, Photonic Structures, Evolution Strategies, Inverse Design, Black-Box Optimization, Domain-Specific Prompting, Photonics Benchmarking

GECCO '25 Companion, July 14–18, 2025, Malaga, Spain

ACM ISBN 979-8-4007-1464-1/2025/07
https://doi.org/10.1145/3712255.3734288

ACM Reference Format:
Haoran Yin, Anna V. Kononova, Thomas Bäck, and Niki van Stein. 2025. Optimizing Photonic Structures with Large Language Model Driven Algorithm Discovery. In *Genetic and Evolutionary Computation Conference (GECCO '25Companion), July 14–18, 2025, Malaga, Spain.* ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3712255.3734288

1 Introduction

Optimization of photonic structures plays a key role in the advancement of technologies in various fields such as telecommunication, solar energy and materials science [1, 17, 22, 23, 39]. However, the complexity and high dimensionality of these optimization problems pose significant challenges. Recent developments in automatic algorithm design, especially those that utilize large language models (LLMs), offer promising solutions.

LLMs have become powerful tools in the field of algorithm discovery and optimization. Several studies have demonstrated the ability of LLMs to automatically generate and refine optimization algorithms through an iterative process [20, 35]. By combining the power of large-scale language models with automatic algorithm design, these approaches open new avenues for developing algorithms for complex problems without relevant application expertise.

In this work, we extend the capabilities of the Large Language Model Evolutionary Algorithm (LLaMEA) framework [35] to deal with real-world photonic problems by addressing two critical limitations:

- **Generic Task Prompts**: Original LLaMEA prompts lacked domain-specific guidance, which can lead to suboptimal algorithm designs.
- **Limited Evolution Strategy Diversity**: Previous studies focused only on simple (1,1) and (1+1) strategies, neglecting the potential of population-based exploration.

To overcome these limitations, we introduce structured task prompts enriched with photonics-specific problem descriptions and algorithmic insights. Furthermore, we systematically evaluate five new evolution strategy configurations-(1,5), (1+5), (2,10), (2+10) and (5+5)-to balance exploration-exploitation trade-offs.

The paper is organized as follows. Section 2 reviews the related work on LLM-driven algorithm discovery and photonic structure optimization problems. Section 3 describes the methodology in detail. Section 4 provides data related to the experimental setup.

Section 5 shows the experimental results. Section 6 summarizes the results of the work.

2 Related Work

The rapid development of LLMs and optimization techniques has stimulated a strong interest in combining LLM-driven algorithm discovery with domain-specific applications. This section provides an overview of related work in two key areas: algorithmic automatic generation tools using LLMs and photonic structure optimization problems.

2.1 Large Language Models for Algorithm Discovery

Recent research has explored how LLMs can contribute to algorithmic discovery [21]. Two state-of-the-art representative frameworks in this area are Evolution of Heuristic (EoH) and LLaMEA [20, 35]. Both frameworks illustrate the potential of LLMs to automate and accelerate algorithmic discovery.

EoH combines LLMs with evolutionary computation to iteratively generate and refine heuristics. These heuristics are small functions that are then inserted into a larger code template before the evaluation of a problem. It evolves both natural-language descriptions of heuristics and executable code representations of these heuristics, allowing for diverse exploration and efficient improvement. EoH demonstrates state-of-the-art performance in heuristic automated design on combinatorial optimization problems, outperforming existing methods such as FunSearch on bin-packing and traveling salesperson problems [20].

LLaMEA also uses LLMs within an evolution algorithm framework, focusing mainly on the (1,1) and (1+1) evolution strategies, to automatically generate, mutate and optimize complete metaheuristics [30, 35]. Unlike EoH's emphasis on 'thought', LLaMEA directly optimizes the algorithm itself, using run-time performance as feedback for iterative improvement. In benchmarking continuous problems, it produces algorithms that are comparable to or better than state-of-the-art optimization methods such as the covariance matrix adaptation evolution strategy (CMA-ES) and differential evolution (DE) [35]. Furthermore, the recently proposed LLaMEA-HPO framework augments LLaMEA with automated hyperparameter tuning, offloading parameter optimization from the LLM and thereby boosting algorithm quality while reducing the overall LLM query cost [36].

While both EoH and LLaMEA are open-source and easy to extend, we chose to base this research on LLaMEA due to it's superior performance for discovering algorithms in the black-box optimization domain and for its ability to generate and optimize large code-bases.

2.2 Earlier Methods for Automated Algorithm Discovery

Beyond EoH and LLaMEA, a variety of other frameworks explore how LLMs can automate algorithm discovery. One of the earliest works is *FunSearch* [31], which applies an LLM within a distributed evolutionary search over function spaces. Starting from seed functions, FunSearch evolves them by prompting the LLM for refined or alternative code.

Algorithm Evolution using Large Language Model (AEL) [19] and *ReEvo* [37], similarly integrate LLMs within an iterative loop, generating small heuristic code snippets and improving them via mutation and crossover prompts. While EoH focuses on heuristics in text and code form, AEL/ReEvo explore a closely related direction by evolving coded modules. Both frameworks have shown promising results on combinatorial benchmarks such as the Traveling Salesperson Problem.

2.3 Optimization of Photonic Structures

Optimization of photonic structures has been an important area of research, as the need for high efficiency and high precision drives the optimization of photonic structures in various applications such as communication, semiconductors, LED displays, materials analysis and solar cells [1, 17, 22, 23, 39]. The following are three real-world problems related to global optimization of multilayered photonic structures, all of which have vital applications.

Bragg Mirror. A Bragg reflector mirror, or Bragg mirror, is a series of two or more semiconductors or dielectric materials stacked in a staggered pattern to achieve high reflectivity in a particular optical band [4]. It has a wide range of applications, such as the construction of acoustic wave reflectors and filters [28], the analysis of the crystal structure of materials [26], the improvement of solar cell efficiency [16], the detection of changes in physical quantities such as temperature and pressure [10] and the enhancement of the interaction between light and matter in quantum computation and quantum information [24].

Ellipsometry Inverse Problem. Ellipsometry is a nondestructive optical measurement technique that can be used to measure the thickness of the thin film, the refractive index and the absorption coefficients [32]. In chip fabrication, the properties of thin films have a critical impact on circuit performance [40]. The solution of the inverse problem of ellipsometry can help to improve the accuracy of the production process [34]. It can also be used in the new energy and chemical industry to study the properties of solar cell materials, nanostructured thin films and chemical coatings [8]. Ellipsometry inverse problem solving can not only improve the data analysis method in materials science, but also improve the applicability of thin film technology in industrial production [9, 14, 33].

Photovoltaic Problem. This refers to the design of a sophisticated antireflection coating for solar cells - a multilayer thin film structure optimized to minimize optical reflections and maximize the absorption of the active layer of the cell. Such coatings typically consist of alternating dielectric layers with different refractive indices. By suppressing surface reflection losses, a well-designed antireflective coating can significantly improve the power conversion efficiency of a solar cell (bare silicon surfaces, for instance, reflect 30% of incident sunlight) [15]. The difficulty lies in achieving broad-band, all-encompassing antireflection: single-layer coatings can only eliminate reflections in narrow bands, so multilayer or graded index designs are needed to reduce reflectivity across the entire broad solar spectrum. Optimizing such coatings is a complex photonic inverse design problem, with many local minima in performance due to wave interference effects at multiple wavelengths [3].

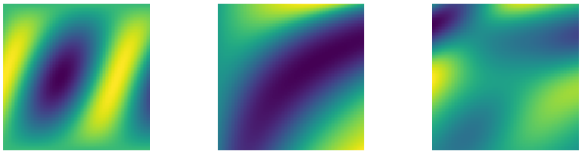

(a) Bragg mirror **(b) ellipsometry** **(c) photovoltaics**

Figure 1: Landscape of photonic structure optimizing problems in 2D. The darker the color, the more fit the structure.

3 Optimizing Photonic Structures with LLaMEA

We apply LLaMEA to real-world problems published by Bennet et al. to automatically discover and implement photonic structure optimization algorithms [2, 35].

3.1 Real-world Problem Simulation

Python-based Multilayer Optics Optimization and Simulation Hub (PyMoosh) is a numerical toolkit for calculating the optical properties of multilayer structures [18]. PyMoosh is especially suited for complex problems involving multilayer optical structures (e.g. structural analysis of metallic or metalized layers) and can be extended to support advanced optimization and inverse problem design.

Building on PyMoosh, Bennet et al. have developed a testbed for defining and solving photonic optimization problems [2]. Based on their work, we migrate the following problems to the IOHexperimenter platform [5]:

(1) optimization of a Bragg mirror,
(2) solving of an ellipsometry inverse problem,
(3) design of a sophisticated anti-reflection coating to optimize solar absorption in a photovoltaic solar cell.

Bennet et al. have completed benchmark work on common optimization algorithms on these three photonic structure optimization problems for reference [3]. The landscapes of these problems in 2D are shown in Fig. 1. This figure is intended to only visualize the problem, since the dimensionality of the problem is directly related to the number of layers of the photonic structure and, therefore, the 2D form of some problems is not of any practical interest.

3.2 Enhanced Task Prompt Design

To improve LLaMEA's ability to address domain-specific problems, we have added two key sections to the original task prompt:

- Problem Description: a concise summary of the photonic structure optimization task, including key parameters (e.g., range of layer thicknesses and dielectric constant constraints) and physical objectives (e.g., reflectivity maximization).
- Algorithmic Insight: domain knowledge guidance, such as 'encourage algorithms to detect and preserve modular structures' or 'encourage periodicity in solutions through customized cost functions or constraints'.

These structured suggestions could lead to more specialized algorithms that perform better for photonics-specific requirements.

3.3 Automatic Generation of Algorithms

We chose LLaMEA as the algorithm discovery method in our experiments for the following reasons:

(1) Targeting continuous optimization problems: Our problems are continuous optimization problems, and LLaMEA provides a solution specifically for such problems. It uses automated generation and iterative optimization based on metaheuristic algorithms suitable for dealing with the optimization of continuous variables.
(2) Reliable performance: LLaMEA has been shown to perform better than EoH in several complex optimization tasks [35].

In our work, we perform *rapid validation and benchmarking* of algorithms generated on small-scale problem instances and apply them to more complex problem instances. Feedback such as convergence speed and anytime performance metrics are used as feedback for LLaMEA to iteratively improve algorithms. After many improvements, LLaMEA will provide a number of usable algorithms. The algorithms validated as the best will participate in benchmarking and comparison with other commonly used state-of-the-art optimization algorithms.

4 Experimental Setup

The experiments are divided into two main parts: **discovery** and **benchmarking**. That is, LLaMEA is used to search for optimization algorithms, and the algorithms discovered with the best performance are benchmarked against other commonly used optimization algorithms on different instances of the problem. The following content details the experimental setup in terms of problem setup, performance metric, prompt setup for LLM and benchmarking.

4.1 Problem Setup

4.1.1 Bragg Mirror. We set the optimization objective to find a photonic structure that has the strongest reflectance for light with a wavelength of 600 nm. The permittivity of the first material is 1.96, the permittivity of the other material is 3.24 and the maximum thickness of each layer is 218 nm. The problem instance with 10 layers of alternating materials will collectively be referred to as *mini-Bragg* and those with 20 layers of alternating materials will collectively be referred to as *Bragg*.

4.1.2 Ellipsometry Inverse Problem. The substrate material is gold and parameters of the reflective material to be solved include thickness and permittivity. The minimum thickness is 50 nm, the maximum thickness is 150 nm, the minimum permittivity is 1.1 and the maximum permittivity is 3.0. The number of layers of material is 1. The problem instance in the latter collectively is referred to as *ellipsometry*.

4.1.3 Photovoltaic Problem. We optimize the photonic structure to improve solar absorption, with the thickness of each material layer ranging from 30 to 250 nm. The permittivity is 2.0 for the first material and 3.0 for the other. The substrate material parameters are set as default. Similarly, we call an instance with 10 layers *photovoltaic*, 20 layers *big-photovoltaic* and 32 layers *huge-photovoltaic*.

All problem instances are set as **minimization** problems. From Fig. 2, according to Exploratory Landscape Analysis (ELA) [25],

Table 1: Parameter settings for different problem instances. Different columns of thickness and permittivity correspond to different materials. The algorithm discovery process for different problems is based on smaller instances marked with ♠.

Instances	*mini-Bragg* ♠		*Bragg*		*ellipsometry* ♠	*photovoltaic* ♠		*big-photovoltaic*		*huge-photovoltaic*	
layers	10		20		1	10		20		32	
materials	2		2		1	2		2		2	
min thickness(nm)	0	0	0	0	50	30	30	30	30	30	30
max thickness(nm)	218	218	218	218	150	250	250	250	250	250	250
min permittivity max permittivity	1.96	3.24	1.96	3.24	1.1 3.0	2.0	3.0	2.0	3.0	2.0	3.0
evaluation budget	10000		20000		1000	5000		10000		16000	

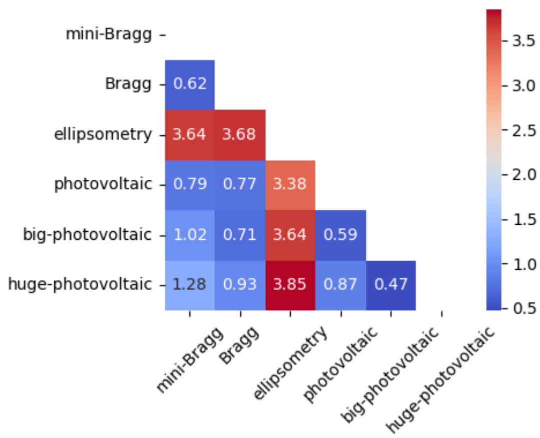

Figure 2: Average Hellinger distance of ELA featuers of different instances [13]. Each ELA feature is calculated 100 times to form a distribution for the calculation of Hellinger distance. The smaller the value, the more similar the instances are. It is indeed more similar between instances of the same problem.

different instances of the same problem are very similar. Thus, algorithm discovery processes are based **only on smaller instances**, while the subsequent benchmarks on the best algorithms are based on **all instances** to save time. Table 1 shows the evaluation budget and summarizes the parameter settings described above.

4.2 Performance Metric

We use the metric suggested by LLaMEA, AOCC, as a feedback to the LLM. Eq. 1 gives the definition of AOCC:

$$AOCC(y_{a,f}) = \frac{1}{B}\sum_{i=1}^{B}\left(1-\frac{\min(\max((y_i),lb),ub)-lb}{ub-lb}\right) \quad (1)$$

where $y_{a,f}$ is a series of log-scaled current-best fitness value of f during the running of optimization algorithm a, B is the evaluation budget, y_i is the i-th element of $y_{a,f}$, ub is the upper bound of f and lb is the lower bound of f. In addition to the AOCC used by LLaMEA, we also add the optimal fitness value found at the end of the algorithm runs, y^*, to the feedback.

To help distinguish problems on *mini-Bragg*, *ellipsometry* and *photovoltaic*, which are used for algorithm discovery, in tables and figures, these problems are marked with ♠. ub is set to 1.0 for *mini-Bragg* and *photovoltaic* and 40.0 for *ellipsometry*.

4.3 Prompt Setup

The high degree of freedom in possible prompt customization is a feature of LLaMEA, which uses the task prompt, the mutation prompt, the feedback prompt and the output prompt–the first three of which are modified and fed to gpt-4o-2024-08-06 with temperature set to 0.8 in our experiments.

Firstly, we state in the task prompt that the goal of the task is to find algorithms suitable for the optimization of multilayer photonic structures:

Task Prompt

The optimization algorithm should be able to find high-performance solutions to a wide range of tasks, which include evaluation on real-world applications such as, e.g., optimization of multilayered photonic structures.
<problem description>
<algorithmic insight>
Your task is to write the optimization algorithm in Python code. The code should contain an '__init__(self, budget, dim)' function and the function 'def __call__(self, func)', which should optimize the black box function 'func' using 'self.budget' function evaluations.
The func() can only be called as many times as the budget allows, not more. Each of the optimization functions has a search space between func.bounds.lb (lower bound) and func.bounds.ub (upper bound). The dimensionality can be varied.
Give an excellent and novel heuristic algorithm to solve this task and also give it a one-line description with the main idea.

where **problem description** and **algorithmic insight** are automatically generated via the following meta-prompts to gpt-4o-2024-08-06:

Meta-prompt for generating problem description and algorithmic insight

problem description: Please read *Illustrated tutorial on global optimization in nanophotonics* first [3]. Then, give me summaries of bragg mirror problem, ellipsometry problem and photovoltaics problem, respectively. The generated summaries need to be used in other chatgpt chats, so make sure they are understood by chatgpt.
algorithmic insight: Please give me algorithmic insights of these three problems respectively. Similarly, these insights are used in other chatgpt chats, prompting chatgpt to generate optimization algorithms applicable to these problems, so make sure that these insights are understood by chatgpt.

The resulting generated problem description for the Bragg mirror is as follows, while all descriptions and insights can be found in the supplementary material:

Problem description for Bragg mirror

The Bragg mirror optimization aims to maximize reflectivity at a wavelength of 600 nm using a multilayer structure with alternating refractive indices (1.4 and 1.8). The structure's thicknesses are varied to find the configuration with the highest reflectivity. The problem involves two cases: one with 10 layers (minibragg) and another with 20 layers (bragg), with the latter representing a more complex inverse design problem. The known optimal solution is a periodic Bragg mirror, which achieves the best reflectivity by leveraging constructive interference. This case exemplifies challenges such as multiple local minima in the optimization landscape.

For each problem, we experiment with (a) the task prompt without problem description and algorithmic insight, (b) with problem description and (c) with both problem description and algorithmic insight. For each task prompt setting, LLaMEA runs 5 times with an evolution strategy (1 + 1), generating 100 algorithms per run. Each algorithm is executed 3 times to eliminate experimental randomness.

Secondly, the dynamic mutation controlling prompt is set to control the LLaMEA mutation process, as this prompt has been validated to have better mutation control when combined with gpt-4o [38] with the fast mutation operator [6]:

Dynamic Mutation Controlling Prompt

Refine the strategy of the selected solution to improve it. Make sure that you only change $x\%$ of the code, which means if the code has 100 lines, you can only change $\lfloor x \rfloor$ lines, and the rest lines should remain the same. For this code, it has n lines, so you can only change $\min(\lfloor n \times x/100 \rfloor, 1)$ lines, the rest $n - \min(\lfloor n \times x/100 \rfloor, 1)$ lines should remain the same. This changing rate $x\%$ is the mandatory requirement, you cannot change more or less than this rate.

The mutation rate x in the dynamic mutation control prompt is sampled from the heavy-tailed distribution known as the fast mutation operator [6].

Finally, in the feedback prompt, we provide both AOCC and y^* information to help LLM better establish the connection between algorithmic implementation and real-world results as a way to improve existing algorithms in a targeted manner:

Feedback Prompt

The algorithm <name> got an average Area over the convergence curve (AOCC, 1.0 is the best) score of <aocc_score> with standard deviation <aocc_std>. And the mean value of best solutions found was <y_best> (0. is the best) with standard deviation <y_best_std>.

4.4 Evolution Strategy Configuration

After confirming which task prompt setting works best, we select per problem kind the task prompt setting in the following experiments. In addition to the initial (1,1) and (1+1) configurations, which are the default for LLaMEA, we investigate five other configurations of evolution strategies:

- Comma strategy: (1,5), (2,10)
- Plus strategy: (1+5), (2+10), (5+5)

These configurations are often used in the black-box optimization literature, as they can balance exploration (by expanding the population of offspring) and exploitation (by retaining the elite), with λ denoting the number of offspring and μ denoting the number of parents and are indicated by (μ, λ) [30].

In these experiments, for each of the configurations, LLaMEA runs 5 times for each problem instance, during each run LLaMEA generates 100 algorithms. Each algorithm is executed 3 times to eliminate experimental randomness.

4.5 Benchmarking

After the algorithm discovery step, for each problem kind, we have 2500 generated algorithms. Of these, the 3 best algorithms in terms of AOCC average values participate in the final benchmarking against five algorithms commonly used in the field of photonic structure optimization, to provide a performance comparison. These algorithms include: DE [7], CMA-ES [11], Broyden–Fletcher–Goldfarb–Shanno (BFGS) [12], quasi-Newton differential evolution (QNDE) [27] and quasi-oppositional differential evolution (QODE) [29]. For each problem instance, each algorithm is executed 15 times. The optimal algorithms discovered through *mini-Bragg* and *photovoltaic* are also applied to higher-dimensional problem instances.

5 Results

This section presents the results of the experiments, for which the raw data and associated code are publicly available [1].

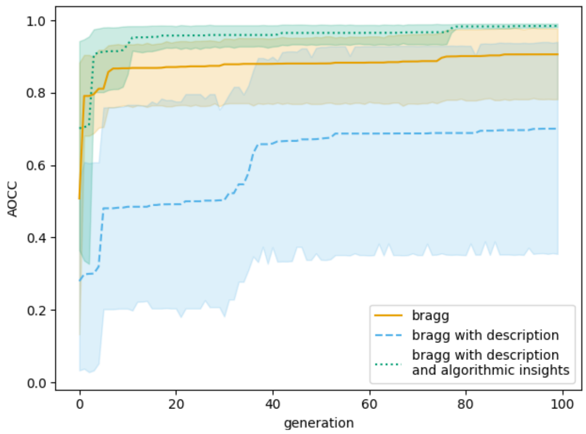

(a) *mini-Bragg* ♠

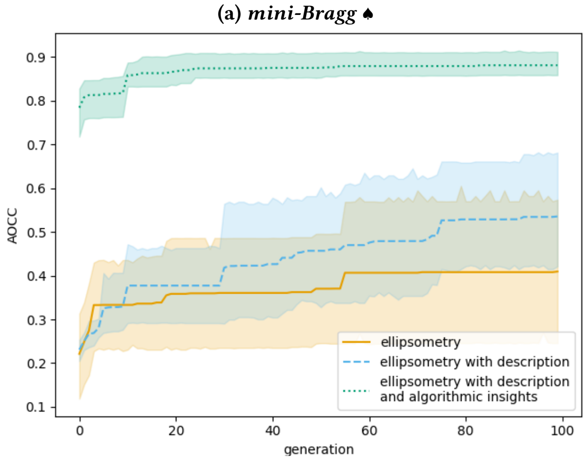

(b) *ellipsometry* ♠

Figure 3: Examples of task prompts that add problem descriptions and algorithm insights that improve the performance of LLaMEA. The higher the AOCC, the better. Prompts with descriptions and insights both outperform other prompt settings for *mini-Bragg* and *ellipsometry* instances.

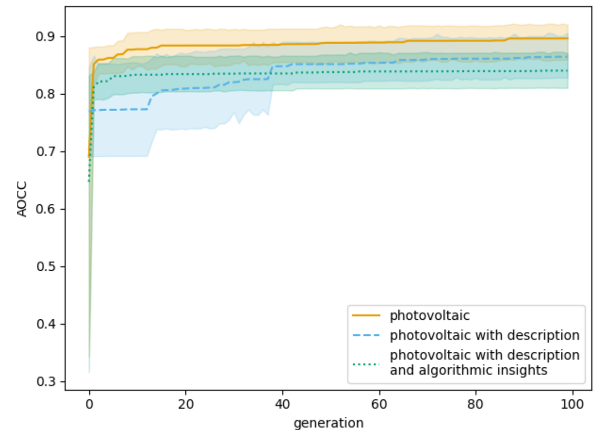

(a) *photovoltaic* ♠

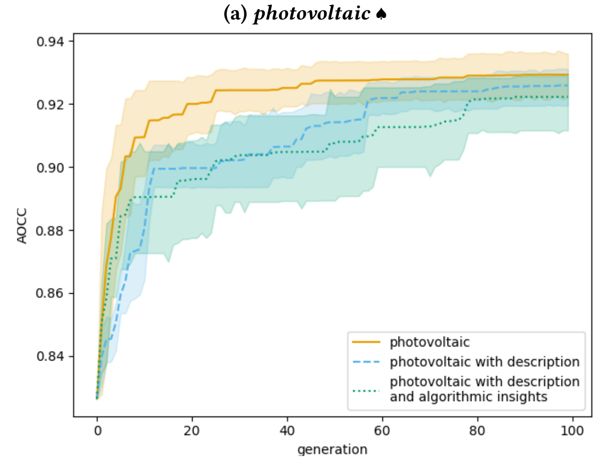

(b) *photovoltaic* ♠, all runs start with the same solution.

Figure 4: Example of a task prompt that adds a problem description and algorithm insight that does not improve the effectiveness of LLaMEA. The higher the AOCC, the better. For *photovoltaic* instance, task prompts without description and insight works best. Considering that the initial AOCC is high for all task prompt settings, Fig. 4b shows the results for each LLaMEA run starting from the same solution to avoid the impact of initial gaps.

5.1 Enhanced Task Prompt Design

The integration of domain-specific knowledge into structured prompts greatly influences the quality of the algorithms generated by LLaMEA. For the Bragg mirror designing and ellipsometry inverse problems, the detailed problem descriptions and algorithmic insights (see the supplementary material) guided LLM to prioritize algorithms that were consistent with physical principles. For example, prompts emphasizing 'periodic solutions' and 'constructive interference' motivate the algorithms to use modular initialization and symmetry constraints, which are essential for Bragg reflector optimization, as shown in Fig. 3a. These prompts explicitly encourage the use of hybrid global-local search methods, e.g. combining DE and BFGS for local refinement, which improves AOCC by 10% compared to general prompts. For ellipsometry, problem description and algorithmic insight provide more significant improvements.

However, in the photovoltaic problem, prompts without domain-specific insights outperformed the augmented variant, as shown in Fig. 4a. This counterintuitive result may be due to the the noisy fitness landscape of the photovoltaic task or gpt-4o lacks related knowledge about photovoltaic problem, where an overly detailed large prompt may prematurely limit exploration. Considering that all prompt settings in Fig. 4a start with a high AOCC, we let each LLaMEA run starting from the same solution to avoid the impact of initial gaps. However, the new and more reliable experimental results in Fig. 4b still show that a large prompt may pose a limitation.

5.2 Evolution Strategy Configuration

The impact of the choice of evolution strategy configuration on the performance of the algorithm varies according to the problem.

[1]https://doi.org/10.5281/zenodo.15073784

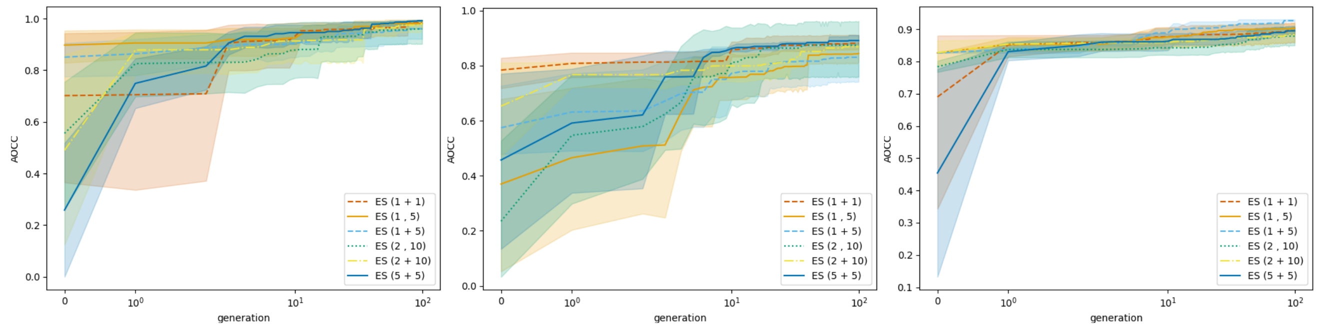

(a) ***mini-Bragg*** ♠ (b) ***ellipsometry*** ♠ (c) ***photovoltaic*** ♠

Figure 5: Impact of different ES configuration choices. The preference for ES strategies is different for each problem and most of the time the difference is not significant.

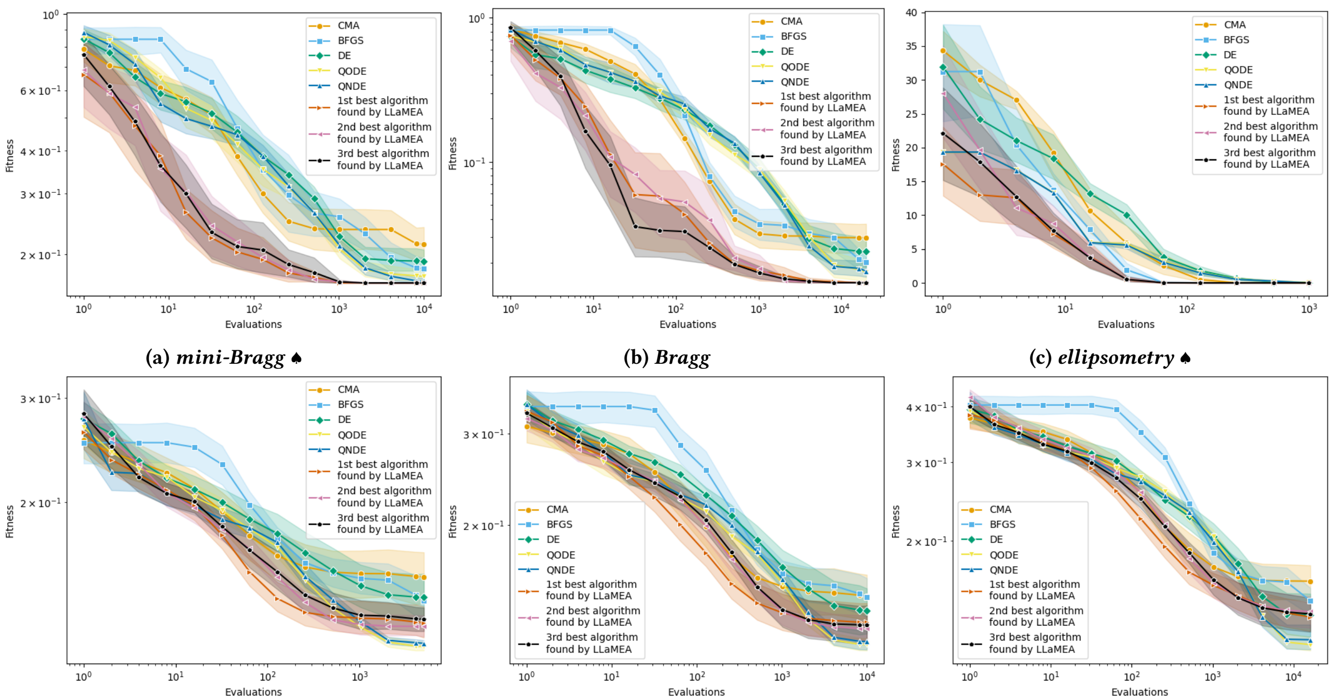

(a) ***mini-Bragg*** ♠ (b) ***Bragg*** (c) ***ellipsometry*** ♠

(d) ***photovoltaic*** ♠ (e) ***big-photovoltaic*** (f) ***huge-photovoltaic***

Figure 6: Convergency curves of best algorithms found by LLaMEA and baselines with different problem instances, averaged over 15 runs. y-axis represents fitness, the smaller, the better. x-axis represents evaluations of problem instances. Each subfigure represents a problem instance. Problems marked with ♠ have been used for both algorithm discovery and benchmarking, while other problems-only for benchmarking.

In the Bragg mirror, as well as the ellipsometry problem, the (5+5) configuration achieves the highest AOCC, as shown in Figs. 5a and 5b, taking full advantage of the population diversity to exploit periodic solutions while retaining elite individuals for local improvement. In contrast, the photovoltaic problem favored the (1+5) configuration, as shown in Fig. 5c. Furthermore, while there are differences between the performance of these configuration, they are not significant.

5.3 Algorithms Benchmarking

LLaMEA-based automatic algorithm discovery for each small instance of the three problem kinds has resulted in 3 best algorithms (different per problem kind). The performance of these 3 algorithms has been compared against the commonly used algorithms both on small and larger instances of each problem kinds, see Fig. 6. It should be noted that automatically discovered algorithms have been "transferred" to larger instances of the problem, in the sense

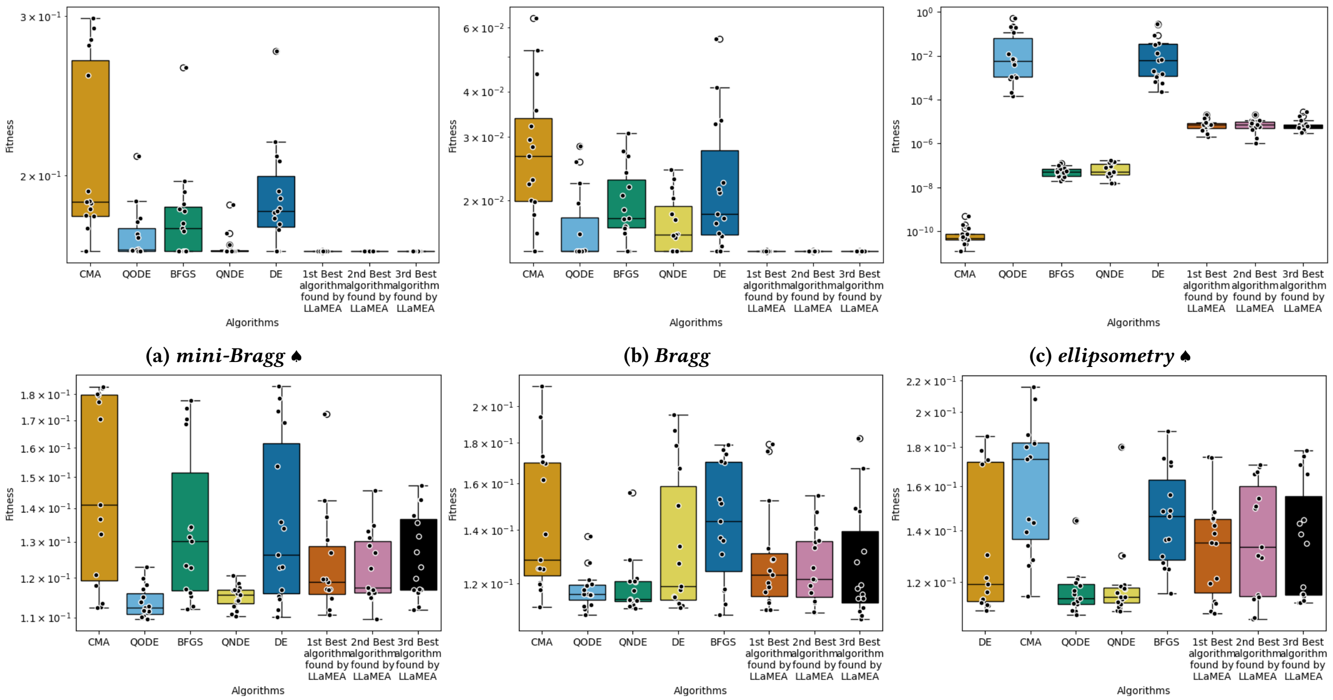

(a) ***mini-Bragg*** ♠ (b) ***Bragg*** (c) ***ellipsometry*** ♠

(d) ***photovoltaic*** ♠ (e) ***big-photovoltaic*** (f) ***huge-photovoltaic***

Figure 7: Distribution of best structure found by different algorithms. y**-axis represent fitness of photonic structure.** x**-axis represent different algorithms. Each subfigure represents a problem instance.**

that they have not been used for automatic algorithm discovery but only used for bechmarking.

From Figs. 6a and 6b, we find that the algorithms found by LLaMEA demonstrate a significantly more rapid convergence trend. However, the best algorithms found by LLaMEA for ellipsometry are all good at finding a local optimum at the initial stage, as shown in Fig. 6c. For the three instances of the photovoltaic problem, algorithms discovered by LLaMEA are not the best, but still show faster convergence.

Fig. 7 shows the distribution of the fitness value of the optimal solutions found by different algorithms for different problem instances after 15 runs. We find that the best algorithms found by LLaMEA perform very well and are basically comparable to the best human-designed algorithms. Especially for Bragg mirror problem instances, the found algorithms not only reliably find the optimal solutions but also the distribution is concentrated, which means that the performances are highly stable. For photovoltaic problem instances, the algorithms found are relatively less stable but still reliable.

6 Conclusions

This study demonstrates the potential of making real-world problem-related descriptions as well as algorithmic insights available to LLM and combining them with different evolution strategies for automatic algorithm discovery in photonic structure optimization. By embedding domain-specific knowledge into the structured task prompt, the algorithms generated by LLaMEA achieve advanced performance on real-world problems in comparison to other widely used successful algorithms in the area of photonic structure optimization. Key findings include:

- Embedding domain-specific photonics knowledge into the LLM prompt, combined with evolutionary search strategies, enables automated discovery of optimization algorithms that achieve high performance in photonic structure design.
- In complex photonic optimization tasks, LLM generates algorithms that can match or even surpass state-of-the-art methods (e.g., DE and CMA-ES) with performance comparable to human-designed heuristics.
- Systematic tests using a variety of evolution strategies (e.g., (1 + 1), (5 + 5) and (2 + 10)) show that while different strategies affect the performance of a given problem (e.g., the (5 + 5) strategy performs well in the Bragg reflector and ellipsometry tasks while the (1+5) strategy performs moderately well in the photovoltaic task), the overall difference in performance is not significant.

In summary, these findings confirm that LLM-driven automatic algorithm design can effectively automate the design of optimization algorithms for realistic photonic problems, paving the way for the fast and cost-effective design of photonic structures without the need for significant human intervention. Crucially, the methodology is broadly applicable and can transform algorithm design across a wide range of real-world domains.

References

[1] Sameh O Abdellatif and Gehad Ali Alsayed. 2019. Optimizing 1D photonic crystal structures for thin film solar cells. In *2019 IEEE Conference on Power Electronics and Renewable Energy (CPERE)*. IEEE, 333–337.

[2] Pauline Bennet, Emmanuel Centeno, Jérémy Rapin, Olivier Teytaud, and Antoine Moreau. 2020. The photonics and ARCoating testbeds in Nevergrad. (May 2020). https://hal.science/hal-02613161 working paper or preprint.

[3] Pauline Bennet, Denis Langevin, Chaymae Essoual, Abdourahman Khaireh-Walieh, Olivier Teytaud, Peter Wiecha, and Antoine Moreau. 2024. Illustrated tutorial on global optimization in nanophotonics. *JOSA B* 41, 2 (2024), A126–A145.

[4] WL Bragg, JJ Thomson, and Herren Friedrich. 1914. Diffraction of short electromagnetic waves. In *Proceedings of the Cambridge Philosophical Society: Mathematical and physical sciences*, Vol. 17. 43.

[5] Jacob de Nobel, Furong Ye, Diederick Vermetten, Hao Wang, Carola Doerr, and Thomas Bäck. 2021. IOHexperimenter: Benchmarking Platform for Iterative Optimization Heuristics. *arXiv e-prints:2111.04077* (nov 2021). arXiv:2111.04077 https://arxiv.org/abs/2111.04077

[6] Benjamin Doerr, Huu Phuoc Le, Régis Makhmara, and Ta Duy Nguyen. 2017. Fast genetic algorithms. In *Proceedings of the genetic and evolutionary computation conference*. 777–784.

[7] Vitaliy Feoktistov. 2006. *Differential evolution*. Springer.

[8] Hiroyuki Fujiwara and Robert W Collins. 2018. *Spectroscopic ellipsometry for photovoltaics*. Vol. 1. Springer.

[9] Débora Gonçalves and Eugene A Irene. 2002. Fundamentals and applications of spectroscopic ellipsometry. *Química Nova* 25 (2002), 794–800.

[10] Michal Gryga, Dalibor Ciprian, and Petr Hlubina. 2022. Distributed Bragg reflectors employed in sensors and filters based on cavity-mode spectral-domain resonances. *Sensors* 22, 10 (2022), 3627.

[11] Nikolaus Hansen, Sibylle D Müller, and Petros Koumoutsakos. 2003. Reducing the time complexity of the derandomized evolution strategy with covariance matrix adaptation (CMA-ES). *Evolutionary computation* 11, 1 (2003), 1–18.

[12] John D Head and Michael C Zerner. 1985. A Broyden—Fletcher—Goldfarb—Shanno optimization procedure for molecular geometries. *Chemical physics letters* 122, 3 (1985), 264–270.

[13] Ernst Hellinger. 1909. Neue begründung der theorie quadratischer formen von unendlichvielen veränderlichen. *Journal für die reine und angewandte Mathematik* 1909, 136 (1909), 210–271.

[14] Karsten Hinrichs and Klaus-Jochen Eichhorn. 2018. *Ellipsometry of functional organic surfaces and films*. Vol. 52. Springer.

[15] Chunxue Ji, Wen Liu, Yidi Bao, Xiaoling Chen, Guiqiang Yang, Bo Wei, Fuhua Yang, and Xiaodong Wang. 2022. Recent applications of antireflection coatings in solar cells. In *Photonics*, Vol. 9. MDPI, 906.

[16] Yajie Jiang, Mark J Keevers, and Martin A Green. 2018. Design of Bragg Reflector in GaInP/GaInAs/Ge Triple-Junction Solar Cells for Spectrum Splitting Applications. In *2018 IEEE 7th World Conference on Photovoltaic Energy Conversion (WCPEC)(A Joint Conference of 45th IEEE PVSC, 28th PVSEC & 34th EU PVSEC)*. IEEE, 0901–0904.

[17] Jungtaek Kim, Mingxuan Li, Oliver Hinder, and Paul Leu. 2024. Datasets and benchmarks for nanophotonic structure and parametric design simulations. *Advances in Neural Information Processing Systems* 36 (2024).

[18] Denis Langevin, Pauline Bennet, Abdourahman Khaireh-Walieh, Peter Wiecha, Olivier Teytaud, and Antoine Moreau. 2024. PyMoosh: a comprehensive numerical toolkit for computing the optical properties of multilayered structures. *JOSA B* 41, 2 (2024), A67–A78.

[19] Fei Liu, Xialiang Tong, Mingxuan Yuan, and Qingfu Zhang. 2023. Algorithm evolution using large language model. *arXiv preprint arXiv:2311.15249* (2023).

[20] Fei Liu, Tong Xialiang, Mingxuan Yuan, Xi Lin, Fu Luo, Zhenkun Wang, Zhichao Lu, and Qingfu Zhang. 2024. Evolution of Heuristics: Towards Efficient Automatic Algorithm Design Using Large Language Model. In *Forty-first International Conference on Machine Learning*.

[21] Fei Liu, Yiming Yao, Ping Guo, Zhiyuan Yang, Zhe Zhao, Xi Lin, Xialiang Tong, Mingxuan Yuan, Zhichao Lu, Zhenkun Wang, et al. 2024. A systematic survey on large language models for algorithm design. *arXiv preprint arXiv:2410.14716* (2024).

[22] Qian Liu, Eric Sandgren, Miles Barnhart, Rui Zhu, and Guoliang Huang. 2015. Photonic nanostructures design and optimization for solar cell application. In *Photonics*, Vol. 2. MDPI, 893–905.

[23] Dang Hoang Long, In-Kag Hwang, and Sang-Wan Ryu. 2009. Design optimization of photonic crystal structure for improved light extraction of GaN LED. *IEEE Journal of Selected Topics in Quantum Electronics* 15, 4 (2009), 1257–1263.

[24] Maurine Malak. 2014. Beyond interferometers based on silicon-air bragg reflectors: Toward on-chip optical microinstruments—A review. *IEEE Journal of Selected Topics in Quantum Electronics* 21, 4 (2014), 49–60.

[25] Olaf Mersmann, Bernd Bischl, Heike Trautmann, Mike Preuss, Claus Weihs, and Günter Rudolph. 2011. Exploratory landscape analysis. In *Proceedings of the 13th annual conference on Genetic and evolutionary computation*. 829–836.

[26] AP Mihai, M Adabi, W Liu, H Hill, N Klein, and PK Petrov. 2015. Metallic multilayers for X-band Bragg reflector applications. *arXiv preprint arXiv:1506.07702* (2015).

[27] Nasimul Noman and Hitoshi Iba. 2008. Accelerating differential evolution using an adaptive local search. *IEEE Transactions on evolutionary Computation* 12, 1 (2008), 107–125.

[28] Pratyasha Priyadarshini, Arnab Goswami, and Bijoy Krishna Das. 2024. Distributed Bragg reflector–based resonance filters for silicon photonics: design and demonstration. *Journal of Optical Microsystems* 4, 4 (2024), 041403–041403.

[29] Shahryar Rahnamayan, Hamid R Tizhoosh, and Magdy MA Salama. 2007. Quasi-oppositional differential evolution. In *2007 IEEE congress on evolutionary computation*. IEEE, 2229–2236.

[30] Ingo Rechenberg. 1978. Evolutionsstrategien. In *Simulationsmethoden in der Medizin und Biologie: Workshop, Hannover, 29. Sept.–1. Okt. 1977*. Springer, 83–114.

[31] Bernardino Romera-Paredes, Mohammadamin Barekatain, Alexander Novikov, Matej Balog, M Pawan Kumar, Emilien Dupont, Francisco JR Ruiz, Jordan S Ellenberg, Pengming Wang, Omar Fawzi, Pushmeet Kohli, and Alhussein Fawzi. 2024. Mathematical discoveries from program search with large language models. *Nature* 625 (01 2024), 468–475. Issue 7995.

[32] Alexandre Rothen. 1945. The ellipsometer, an apparatus to measure thicknesses of thin surface films. *Review of Scientific Instruments* 16, 2 (1945), 26–30.

[33] Mathias Schubert. 2004. *Infrared ellipsometry on semiconductor layer structures: phonons, plasmons, and polaritons*. Vol. 209. Springer Science & Business Media.

[34] Ryan G Toomey. 2024. Tackling the inverse problem in ellipsometry: analytic expressions for supported coatings with nonuniform refractive index profiles in the thin film and weak contrast limits. *Physica Scripta* 99, 3 (2024), 035529.

[35] Niki van Stein and Thomas Bäck. 2024. LLaMEA: A Large Language Model Evolutionary Algorithm for Automatically Generating Metaheuristics. *IEEE Transactions on Evolutionary Computation* (2024), 1–1. doi:10.1109/TEVC.2024.3497793

[36] Niki van Stein, Diederick Vermetten, and Thomas Bäck. 2024. In-the-loop hyperparameter optimization for llm-based automated design of heuristics. *arXiv preprint arXiv:2410.16309* (2024).

[37] Haoran Ye, Jiarui Wang, Zhiguang Cao, Federico Berto, Chuanbo Hua, Haeyeon Kim, Jinkyoo Park, and Guojie Song. 2024. Reevo: Large language models as hyper-heuristics with reflective evolution. *arXiv preprint arXiv:2402.01145* (2024).

[38] Haoran Yin, Anna V Kononova, Thomas Bäck, and Niki van Stein. 2024. Controlling the Mutation in Large Language Models for the Efficient Evolution of Algorithms. *arXiv preprint arXiv:2412.03250* (2024).

[39] Yan Zhan, Chang Li, Zhigang Che, Ho Cheung Shum, Xiaotian Hu, and Huizeng Li. 2023. Light management using photonic structures towards high-index perovskite optoelectronics: fundamentals, designing, and applications. *Energy & Environmental Science* (2023).

[40] Stefan Zollner. 2013. Spectroscopic ellipsometry for inline process control in the semiconductor industry. In *Ellipsometry at the Nanoscale*. Springer, 607–627.