Enabling high-volume production of photonics chips with machine learning

K. Yadav*[a], S. Bidnyk[a], A. Balakrishnan[a]
[a]Enablence Technologies Inc., 390 March Road, Suite 119, Ottawa, Canada, K2K 0G7

ABSTRACT

Leveraging the power of machine learning, we introduce a breakthrough approach in high-volume manufacturing of photonics chips for advanced applications. Despite the transformative potential of photonics in many industries, its widespread adoption has been hindered by multiple challenges in the fabrication of complex integrated chips. We deployed machine learning models with diverse architectures at every stage of our manufacturing process to overcome these challenges. Inevitable variations in the fabrication process often lead to performance variability among photonics chips on a single wafer and across different wafers. We effectively overcome this challenge by employing a deep neural network to study the variability in the performance of individual chips, enabling us to predict the precise optimizations necessary to compensate for inevitable process variations. We describe our selection of the deep neural network architecture that addresses this challenge, our methodology for obtaining a high-quality dataset for training, and the enhancements in performance uniformity achieved through machine learning-enhanced production masks. Moreover, our use of machine learning has allowed us to bypass the time-consuming and labour-intensive process of optical chip testing, which significantly limits the scalability of photonic deployments in high-volume applications. As a powerful alternative to such testing, we developed a new technology that relies on a wafer probe that collects metrology data from multitude of locations on an undiced wafer. Utilizing a support vector machine (SVM), we analyze this metrology data and employ nonlinear binary classification to accurately predict the performance of hundreds of chips on a wafer across various metrics. We describe the approach employed for data collection to train the model, the trade-offs involved in hyperparameter tuning, and our methodology for evaluating the predictive quality of the binary classifiers. Additionally, we highlight the new capability of in-situ monitoring of wafer fabrication, which enables high-volume production and deployment of photonic solutions.

Keywords: machine learning, artificial intelligence, deep learning, AI/ML, SVM, photonic integrated circuits, integrated optics, fabrication

1. INTRODUCTION

A key driving force in the remarkable surge in information exchange over the past few decades has been the widespread adoption of optical communication systems.[1] Optical fibers enable long-distance communication with bandwidths unachievable by any other technology, while integrated optical devices allow the development of optical networks with advanced routing and multiplexing features.[2] The next frontier for photonic integrated circuits lies in transforming short-reach links, promising advanced optical interconnect solutions with low latency and low power requirements. However, despite their dominance in long-distance communication, the intricacies and cost-effectiveness of implementing optical solutions for shorter range links remain challenging,[3] impeding their widespread adoption.

In this paper, we present our progress in leveraging the power of machine learning to overcome the biggest hurdles for the widespread adoption of photonic integrated circuits. We describe our use of multi-path neural networks as a key tool for transitioning from low-volume prototype designs to consistently high-performing chips in volume production. We also present how we utilize machine learning to predict the performance of optical devices, thus eliminating chip testing, a time-consuming and labor-intensive process that is the central bottleneck in the production of photonic chips. These approaches not only accelerate the adoption of photonic solutions in existing domains but also facilitate their expansion into emerging high-volume applications such as LiDAR and optical computing.

*ksenia.yadav@enablence.com

2. MULTI-PATH DEEP LEARNING FOR DESIGN OPTIMIZATIONS

The transition from prototype design to achieving uniform manufacturing in high volumes is a pivotal phase in any product development process. In photonics, a standard production mask contains hundreds of photonic devices. Figure 1(a) illustrates an example of a production mask featuring 600 four-channel multiplexer devices, while Figure 1(b) shows the typical transmission spectra anticipated from each device on the wafer. The challenge in reproducing a successful prototype in a volume production environment lies in the fact that even minute variations in the physical parameters within the wafer can lead to substantial performance differences among identically designed devices. Such physical variations remain even after stringent process control measures and are often too subtle to detect by process metrology instruments, and thus are challenging or impossible to rectify.

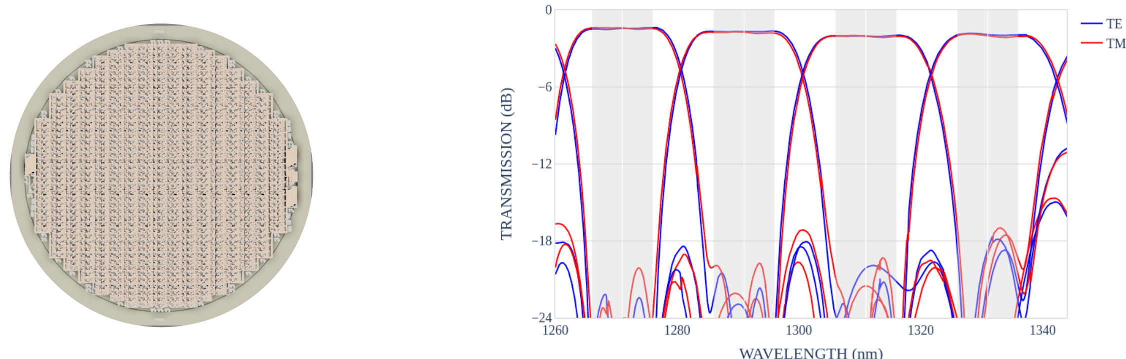

Figure 1. (a) A production mask with 600 four-channel multiplexer devices for CWDM applications. (b) Typical transmission spectra expected from each of the multiplexer devices, for TE and TM polarized light.

Regardless of the particular fabrication process, manufacturing variability presents a substantial challenge in achieving reproducible performance of photonic chips.[4] While traditional statistical methods can help compensate for some systematic non-uniformities, their impact is limited since they can only address a few known dependencies. In contrast, machine learning offers an entirely new approach to address this challenge due to its ability to handle high-dimensional data and identify hidden correlative patterns.

To solve this problem using deep learning, and to do so with reasonable resource constraints, we transitioned away from employing generic, fully interconnected neural networks. Instead, our focus shifted towards creating a task-specific neural network with custom layer definitions tailored to address our particular problem. Our solution involved the development of a custom multi-path neural network with a significantly reduced number of trainable parameters, allowing it to be trained with a smaller training dataset.

In a multi-path neural network, each pathway processes input data independently, learning distinct features or representations. The outputs from these pathways are then aggregated to make a final prediction. Each pathway can be architecturally different, designed to focus on a specific subtask within the overall problem. In our case, the multiplexer is implemented as a binary tree of cascaded lattice filters,[5,6] so each path corresponds to a particular wavelength splitter as indicated by the colors in Figure 2. This multi-path approach not only speeds up training but also enables effective training on a relatively small dataset, utilizing fewer computational resources without compromising predictive power.

The multi-path neural network is trained to study variations in the performance of individual chips based on a synthetic training dataset. Each training instance $\mathbf{x}^{(i)}$ includes a normalized four-channel spectrum obtained through simulation and a label that consists of 17 design parameters that are used to construct the studied multiplexer device, as shown in Figure 2(c). The training dataset contains 900,000 training instances, and we employ a standard iterative approach to train the model. Once trained, the model is utilized to infer the vector of actual design parameters $\hat{\mathbf{d}}^{(i)}$ from a real measured spectrum of a fabricated multiplexer chip. The difference between these as-fabricated design parameters and the intended design parameters, $\hat{\mathbf{d}}^{(i)} - \mathbf{d}^{(i)}$, is captured for each chip on a wafer, producing a multi-dimensional variability map that represents the true variability of the design parameters for a specific fabrication process. Once the variability map is established, we use it to compensate for systematic process variations in a new, optimized production mask. In this machine learning-enhanced version of the mask, the devices are no longer identical but vary as dictated by the neural network model.

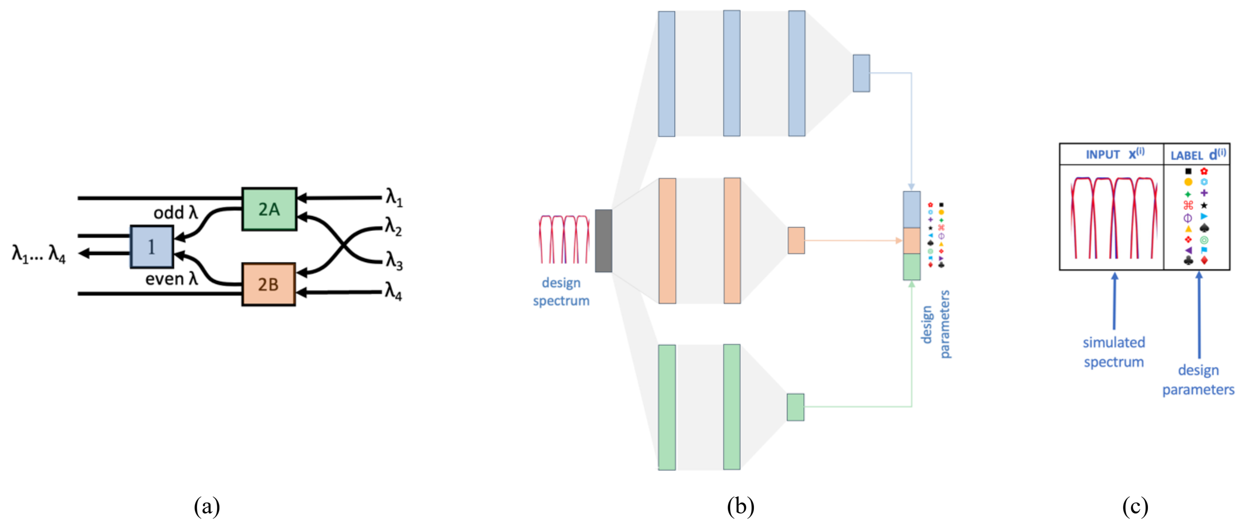

Figure 2. (a) A schematic of a four-channel multiplexer based on a two-stage binary tree of filters. (b) A multi-path neural network developed to infer the design parameters of a multiplexer from the provided spectra. (c) A training instance $\mathbf{x}^{(i)}$ that consists of a normalized spectrum and a label $\mathbf{d}^{(i)}$ with 17 design parameters.

To validate the approach and assess its impact on achieving improved uniformity of performance across the wafer, we applied the steps above to the production mask shown in Figure 1(a). Despite employing traditional statistical methods to correct for systematic process variabilities, the initial mask's devices exhibited large performance variations. While the typical performance was previously shown in Figure 1(b), the extensive variations are evident when we overlay the 20 worst-performing chips on the mask, as depicted in Figure 3(a).

In contrast, with the machine learning-enhanced version of the mask, an exceptional uniformity of performance was achieved, as shown in Figure 3(b). By adopting this strategy, we have gained the capability to fine-tune the various design parameters, ensuring consistent performance across hundreds of devices produced on a single wafer, despite imperfect control over the fabrication process.

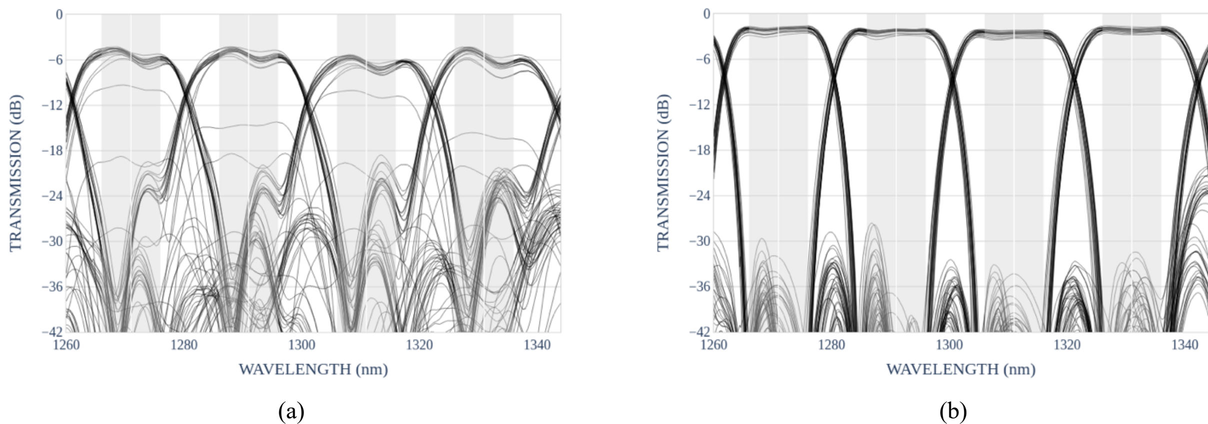

Figure 3. (a) Overlaid transmission spectra for 20 worst-performing chips in a production mask with about 600 devices. (b) Overlaid transmission spectra for the same 20 chips after a multi-path neural network was used to optimize the design parameters and create a machine learning-enhanced version of the production mask.

3. SUPPORT VECTOR MACHINE FOR PREDICTION OF OPTICAL PERFORMANCE

The preceding section described how deep neural networks can be used to reduce non-uniformity of performance across many devices. However, if the aim is to accelerate the adoption of photonics, it is important to realize that the primary bottleneck in photonics chip production is chip testing. Traditionally, the photonics industry relies on 100% chip testing – a labor-intensive and time-consuming process involving manual handling and precise alignment of individual chips that is not cost-effective for high-volume production. Even with experienced personnel and some automation, several minutes are required to test a single chip, and multiple days are necessary to fully characterize a single production wafer.

In our pursuit of a more efficient solution, we developed an automatic probing technique for undiced wafers. The probe measurement collects spectroscopic signatures from 64 locations on the wafer, as shown in Figure 4(a). The primary challenge of this technique lies in the fact that the signal collected by the wafer probe is weak and noisy (Figure 4(b)), yet it must be relied on to accurately predict the performance of hundreds of chips on a wafer based on stringent optical specifications.

Recognizing that finding correlations within weak and noisy signals is one of the primary strengths of machine learning, we developed a support vector machine (SVM) that performs nonlinear binary classification (pass / fail) based on the automatic probe measurement. The SVM's prediction for a particular wafer is shown in Figure 4(c), while the actual performance of the same wafer – obtained after dicing the wafer and characterizing each of the chips by traditional chip testing – is shown in Figure 4(d). We use the receiver operating characteristic (ROC) curve to evaluate the performance of the binary classification algorithm, and the area under the curve (ROC AUC) for the selection and fine-tuning of hyperparameter configurations. An incremental learning algorithm enables us to improve prediction accuracy as wafer production continues and more training data becomes available. Once a product enters the volume production stage, the predictive capabilities of the SVM become robust enough to reduce chip testing to only a small fraction of the fabricated devices, effectively alleviating the primary bottleneck of the production process.

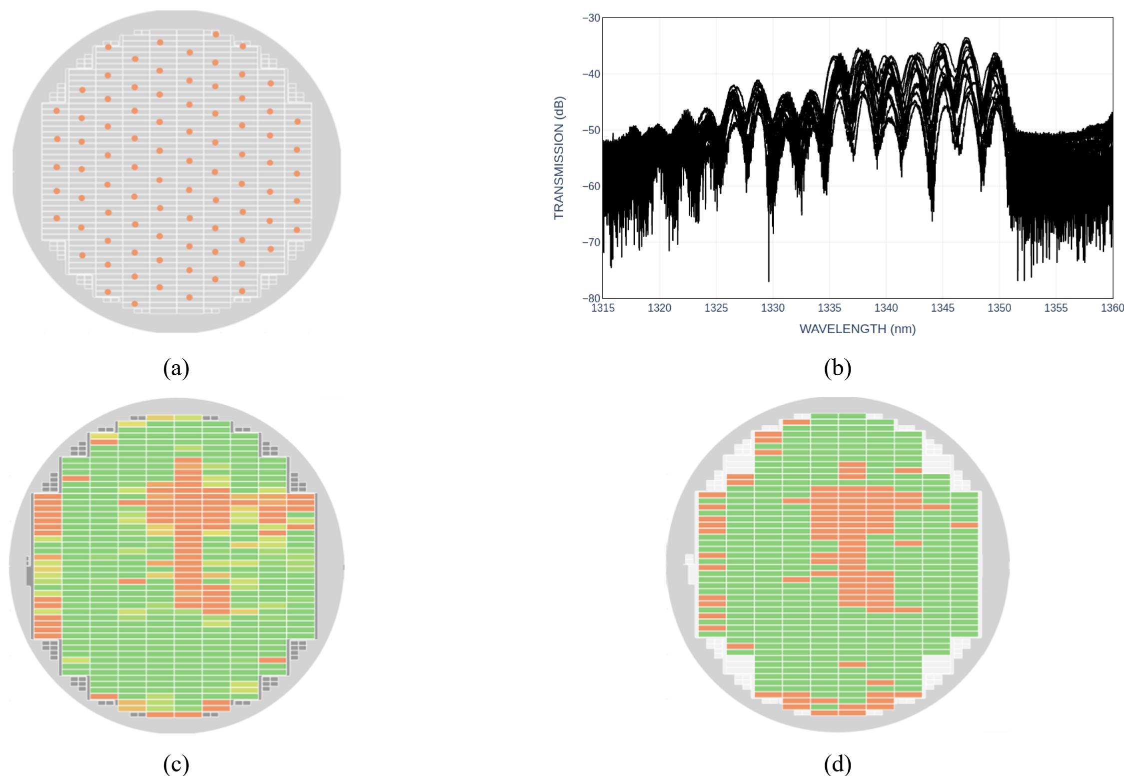

Figure 4. (a) Metrology data is obtained from an automatic probe measurement from 64 locations around the wafer. (b) Typical spectroscopic signature obtained from the wafer probe measurement. (a) SVM prediction of the wafer performance based on the probe measurement (green: predicted pass, orange: predicted fail, colors in between indicate the degree of confidence of the prediction). (b) Actual wafer performance after traditional optical chip testing (green: pass, orange: fail).

Relying on an SVM to predict the performance of entire wafers based only on a wafer probe measurement revolutionizes the manufacturing of photonic chips for two compelling reasons. The first reason is economic. Obtaining the pass / fail map through traditional chip testing, as shown in Figure 4(d), is a multi-day effort, whereas the total time required to mount, align, obtain the spectroscopic signatures and infer the predicted map, as shown in Figure 4(c), is only 12 minutes. The second reason stems from a fundamentally new capability introduced by this approach: as the wafer remains intact and the probe measurement requires minimal effort, the probe becomes an invaluable in-situ monitoring tool for wafer fabrication. In contrast to traditional fabrication, where chip performance evaluation occurs only after the process is considered final, our approach facilitates continuous monitoring of the wafer's state at crucial stages throughout the fabrication process. This capability enables real-time adjustments to optimize the performance of entire wafers, leading to consistently high-performing devices at high production volumes.

4. CONCLUSIONS

In this paper, we described how we leveraged the power of machine learning to facilitate high-volume manufacturing of photonics chips. Employing a custom multi-path deep neural network, we fine-tuned the design parameters of devices on a mask, compensating for inevitable process variations and achieving unprecedented uniformity of performance across hundreds of devices on a wafer. Additionally, a support vector machine was trained to predict the multi-dimensional performance of photonic chips on a wafer based on weak spectroscopical signatures obtained from an automatic wafer probe measurement. This approach not only eliminates the need for the labour-intensive process of optical chip testing but also allows in-situ monitoring of wafers and real-time process adjustments. We demonstrate that the use of machine learning in the design and the manufacturing of photonic chips in high-volume production is critical for accelerating the adoption of photonic solutions across various emerging applications.

REFERENCES

[1] Doerr, C., "Silicon photonic integration in telecommunications," Frontiers in Physics, vol. 3, p. 37 (2015).

[2] Cheben, P., Halir, R., Schmid, J. H., Atwater, H. A., Smit, D. R., "Subwavelength integrated photonics," Nature, vol. 560, p. 565 (2018).

[3] McMahon, P. L., "The physics of optical computing", Nature Reviews Physics, vol. 5, p. 717 (2023)

[4] Lu, Z., Jhoja, J., Klein, J., Wang, X., Liu, A., Flueckiger, J., Pond, J., Chrostowski, L., "Performance prediction for silicon photonics integrated circuits with layout-dependent correlated manufacturing variability," Optics Express, vol. 25, no.9, p. 9712 (2017).

[5] Horst, F., Green, W. M. J., Assefa, S., Shank, S. M., Vlasov, Y. A., and Jan Offrein, B., "Cascaded Mach-Zehnder wavelength filters in silicon photonics for low loss and flat pass-band WDM (de-) multiplexing," Opt. Express, vol. 21, no. 10, pp. 11652–11658 (2013).

[6] Bidnyk, S., Yadav, K., Balakrishnan, A., "Synthesis of ultra-dense interferometric chains in planar lightwave circuits," in Proc. SPIE 12004, Integrated Optics: Devices, Materials, and Technologies XXVI (2022).