

# Inverse Design of Photonic Crystal Surface Emitting Lasers is a Sequence Modeling Problem

Ceyao Zhang<sup>1,2\*</sup>, Renjie Li<sup>2,3\*</sup>, Cheng Zhang<sup>2\*</sup>, Zhaoyu Zhang<sup>2,3†</sup>, Feng Yin<sup>2†</sup>

<sup>1</sup>Future Network of Intelligence Institute (FNii), The Chinese University of Hong Kong, Shenzhen

<sup>2</sup>School of Science and Engineering (SSE), The Chinese University of Hong Kong, Shenzhen

<sup>3</sup>Shenzhen Key Laboratory of Semiconductor Lasers

{ceyaozhang2@link., renjieli@link., zhangzy, yinfeng}@cuhk.edu.cn

## Abstract

Photonic Crystal Surface Emitting Lasers (PCSEL)'s inverse design demands expert knowledge in physics, materials science, and quantum mechanics which is prohibitively labor-intensive. Advanced AI technologies, especially reinforcement learning (RL), have emerged as a powerful tool to augment and accelerate this inverse design process. By modeling the inverse design of PCSEL as a sequential decision-making problem, RL approaches can construct a satisfactory PCSEL structure from scratch. However, the data inefficiency resulting from online interactions with precise and expensive simulation environments impedes the broader applicability of RL approaches. Recently, sequential models, especially the Transformer architecture, have exhibited compelling performance in sequential decision-making problems due to their simplicity and scalability to large language models. In this paper, we introduce a novel framework named *PCSEL Inverse Design Transformer* (PiT) that abstracts the inverse design of PCSEL as a sequence modeling problem. The central part of our PiT is a Transformer-based structure that leverages the past trajectories and current states to predict the current actions. Compared with the traditional RL approaches, PiT can output the optimal actions and achieve target PCSEL designs by leveraging offline data and conditioning on the desired return. Results demonstrate that PiT achieves superior performance and data efficiency compared to baselines.

## Introduction

Photonic Crystal Surface Emitting Lasers (PCSELS; Hirose et al. 2014; Yoshida et al. 2019; Noda et al. 2017) are a type of nanoscale laser that combines the benefits of photonic crystals (PhC; Quan, Deotare, and Loncar 2010) and Vertical Cavity Surface Emitting Lasers (VCSELS; Chang-Hasnain 2000). PhCs are artificial structures that have a periodic refractive index modulation in semiconductor materials. This periodicity creates a photonic bandgap that inhibits the propagation of light in certain frequency ranges to amplify the optical resonance effect. VCSELS are lasers that emit light perpendicular to the surface of the semiconductor structure, which allows for efficient coupling to optical fibers and other

optical components. PCSELS combine these two technologies to create advanced lasers that have several advantages over traditional ones and therefore enjoy the best of both worlds. PCSELS have great potential for important applications in sensing, autonomous driving, medicine, machining, and telecommunication. Fundamentally, when an electrical pumping current is injected into the active layer (that is, the core layer of PCSEL), it emits laser light which is then confined and amplified within the PhC resonant cavity (Sze, Li, and Ng 2021). Additionally, the active layer may contain quantum wells that increase the recombination rate of spontaneous photon emission (which was first predicted by Albert Einstein in his quantum mechanics papers (Hilborn 1982)) and thus substantially enhance the lasing effect (Sze, Li, and Ng 2021). So the bottom line is that the PhC layer is used to control the amplitude and direction of the emitted light, while the active layer is what actually generates the light. Therefore, proper design of the PhC layer and the active layer plays a central role in the overall quality of a PCSEL, which brings to us the critical yet challenging PCSEL inverse design problem. This problem is a combination of disciplines including physics, materials science, and quantum mechanics which generally demands prohibitively high-level domain knowledge.

Recent developments in Transformer (Vaswani et al. 2017) architectures witness the blooming performance from prediction tasks into decision-making problems (Wen et al. 2023). This advancement, namely, sequential decision-making as sequence modeling, is notably distinct from the traditional approaches in reinforcement learning (RL), which are typically focused on learning a singular policy for a specific, narrowly defined behavior distribution. Given the broad range of successful implementations of these models, this study aims to explore their applicability to the inverse design of PCSEL.

In this paper, we introduce a novel framework named PCSEL Inverse Design Transformer (PiT) that abstracts the inverse design of PCSEL as a sequence modeling problem. The central part of our PiT is a Transformer-based structure that leverages the past trajectories and current states to predict the current actions. Compared with the traditional RL approaches, PiT can output the optimal actions and achieve target PCSEL designs by leveraging offline data and condi-

\*These authors contributed equally.

†Corresponding author




Figure 1: PCSEL structure schematic, with the active layer and PhC cavity layer shown. Based on the band-edge mode, a large area lasing resonance mode is formed within the PhC resonance cavity. Lasing arises from the evanescent coupled MQW gain medium in the active layer. The inverse design problem therefore focuses on optimizing the designs of the active layer and the PhC cavity as a whole. The size of the device is about 200 microns in side lengths and 400 microns in height and solely manufactured with semiconductor materials.

tioning on the desired return. Results demonstrate that PiT achieves superior performance and data efficiency compared to baselines.

## Related Works

**Optimizing the PCSEL through AI** Overall, recent advancements in machine learning (Goodfellow, Bengio, and Courville 2016) and optimization algorithms (Wiecha et al. 2017) have propelled the progress of inverse designs in science. Early in the 90s, heuristic, evolutionary (Hegde 2019), and gradient-based (Zhang et al. 2020) optimization algorithms began to emerge prolifically. Key algorithms include simulated annealing (Bertsimas and Tsitsiklis 1993), Newton’s method (Milzarek and Ulbrich 2014), Bayesian optimization (Shahriari et al. 2015), Monte Carlo method (Rubinstein and Kroese 2016), particle swarm (Ma and Li 2020) and genetic algorithm (Ren et al. 2021) etc. These algorithms provide a new way of thinking when facing hard non-convex optimization problems, which act as a solid foundation for scientific problems. However, the issue remains of heavy human involvement due to sophisticated trial-and-error iterations. To solve this, in around 2012, researchers proposed deep learning (DL) (Krizhevsky, Sutskever, and Hinton 2012; Goodfellow, Bengio, and Courville 2016) frameworks to construct a mapping relationship between input data and output targets through Deep Neural Networks. In particular, DL consists of supervised, unsupervised, and reinforcement learning (RL) (Sutton and Barto 2018). These DL models greatly bolstered the efficiency of inverse design in science, pushing the possibility of automated design into a new era (So et al. 2020; Jiang, Chen, and Fan 2021; Li et al. 2021; Mirhoseini et al. 2021; Li et al. 2022; Degraive et al. 2022; Li et al. 2023a; Kuprikov et al. 2022). Circa 2023, a novel framework based on RL, called Learning to De-

sign Optical-Resonators (L2DO) (Li et al. 2023b), provides the solution for autonomous inverse design of nanophotonic chips without human intervention. With two orders of magnitude higher sample efficiency compared to supervised learning, L2DO has preliminarily realized RL-driven chip inverse design on an algorithmic level. However, to the best of our knowledge, there aren’t any published results on the inverse design of PCSELS via AI methods to this day. Due to the strategic significance of PCSELS for a host of key industries, we believe there is an urgent need to develop an RL-based approach for PCSEL’s rapid inverse design.

**Sequential Models for Sequential Decision-Making** Recent research on transformer models for sequential decision-making has revolutionized the field by pushing boundaries and introducing novel capabilities. Transformers (Wen et al. 2023) excel in capturing long-range dependencies, efficiently assigning credit, and modeling temporal dynamics. These have significantly advanced the field and hold promise for applications in reinforcement learning. Foundational works in the application of transformers to sequential decision-making include Decision Transformer (DT) (Chen et al. 2021a) and Trajectory Transformer (TT) (Janner, Li, and Levine 2021). These works have made a significant impact on the field by introducing transformative methodologies. DT, for instance, revolutionized traditional offline RL approaches by employing a transformer decoder as the backbone model. It uses (return-to-go, state) as input data, and chooses action as output. This method outperforms the offline RL baseline CQL. TT is another foundational application of transformers in sequential decision-making, which utilizes the GPT model (Radford et al. 2018) as its backbone, demonstrating remarkable performance in long-range planning tasks. Other efforts like UPDeT (Hu et al. 2021) for multi-agent planning and MGDT (Lee et al. 2022) for multi-game decision-making highlight the versatility of transformers. Looking ahead, transformer models show great promise to transform decision-making in the coming years by enabling a more nuanced understanding of situational context and more predictive recommendations.

## Background

### Modeling Inverse Design of PCSEL as a Sequential Decision-making problem

Sequential decision-making describes a situation where the decision-maker makes successive observations of a process before a final decision is made. In most sequential decision problems, there is an implicit or explicit cost/regret/reward associated with each observation or action. The procedure to decide when to stop taking observations and when to continue is called the ‘stopping criteria’. The objective of sequential decision-making is to find a stopping criterion that optimizes the decision in terms of minimizing losses or maximizing returns, including observation costs. The optimal stopping criteria are also called optimal strategy and optimal policy, which are commonly adopted in classic RL algorithms.

To the best of our knowledge, Li et al. (2023b) is the first work on modeling the inverse design of PCSELS as a

Sequential decision-making problem, which offers a structured and efficient approach to navigating the complex design space. Sequential Decision-Making involves breaking down the design process into a series of decisions, each contingent upon the outcomes of preceding ones. This approach is particularly suited to PCSEL design due to the layered nature of their construction and the interdependence of various design parameters. There are a few key components of the sequential decision-making model.

- State Space represents the current status of the PCSEL design, encompassing all relevant parameters (e.g., layer thicknesses, refractive indices, geometric patterns).
- Action space is discrete and consists of 16 actions. Each action represents an increase or decrease in geometric parameters by one unit according to the scale of the parameters. For example, action 0 is to increase parameter 0 by 25, and action 3 is to decrease parameter 1 by 2.5.
- Reward: As to the quality of a PCSEL, there are a few indicators to determine its performance, which we listed in Table 1. Unify the dimensions between different indicators through weighting, we can calculate the score from the parameters returned at each step, as in Eqn 6. In PCSEL, the reward function is the difference between the score obtained in the current step and the score obtained in the previous step. As the first step, we take the score as a reward. To analyze the composition of the score in more detail, we can examine the parameters returned by the environment and the method of calculation.

notations	Indicators
$Q$	Q-factor
$lam$	lambda
$power$	power
$area$	area
$div\_angle$	divergence angle

Table 1: Indicators that determine PCSEL performance.

$$r_1 = 1 - (Q_{goal} - Q)/Q_{goal} \quad (1)$$

$$r_2 = 1 - |lam_{goal} - lam|/lam_{goal} \quad (2)$$

$$r_3 = 1 - (area_{goal} - area)/area_{goal} \quad (3)$$

$$r_4 = 1 - (power_{goal} - power)/power_{goal} \quad (4)$$

$$r_5 = 1 + (div\_angle_{goal} - div\_angle)/div\_angle_{goal} \quad (5)$$

$$score = \gamma * r_1 + \epsilon * r_2 + \beta * r_3 + \alpha * r_4 + \eta * r_5 \quad (6)$$

The designer continues to take actions to set/update the layout until a pre-set stopping criteria is met, which in our case is the target PCSEL performance characteristics listed in Table 1. When the target characteristics are met, one can consider the cumulative return is maximized or an optimal policy is found. Due to the high-level domain knowledge and labor intensity required by PCSEL inverse design, the authors believe modeling it as a sequential decision problem can largely alleviate human labor and accelerate the R&D of advanced PCSEL lasers.

## Proposed Methods

### Inverse Design of PCSEL is a Sequential Modeling Problem




Figure 2: The architecture of PiT.

The PCSEL inverse design can be conceptualized as a sequential decision model, where the goal is to maximize the score within a limited number of steps through strategic adjustments. This model is characterized by a clear reward target, a small and discrete action space, and the Markov property. Given these properties, it is naturally suited for RL algorithms.

However, unlike games that can be quickly simulated, data collection for PCSEL inverse design is challenging due to its complexity. Therefore, considering data efficiency, offline RL is a more suitable choice. Our PiT model is based on sequence models, for example, DT (Chen et al. 2021a) or TT (Janner, Li, and Levine 2021). Among the existing sequence models for sequential decision-making, we choose DT for its simplicity and superior performance.

### Dataset




Figure 3: The frequency of the total return of the trajectories used to train PiT.

The dataset of offline trajectories for training our PiT is extracted from replay buffers that were saved during previous online RL experiments associated with PCSELS. The

offline dataset contains roughly 16,057 samples (each sample includes state, action, reward, and next state). Specifically, the state is defined as a vector of design parameters of the PCSEL, the action is a change/update in the state, and the reward is defined as how close we are to the target PCSEL design. For instance, if we wish to achieve a target PCSEL design with a Q-factor of 5 million, a wavelength ( $\lambda$ ) of 1310 nm, a power conversion rate of 80%, a lasing area of  $3.0 \times 10^{-13}$  square meters, and a divergence angle of 1.0 degrees, we would train PiT to inverse design a PCSEL that meets these target characteristics. The actual criteria is a weighted sum of these five reward of merit according to Eqn. 1-6 and is referred to as a score in the remaining text.

## Experiments

In this section, we performed a comprehensive performance comparison involving our PiT and behavior cloning. Models were trained on both the whole dataset and a subset comprising exclusively of the rising dataset, where the score at the last timestep exceeded the initial timestep.

### Baseline

We choose behavior cloning as our baseline. Behavior cloning involves training a network to imitate the behavior of a demonstrated or an expert dataset. The network is an MLP with 3 layers and 256 embedding dimensions. It is a form of supervised learning where the model learns to map observations to actions directly by mimicking the expert’s actions.




Figure 4: Training loss curve of both BC and PiT under the whole and rising dataset.

From Figure 4, it is evident that the PiT exhibits a lower training loss compared to baseline BC. This lower training loss indicates that PiT has a greater ability to predict the next action accurately on the training set.

### Results

In Table 2, we compare PiT’s results to the best literature baselines. The results are quantified by a score which is defined in Eqn. 6. Moreover, thanks to the offline RL framework, PiT obtained better data efficiency than the baselines as it doesn’t require online environment interactions or labor-intensive manual optimizations by human designers.

Therefore, we conclude that PiT has acquired state-of-the-art capabilities for PCSEL inverse design with consistently better performance than baselines. Nonetheless, we recognize that PiT’s score is still quite some distance away from the target, which means there’s room for improvement in future research.

Imada et al. (1999)	9.19709
Ohnishi et al. (2004)	19.2290
Sakai et al. (2005)	21.23469
Hirose et al. (2014)	18.95918
Hsu, Lin, and Pan (2017)	10.01160
Chen et al. (2021b)	2.117346
Wang et al. (2021)	41.8554
Itoh et al. (2022)	20.2335
<b>BC</b>	71.28
<b>PiT</b>	73.95

Table 2: PiT’s results compared to baselines in the literature.

### Discussion

In this section, we conducted a performance comparison between PiT and BC trained on both the whole dataset and the rising dataset. The results are listed in Table 3. According to the table, we can see that regardless of whether BC or PiT methods are used, the results based on the rising dataset are better than the results of the whole dataset. This phenomenon indicates that the offline dataset we choose has a large impact on performance. More specifically, the better the data set we use (i.e. the larger the reward), the better the final performance will be. Therefore, if we can train a better RL policy and collect a dataset with higher returns, it is possible to further improve the performance of PCSEL by a large margin.

Dataset \ Method	BC	PiT
Whole	68.51	70.54
Rising	71.28	73.95

Table 3: Performance Comparison for PiT and BC under the whole and rising dataset.

Finally, we discuss the potential limitations of our method as we as future work. First, the effectiveness of the PiT model is heavily dependent on the availability and quality of offline data, which could be a limiting factor in some scenarios where data is scarce. Second, while effective in the specific context of PCSEL inverse design, it’s unclear how well this approach generalizes to other types of photonic devices or design objectives. In the future, we will expand the realm of PiT by demonstrating its effectiveness and versatility for different types of devices. We will also generate a diverse pool of data for photonic inverse design problems and make it open-source.

## Conclusion

In this paper, we investigate the PCSEL inverse design problem through sequential modeling that facilitates the use of

offline data and eliminates the need for online interactions. Our simulation experiments show the effectiveness of our proposed framework. We believe that our work points out promising future avenues to design advanced PCSEL lasers and photonics in general. We further analyze the impact of Transformer structure selection and highlight several potential solutions to improve the performance in the future.

### Acknowledgement

The work was supported in part by the National Natural Science Foundation of China (NSFC) under Grant No. 62271433, by the Basic Research Project No. HZQB-KCZYZ-2021067 of Hetao Shenzhen-HK S&T Cooperation Zone, by the National Key R&D Program of China with grant No. 2018YFB1800800, by NSFC with Grant No. 62293482, by the Shenzhen Outstanding Talents Training Fund 202002, by the Guangdong Research Projects No. 2017ZT07X152 and No. 2019CX01X104, by the Guangdong Provincial Key Laboratory of Future Networks of Intelligence (Grant No. 2022B1212010001), by the Shenzhen Key Laboratory of Big Data and Artificial Intelligence (Grant No. ZDSYS201707251409055), and by Shenzhen Science and Technology Program under Grant No. JCYJ20220530143806016 and No. RCJC20210609104448114, by NSFC under Grant No.62174144, Shenzhen Science and Technology Program under Grant No.JCYJ20210324115605016, No.JCYJ20210324120204011, No.JSGG20210802153540017, and No.JCYJ20220818102214030, Guangdong Key Laboratory of Optoelectronic Materials and Chips under Grant No.2022KSYS014, Shenzhen Key Laboratory Project under Grant No.ZDSYS201603311644527; Longgang Key Laboratory Project under Grant No.ZSYS2017003 and No.LGKCZSYS2018000015; Shenzhen Research Institute of Big Data; President's Fund (PF01000154).

### References

Bertsimas, D.; and Tsitsiklis, J. 1993. Simulated annealing. *Statistical science*, 8(1): 10–15.

Chang-Hasnain, C. J. 2000. Tunable vcsel. *IEEE Journal of Selected Topics in Quantum Electronics*, 6(6): 978–987.

Chen, L.; Lu, K.; Rajeswaran, A.; Lee, K.; Grover, A.; Laskin, M.; Abbeel, P.; Srinivas, A.; and Mordatch, I. 2021a. Decision transformer: Reinforcement learning via sequence modeling. *Advances in neural information processing systems*, 34: 15084–15097.

Chen, L.-R.; Hong, K.-B.; Huang, K.-C.; Yen, H.-T.; and Lu, T.-C. 2021b. Improvement of output efficiency of p-face up photonic-crystal surface-emitting lasers. *Optics Express*, 29(7): 11293–11300.

Degrave, J.; Felici, F.; Buchli, J.; Neunert, M.; Tracey, B.; Carpanese, F.; Ewalds, T.; Hafner, R.; Abdolmaleki, A.; de Las Casas, D.; et al. 2022. Magnetic control of tokamak plasmas through deep reinforcement learning. *Nature*, 602(7897): 414–419.

Goodfellow, I.; Bengio, Y.; and Courville, A. 2016. *Deep learning*. MIT press.

Hegde, R. S. 2019. Photonics inverse design: pairing deep neural networks with evolutionary algorithms. *IEEE Journal of Selected Topics in Quantum Electronics*, 26(1): 1–8.

Hilborn, R. C. 1982. Einstein coefficients, cross sections, f values, dipole moments, and all that. *American Journal of Physics*, 50(11): 982–986.

Hirose, K.; Liang, Y.; Kurosaka, Y.; Watanabe, A.; Sugiyama, T.; and Noda, S. 2014. Watt-class high-power, high-beam-quality photonic-crystal lasers. *Nature photonics*, 8(5): 406–411.

Hsu, M.-Y.; Lin, G.; and Pan, C.-H. 2017. Electrically injected 1.3- $\mu\text{m}$  quantum-dot photonic-crystal surface-emitting lasers. *Optics Express*, 25(26): 32697–32704.

Hu, S.; Zhu, F.; Chang, X.; and Liang, X. 2021. Updet: Universal multi-agent reinforcement learning via policy decoupling with transformers. *arXiv preprint arXiv:2101.08001*.

Imada, M.; Noda, S.; Chutinan, A.; Tokuda, T.; Murata, M.; and Sasaki, G. 1999. Coherent two-dimensional lasing action in surface-emitting laser with triangular-lattice photonic crystal structure. *Applied physics letters*, 75(3): 316–318.

Itoh, Y.; Kono, N.; Inoue, D.; Fujiwara, N.; Ogasawara, M.; Fujii, K.; Yoshinaga, H.; Yagi, H.; Yanagisawa, M.; Yoshida, M.; et al. 2022. High-power CW oscillation of 1.3- $\mu\text{m}$  wavelength InP-based photonic-crystal surface-emitting lasers. *Optics Express*, 30(16): 29539–29545.

Janner, M.; Li, Q.; and Levine, S. 2021. Offline reinforcement learning as one big sequence modeling problem. *Advances in neural information processing systems*, 34: 1273–1286.

Jiang, J.; Chen, M.; and Fan, J. A. 2021. Deep neural networks for the evaluation and design of photonic devices. *Nature Reviews Materials*, 6(8): 679–700.

Krizhevsky, A.; Sutskever, I.; and Hinton, G. E. 2012. ImageNet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25.

Kuprikov, E.; Kokhanovskiy, A.; Serebrennikov, K.; and Turtitsyn, S. 2022. Deep reinforcement learning for self-tuning laser source of dissipative solitons. *Scientific Reports*, 12(1): 1–9.

Lee, K.-H.; Nachum, O.; Yang, M. S.; Lee, L.; Freeman, D.; Guadarrama, S.; Fischer, I.; Xu, W.; Jang, E.; Michalewski, H.; et al. 2022. Multi-game decision transformers. *Advances in Neural Information Processing Systems*, 35: 27921–27936.

Li, R.; Gu, X.; Li, K.; Huang, Y.; Li, Z.; and Zhang, Z. 2021. Deep learning-based modeling of photonic crystal nanocavities. *Optical Materials Express*, 11(7): 2122–2133.

Li, R.; Gu, X.; Shen, Y.; Li, K.; Li, Z.; and Zhang, Z. 2022. Smart and Rapid Design of Nanophotonic Structures by an Adaptive and Regularized Deep Neural Network. *Nanomaterials*, 12(8): 1372.

Li, R.; Zhang, C.; Mao, S.; Huang, H.; Zhong, M.; Cui, Y.; Zhou, X.; Yin, F.; Theodoridis, S.; and Zhang, Z. 2023a. From English to PCSEL: LLM helps design and optimize photonic crystal surface emitting lasers. 14 pages, 9 graphics.

- Li, R.; Zhang, C.; Xie, W.; Gong, Y.; Ding, F.; Dai, H.; Chen, Z.; Yin, F.; and Zhang, Z. 2023b. Deep reinforcement learning empowers automated inverse design and optimization of photonic crystals for nanoscale laser cavities. *Nanophotonics*, 12(2): 319–334.
- Ma, Z.; and Li, Y. 2020. Parameter extraction and inverse design of semiconductor lasers based on the deep learning and particle swarm optimization method. *Optics Express*, 28(15): 21971–21981.
- Milzarek, A.; and Ulbrich, M. 2014. A semismooth Newton method with multidimensional filter globalization for L1-optimization. *SIAM Journal on Optimization*, 24(1): 298–333.
- Mirhoseini, A.; Goldie, A.; Yazgan, M.; Jiang, J. W.; Songhori, E.; Wang, S.; Lee, Y.-J.; Johnson, E.; Pathak, O.; Nazi, A.; et al. 2021. A graph placement methodology for fast chip design. *Nature*, 594(7862): 207–212.
- Noda, S.; Kitamura, K.; Okino, T.; Yasuda, D.; and Tanaka, Y. 2017. Photonic-crystal surface-emitting lasers: Review and introduction of modulated-photonic crystals. *IEEE Journal of Selected Topics in Quantum Electronics*, 23(6): 1–7.
- Ohnishi, D.; Okano, T.; Imada, M.; and Noda, S. 2004. Room temperature continuous wave operation of a surface-emitting two-dimensional photonic crystal diode laser. *optics express*, 12(8): 1562–1568.
- Quan, Q.; Deotare, P. B.; and Loncar, M. 2010. Photonic crystal nanobeam cavity strongly coupled to the feeding waveguide. *Applied Physics Letters*, 96(20): 203102.
- Radford, A.; Narasimhan, K.; Salimans, T.; and Sutskever, I. 2018. Improving language understanding by generative pre-training.
- Ren, Y.; Zhang, L.; Wang, W.; Wang, X.; Lei, Y.; Xue, Y.; Sun, X.; and Zhang, W. 2021. Genetic-algorithm-based deep neural networks for highly efficient photonic device design. *Photonics Research*, 9(6): B247–B252.
- Rubinstein, R. Y.; and Kroese, D. P. 2016. *Simulation and the Monte Carlo method*. John Wiley & Sons.
- Sakai, K.; Miyai, E.; Sakaguchi, T.; Ohnishi, D.; Okano, T.; and Noda, S. 2005. Lasing band-edge identification for a surface-emitting photonic crystal laser. *IEEE Journal on Selected Areas in Communications*, 23(7): 1335–1340.
- Shahriari, B.; Swersky, K.; Wang, Z.; Adams, R. P.; and De Freitas, N. 2015. Taking the human out of the loop: A review of Bayesian optimization. *Proceedings of the IEEE*, 104(1): 148–175.
- So, S.; Badloe, T.; Noh, J.; Bravo-Abad, J.; and Rho, J. 2020. Deep learning enabled inverse design in nanophotonics. *Nanophotonics*, 9(5): 1041–1057.
- Sutton, R. S.; and Barto, A. G. 2018. *Reinforcement learning: An introduction*. MIT press.
- Sze, S. M.; Li, Y.; and Ng, K. K. 2021. *Physics of semiconductor devices*. John wiley & sons.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Wang, Z.; Tong, C.; Wang, L.; Lu, H.; Tian, S.; and Wang, L. 2021. Photonic crystal surface emitting laser operating in pulse-periodic regime with ultralow divergence angle. In *Photonics*, volume 8, 323. MDPI.
- Wen, M.; Lin, R.; Wang, H.; Yang, Y.; Wen, Y.; Mai, L.; Wang, J.; Zhang, H.; and Zhang, W. 2023. Large sequence models for sequential decision-making: a survey. *Frontiers of Computer Science*, 17(6): 176349.
- Wiecha, P. R.; Arbouet, A.; Girard, C.; Lecestre, A.; Larrieu, G.; and Paillard, V. 2017. Evolutionary multi-objective optimization of colour pixels based on dielectric nanoantennas. *Nature nanotechnology*, 12(2): 163–169.
- Yoshida, M.; De Zoysa, M.; Ishizaki, K.; Tanaka, Y.; Kawasaki, M.; Hatsuda, R.; Song, B.; Gellesta, J.; and Noda, S. 2019. Double-lattice photonic-crystal resonators enabling high-brightness semiconductor lasers with symmetric narrow-divergence beams. *Nature materials*, 18(2): 121–128.
- Zhang, J.; Xiao, P.; Sun, R.; and Luo, Z. 2020. A single-loop smoothed gradient descent-ascent algorithm for nonconvex-concave min-max problems. *Advances in neural information processing systems*, 33: 7377–7389.