Fine-Tuning Multi-Modal LLMs for Materials Discovery

Daniel Ceballos, Renjie Li, Lynford Goddard

***Abstract*—The discovery of new materials is important for technological progress, but is usually a slow and laborious process. Advancements in artificial intelligence, such as large language models (LLMs) and vision transformers, are resulting in faster and more efficient approaches to polymer design. Here, we present a viable pipeline assisted by transformer-based LLMs and vision transformers that accelerate the development of novel photoresists for photonic integrated circuits (PICs).**

***Index Terms*—Large language models, semiconductor lasers, inverse design, PCSEL, photonic crystal, CSL, machine learning, photonic integrated circuits, polymer informatics**

I. Introduction

With global data usage growing at approximately 20% annually, we need higher bandwidth methods of data transfer [1]. Photonic integrated circuits (PICs) present a unique solution. Using light rather than electricity to transmit information and integrating lasers, waveguides, modulators, photodetectors, and other components on a single substrate to maximize efficiency and miniaturization [2, 3]. Fundamentally, light is described as electromagnetic waves that oscillate through space according to Maxwell's equations. First described by Einstein, at the quantum level, light is made up of photons, massless bosons carrying both energy and momentum, acting as both particles and electromagnetic waves. [4].

This allows them to transmit information efficiently via total internal reflection in fiber optics. [2]. PICs are a promising alternative to electronics; no ohmic heating (less energy loss), no inertia (photons are massless), immunity to electromagnetic interference, and bandwidths that can exceed electronic transmission by one to two orders of magnitude [1, 2, 5]. For example, state-of-the-art systems have achieved petabit-per-second data rates. [5]. Electronic circuits typically max out at 40–400 Gbps per channel, which is insufficient for the future demands of science, finance, medicine, and AI.

In contrast to electronic circuits, which rely on transistors for logic, PICs use nanoscale photonic devices to focus, split, isolate, polarize, couple, and modulate light [2, 3]. The main "active devices" in a PIC are: lasers (generate photons at specific wavelengths for signal transmission), waveguides (channeling photons through the circuit, analogous to wire in electronic circuits), photodetectors (convert light into electrical signals), and modulators (manipulate phase, amplitude, or polarization to encode information) [2, 6]. Current progress in photonic-crystal surface-emitting lasers (PCSELs), the on-chip laser, is further improving PIC scalability and performance [4, 6].

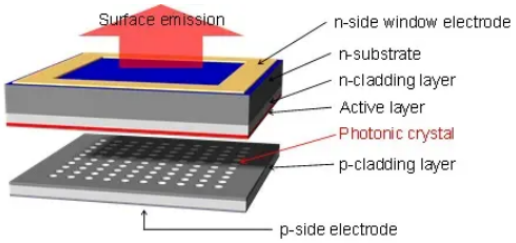

Fig. 1: Schematic structure of a photonic crystal surface-emitting laser (PCSEL). [**rp-photonics**].

The development of new materials for these devices, especially novel photoresists suitable for advanced two-photon polymerization (2PP), is severely limited by the vast chemical space of possible polymers [7]. Traditional polymer-chemistry R & D can take five or more years to deliver meaningful results. To expedite the process, advanced machine learning and LLM/AI methods present a potentially significant improvement in the speed and efficiency with which new polymers could be discovered [7–9].

II. Machine Learning for Accelerated Materials Discovery

The past couple years we have seen massive improvements in materials discovery with AI and machine learning [7]. Large language models (LLMs) for chemistry like ChemCrow [10] utilizing transformer architectures, have allowed faster and cheaper prediction of chemical properties and synthesis routes. [10, 11]. Transformers, first introduced by Vaswani et al. [12], use self-attention to model relationships within sequences and are effective for SMILES string representations. SMILES was the chosen route for monomer/polymer representation because of their ease of understanding and interpretation by transformer networks, which allowed our models to more efficiently learn and predict specific chemical properties. Unlike convolutional networks, transformers capture both local and distant dependencies, making them valuable for chemistry prediction [11, 13].

Vision transformers (ViTs) use these principles but for visual data, splitting images into arrays and processing them as tokens, which then extract morphological and structural information

(*Corresponding authors: Daniel Ceballos, email: deceba2@illinois.edu ;Dr. Renjie Li, email: renjie2@illinois.edu; and Lynford Goddard, email: lgoddard@illinois.edu*)

Dr. Renjie Li and Dr. Lynford Goddard are with the Electrical and Computer Engineering Department at the University of Illinois Urbana-Champaign.

Daniel Ceballos is with the Siebel Center for Computer Science at the University of Illinois Urbana-Champaign.

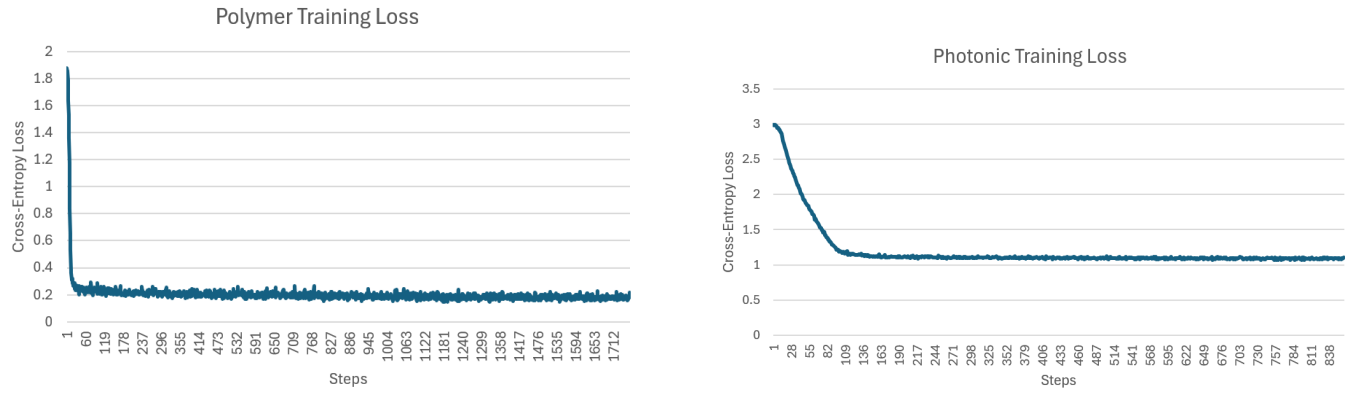

(a) Polymer-property prediction

(b) Photonic-device performance prediction

Fig. 2: Training-loss convergence of the fine-tuned Qwen models. Both tasks reach a stable minimum after ∼5 epochs, with the photonic model requiring slightly more iterations to plateau.

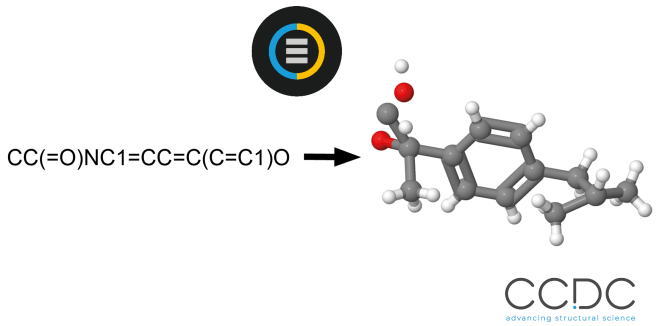

Fig. 3: CC(=O)NC1=CC=C(C=C1)O) into its 3-D, energy-minimized ball-and-stick model. The structure was generated and rendered with CCDC Mercury.

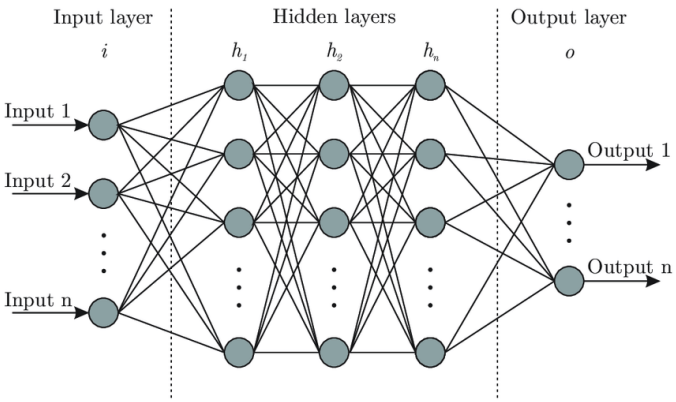

Fig. 4: Feedforward neural network. The input layer receives n inputs, which are processed through multiple hidden layers before producing n outputs.

and can make predictions based on these. [10, 14]. Our ViTs were trained on labeled image datasets to identify characteristics necessary for photonic device performance [4, 14].

Our approach integrates both modes: transformers for chemical structure prediction and ViTs for device image analysis, enabling a multi-modal AI pipeline for identifying new photoresist resins for two-photon polymerization in PIC fabrication [7, 10].

III. Methods

We began by reviewing current literature on neural networks, vision transformers, transformers, AI/LLMs, polymer chemistry, and multi-photon lithography [6–8, 14]. We eventually began our first phase which was utilizing ChemCrow. ChemCrowhttps://github.com/ur-whitelab/chemcrow-public, version v0.3.24, commit 9a18c01) was deployed into the UIUC ICRN JupyterHub environment, which provided access to PyTorch 2.x (preinstalled), eight NVIDIA A100 GPUs, and dual Intel Xeon CPUs. ChemCrow's published success in automated synthesis planning was why it was selected [10]. Alternative models with future implementation like the Molecular Transformer and trained multi-modal LLMs have further demonstrated high accuracy in chemical properties and thus reaction prediction [13] [11].

The next phase attempted to implement various curated datasets of structured chemical information into the open source LLM Qwen. The implementation we built in was UIUC's ICRN JupyterHub environment. Utilizing LoRa adapters, unsloth, ollama, hugging face, and numerous other features and libraries like RDKit, SciPy, semantic scholar, paperscraper, and PubChem Search. We extracted, parsed, and collected useful chemical information and academic standard methods of synthesis of these polymers. [8, 9]. These coding pipelines sufficiently provide the data for large-scale model training and fine-tuning.

IV. Results

We generated and evaluated three distinct synthesis plans, each of which proposed different properties for the target photoresist resin. The best combination identified so far uses Lucirin TPO, pentaerythritol tetraacrylate (PETA) and rhodamine 6G.

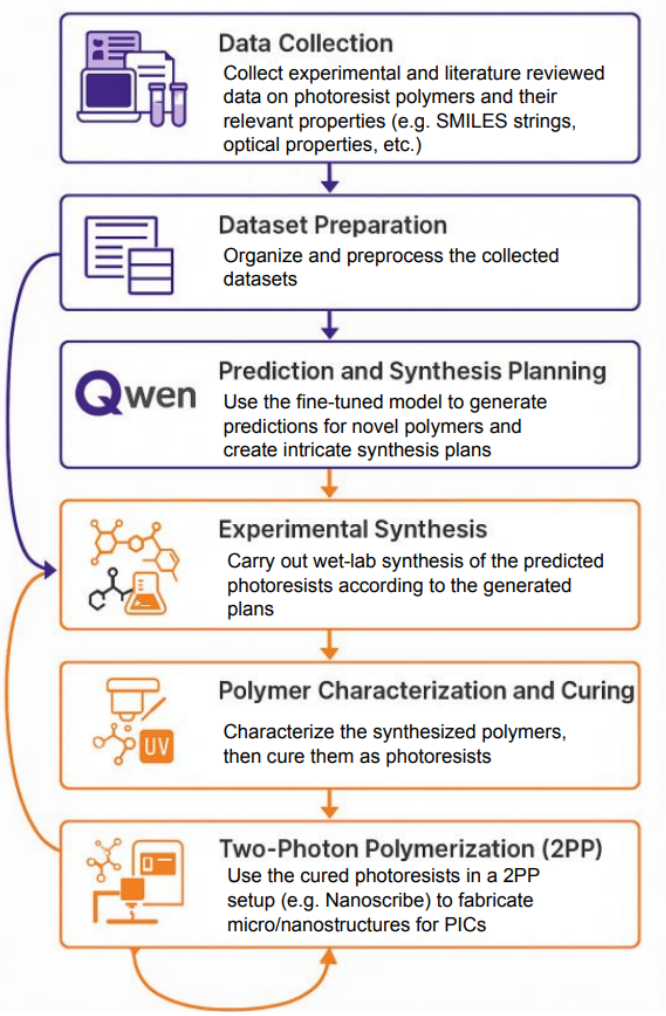

Fig. 5: Workflow interface. This figure describes the iterative workflow of our procedure in training our LLM and selecting possible polymers.

Our preliminary synthesis included the mixture of the photoinitiator Lucirin TPO-L and the monomer PETA (Pentaerythritol Tetraacrylate) with the solvent ethyl acetate and cosolvent PGMEA (propylene glycol Methyl ethyl acetate). Lucirin was chosen as the most optimal photoinitiator because of its excellent two-photon absorption cross-section. [8, 9].

V. Discussion

This work demonstrates the transformer-based models, multimodal LLM integration, and simulation in materials discovery. As transformer and vision transformer-based LLMs improve in accordance with Moore's Law, our work here will only improve the natural language processing, vision processing, and workflow demands required for complex polymer chemistry prediction to design novel photoresists for laser components. On-chip laser sources and nanophotonic circuits will allow us to gain even better performance out of PICs. Ongoing improvements in active device components on PICs through LLM-powered materials discovery will assist in scaling PICs to meet the demands for data, AI, and quantum-technology. LLM-powered materials science research are promising in R&D in industry and academia. [15] [4, 6] [4] [7]

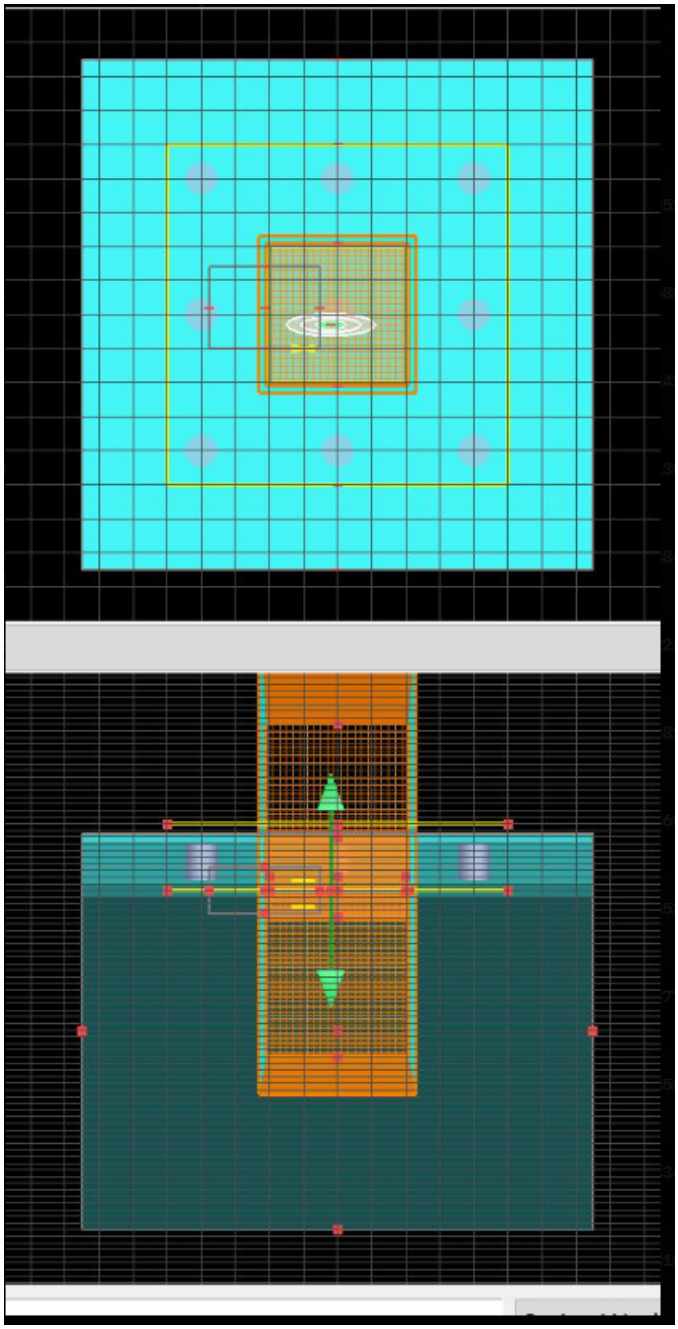

Fig. 6: Visualization of a photonic crystal surface-emitting laser (PCSEL) structure used for Qwen model training and evaluation. The upper panel shows the top view of the PCSEL. The lower panel shows a cross-sectional view,. These inputs were used for model learning of device property prediction.

VI. Conclusion

Our work here illustrates a novel approach to discovering functional polymers through LLM/AI training and augmentation. Our pipeline presents possible AI workflow for accelerating polymer-chemistry research, making discovery

exponentially faster, more efficient, and more accessible to beginners. This shift in academia and industry will continue to revolutionize the photonic-integrated-circuit industry.

Acknowledgement

This research was supported by the IBM-Illinois Discovery Accelerator Institute (IIDAI), the University of Illinois Urbana-Champaign, including the Grainger College of Engineering and the Department of Electrical and Computer Engineering, as well as the GEAR UP program and its constituents.

Conflict of interest

The authors declare no conflict of interest.

Data and Code availability

The data and code are available at https://github.com/Arcadianlee/LLM-for-Photonics. Requests for code, data, and files should be addressed to Dr. Renjie Li, email: renjie2@illinois.edu; and Lynford Goddard, email: lgoddard@illinois.edu.

Supplementary material

The supplementary material is available in the online version of this manuscript.

References

[1] Cisco. *Cisco Annual Internet Report (2018–2023), Executive Perspectives White Paper*. https://www.cisco.com/c/en/us/solutions/collateral/executive-perspectives/annual-internet-report/white-paper-c11-741490.html. Accessed: 2024-08-07. 2020.

[2] M. A. Butt, B. Janaszek, and R. Piramidowicz. "Lighting the way forward: The bright future of photonic integrated circuits". In: *Sensors International* 6 (2025), p. 100326. DOI: 10.1016/j.sintl.2025.100326.

[3] R. Chandrasekar et al. "Photonic integrated circuits for Department of Defense-relevant chemical and biological sensing applications: state-of-the-art and future outlooks". In: *Optical Engineering* 58.2 (2019), p. 020901. DOI: 10.1117/1.OE.58.2.020901.

[4] Y. Shen et al. "Deep learning with coherent nanophotonic circuits". In: *Nature Photonics* 11.7 (2017), pp. 441–446. DOI: 10.1038/nphoton.2017.93.

[5] A. A. Jørgensen et al. "Petabit-per-second data transmission using a chip-scale microcomb ring resonator source". In: *Nature Photonics* 16 (2022), pp. 798–802. DOI: 10.1038/s41566-022-01082-z.

[6] Y. Yang et al. "From past to future: on-chip laser sources for photonic integrated circuits". In: *Light: Science & Applications* 12.1 (2023), p. 19. DOI: 10.1038/s41377-022-01006-0.

[7] K. T. Butler et al. "Machine learning for molecular and materials science". In: *Nature* 559 (2018), pp. 547–555.

[8] T. Kudo, H. Kondo, and S. Kuwajima. "RadonPy: automated physical property calculation using all-atom classical molecular dynamics simulations for polymer informatics". In: *npj Computational Materials* 10 (2024), p. 148. DOI: 10.1038/s41524-024-01492-3.

[9] M. Suzuki, H. Kondo, and S. Kuwajima. "Procedural Construction of Atomistic Polyurethane Block Copolymer Models for High Throughput Simulations". In: *arXiv preprint* (2024). arXiv: 2506.02129 [cond-mat.soft]. URL: https://arxiv.org/abs/2506.02129.

[10] A. M. Bran et al. "Augmenting large language models with chemistry tools". In: *Nature Machine Intelligence* 6.5 (2024), pp. 525–535. DOI: 10.1038/s42256-024-00832-8.

[11] M. H. S. Segler, M. Preuss, and M. P. Waller. "Planning chemical syntheses with deep neural networks and symbolic AI". In: *Nature* 555.7698 (2018), pp. 604–610. DOI: 10.1038/nature25978.

[12] A. Vaswani et al. "Attention Is All You Need". In: *Advances in Neural Information Processing Systems*. Vol. 30. 2017. URL: https://arxiv.org/abs/1706.03762.

[13] P. Schwaller et al. "Molecular Transformer: A Model for Uncertainty-Calibrated Chemical Reaction Prediction". In: *ACS Central Science* 5.9 (2019), pp. 1572–1583.

[14] A. Dosovitskiy et al. "An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale". In: *International Conference on Learning Representations*. 2021. URL: https://arxiv.org/abs/2010.11929.

[15] D. A. B. Miller. "Attojoule Optoelectronics for Low-Energy Information Processing and Communications". In: *Journal of Lightwave Technology* 35.3 (2017), pp. 346–396. DOI: 10.1109/JLT.2016.2638458.