

THE NOBLE ART OF DEFECT DETECTION

In the rapidly evolving field of photonics, the integrity of photonic wafers is paramount to the performance and reliability of optical devices. These devices form the backbone of technologies ranging from high-speed communication systems to cutting-edge sensors. Defect inspection on photonic wafers involves sophisticated techniques and technologies designed to identify and mitigate imperfections that could compromise device functionality. This article explores the various techniques used in defect inspection, the challenges encountered, and the latest advancements that are enhancing the accuracy and efficiency of these processes.

YANICK DOUVEN, RIK WETZELS, STEFAN VAN DER PALEN, RINUS HOOGESTEGER AND WOUTER HOUTMAN

Introduction

The rapid growth of the photonic integrated circuit (PIC) market has introduced new challenges to the manufacturing process, particularly in the inspection and quality control of photonic wafers. Photonic chips are built on different substrates compared to traditional chips and feature unique geometries, such as the absence of sharp 90° angles, as light cannot bend at such angles. These differences render established inspection methods ineffective for photonic wafers.

The absence of established libraries of standardised components forces manufacturers to experiment with both the design and manufacturing processes, leading to a complex interplay between the two. This experimentation often results in yield issues and forces smaller companies to spend significant time on testing the products they create. The time required to test a new design, and the absence of tested and proven designs make it difficult for such companies to scale their operations efficiently.

To address these issues, advanced surface-inspection solutions have been developed within the MEKOPP [1] NXTGEN [2] project, combining cutting-edge hardware, sophisticated machine learning algorithms, and innovative software frameworks. From cost-effective offline solutions that support early-stage product development to high-end inline systems for large-scale production, these tools provide a comprehensive approach to defect detection and are designed to meet the unique demands of photonic wafer inspection, enabling manufacturers to overcome the limitations of traditional methods.

Nobleo has played an active role in the development of a photonic wafer inspection system, contributing to two key areas:

1. Developing camera-control software to enable high-speed surface scanning of photonic wafers.
2. Implementing machine-learning-based approaches to detect defects on photonic wafers.

This article provides an overview of the technical challenges and shares the most important insights gained throughout the development process.



The Helios system, a visual inspection tool for performing defect inspection on photonic integrated circuits.

AUTHORS' NOTE

Yanick Douven, Rik Wetzels, Stefan van der Palen, Rinus Hoogesteger and Wouter Houtman work at the Nobleo Autonomous and Intelligent Solutions group of Nobleo Technology in Eindhoven (NL), forming a mixed team of vision, automation, mechatronic and AI experts.

yanick.douven@nobleo.nl
www.nobleo-technology.nl

About Nobleo

Nobleo is an Eindhoven-based engineering company with 100+ employees and three business units:

1. Autonomous and Intelligent Solutions
2. Intelligent Embedded Solutions
3. High Tech Systems

The Nobleo Autonomous and Intelligent Solutions business unit is involved in developing machine-learning solutions for high-tech systems and autonomous robotics.

The strategic focuses are:

- 3D object tracking (for outdoor applications):
 - multi-modal sensor platform;
 - AI object detection;
 - sensor fusion (stereo cameras, LiDAR, encoders).
- 3D image processing (for outdoor applications):
 - sensor design & integration;
 - edge processing;
 - integration with motion control.
- Surface inspection (for high-mix, low-volume production):
 - data annotation;
 - AI anomaly detection;
 - AI object detection.
- Precision alignment (for high-tech systems):
 - sensor selection & calibration;
 - feature-matching networks;
 - system integration & calibration.

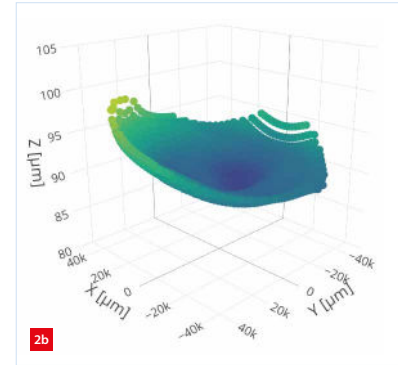
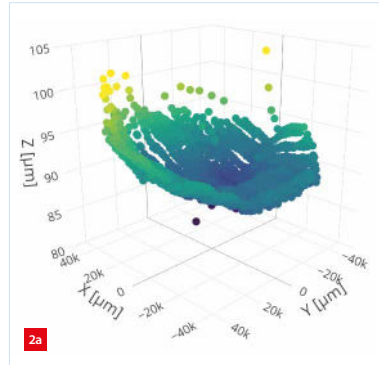
It is Nobleo's ambition to create machine-learning solutions by leveraging know-how of mathematics, physics, optics, electronics and mechatronics to come up with innovative solutions for their customers.

Helios photonic wafer-inspection tool

The development of wafer-defect-inspection solutions for the photonics industry has brought together expertise of Nobleo, Integrated Mechanization Solutions (IMS), Workfloor, and Settels Savenije. An outcome of this collaboration is the Helios system (Figure 1), a photonic visual inspection tool, which works alongside the Sirius active probing system developed by IMS [3]. These tools address the stringent requirements of defect detection in photonic integrated circuits (PICs). LioniX International is actively involved as a representative of the photonics industry.



Key system features developed by Nobleo include:

1. Scanning areas of interest (AOI):
 - Targeting specific regions of the wafer for inspection.



Height profile of a full wafer. The residual RMS error of $1\text{ }\mu\text{m}$ is smaller than the $2\text{ }\mu\text{m}$ depth of field of the camera.

- (a) Height measurements.
(b) Resulting height map.



The output of the inspection tool.

- (a) A heatmap showing anomalous regions of the full wafer can be created to highlight polluted areas.
(b) Detected defects are also overlaid on top of the design.

2. Wafer alignment:

Ensuring accurate knowledge of the position of the wafer within the machine.

3. Height mapping:

Creating a height profile to guide autofocus adjustments.

4. Defect detection:


Applying algorithms to identify surface imperfections.

5. Defect reporting:

Documenting defect locations in wafer coordinates for further analysis.

Helios is built for efficiency and precision, inspecting up to ten 4-inch wafers per hour. It supports resolutions of $0.125\text{ }\mu\text{m}$ and $0.5\text{ }\mu\text{m}$ per pixel, enabling the detection of small defects. Brightfield and darkfield illumination modes ensure adaptability to various inspection needs.

Users can customise test recipes by defining AOIs, inspection types, illumination settings, resolution and alignment markers. The system generates detailed test reports, including defect positions and types, providing essential data for quality control and process optimisation.



Defect data is organised by creating maps linked to high-resolution scan images.

Benefits

The photonic visual inspection tool Helios offers several advantages that enhance the efficiency and accuracy of photonic wafer-inspection processes:

- **Reduced inspection time and improved consistency:**
Helios streamlines the inspection process, ensuring faster and more reliable results.
- **Detailed defect visualisation:**
Wafer defects as small as 100 nm can be visualised directly on top of the design file (GDSII format, GDS = Graphic Design System), significantly reducing the time required to determine whether a wafer is acceptable (OK) or not acceptable (NOK).
- **Web-based annotation:**
Process experts can collaborate remotely to establish a baseline definition of defects, leveraging a web-based annotation platform.
- **Rapid anomaly detection:**
An anomaly detection network can be developed for each wafer design within just 0.5 hour, enabling quick adaptation to new designs.
- **Custom defect classification:**
Machine-learning tools are used to co-develop tailored defect-classification algorithms, improving the accuracy of defect identification.
- **Robust data storage:**
A 64-TB data storage solution provides immediate access to wafer data files, supporting efficient analysis and traceability.
- **High throughput:**
The system can inspect up to ten wafers per hour, making it suitable for industrial-scale operations.

Height mapping

One of the primary challenges in this field is scanning photonic wafers with high-magnification objectives. The limited depth of field of the imaging system compared to the local height variations on the wafer poses a significant

obstacle. To maintain image clarity during scanning, the microscope objective, designed by Settels Savenije, incorporates a high-speed autofocus sensor. To filter out height-measurement errors on sharp edge transitions, the sensor is used to measure the wafer topography, generating a detailed height profile of the wafer; see Figure 2. This enables the autofocus system to dynamically adjust during scanning.

As a result, the system achieves high-speed scanning without sacrificing image quality. This allows the user to scan surfaces with a resolution of up to 125 nm per pixel, producing the detailed images necessary for defect detection.

Defect maps

The inspection tool leverages state-of-the-art AI algorithms to generate detailed defect maps; see Figure 3. These maps provide critical insights into wafer quality by pinpointing the position and type of detected defects, and are seamlessly integrated with a Manufacturing Execution System (MES), enabling further analysis and streamlining the inspection process.

For high-mix, low-volume production, where visual inspection is a common practice, these maps significantly reduce the time required to identify issues. By cross-referencing defect positions with the wafer design file, non-critical defects can be filtered out, allowing engineers to focus on the most relevant problems. Additionally, the maps facilitate root-cause analysis by enabling the comparison of defects across consecutive processing steps, helping to identify and address underlying issues in the production process.

Beyond defect detection, the tool supports process quality evaluation by linking defect statistics with process parameters, offering valuable insights for optimising manufacturing workflows. Engineers can use the defect

maps to assess wafer-level issues quickly, while high-resolution images from the scans can be loaded for a more detailed analysis when needed; see Figure 4. This combination of defect mapping and in-depth imaging ensures a more efficient and accurate approach to quality control in photonic wafer production.

AI-driven defect detection: a two-phase approach

The integration of artificial intelligence (AI) into defect detection for photonic wafers follows a carefully structured two-phase approach, designed to address the challenges of both early development and large-scale production.


In the initial phase, the focus is on supporting early-stage

development through an offline scanning solution. During this stage, algorithms requiring minimal annotation effort, such as autoencoders, are employed. These algorithms are designed to deliver results quickly, even if they are not yet perfect or fully optimised. The primary goal of this phase is to gain a deeper understanding of the defect-detection problem, build an initial dataset, and provide support for visual inspection. This phase is particularly valuable for identifying patterns and laying the groundwork for more advanced solutions, enabling teams to iterate rapidly and refine their approach.

As the system transitions to the scale-up phase, the focus shifts to inline scanning solutions that are fully integrated into the production line. This phase leverages the insights and datasets developed during the first phase, but it requires a larger and more thoroughly annotated dataset to achieve the desired level of accuracy and autonomy. By building on the foundation established earlier, the system evolves into a fully autonomous solution capable of detecting defects in real time during production. This ensures that the inspection process is not only faster but also more reliable, meeting the demands of industrial-scale operations.

Together, these two phases create a robust pathway for integrating AI into defect detection, starting with exploratory development and culminating in a mature, production-ready solution. This approach ensures that the system is both adaptable and scalable, addressing the needs of the photonics industry at every stage of the production lifecycle.

NoblAI



Nobleo's proprietary machine-learning framework, NoblAI.

Nobleo has developed its own machine-learning framework, NoblAI, designed to help customers create and maintain machine-learning networks with ease; see Figure 5. The framework simplifies the process of managing users, creating projects, and adding datasets. It also supports the integration of various machine-learning networks, such as object detection, anomaly detection, and segmentation, into these projects. Once the networks are added, users can train, evaluate and deploy them seamlessly.

To streamline these tasks, NoblAI relies on blueprints: predefined sets of parameters that cover the majority of use cases. These blueprints allow most tasks to be executed without requiring parameter adjustments, though customers can modify them when fine-tuning is necessary. This flexibility ensures that the framework can adapt to a wide range of applications while maintaining ease of use.

One of the standout features of NoblAI is its ability to reduce the annotation effort required to train machine-learning networks. By sorting images in a dataset based on anomaly scores, the framework prioritises the most critical images for annotation. Images with the highest anomaly scores are annotated first, ensuring that the most relevant data is addressed early in the training process.

Additionally, the framework's anomaly-detection networks can suggest bounding boxes for object-detection networks, allowing annotators to focus on refining these suggestions rather than starting from scratch. This approach not only saves time but also improves the consistency of the annotation process.

Defect-detection algorithms

The defect-detection process for photonic wafers involves handling images with extremely high resolution, each measuring 4,504x4,504 pixels. These images are too large for standard object-detection networks to process directly, so they are divided into smaller patches. The size of these patches is carefully chosen based on the specific requirements of the detection network, ensuring an optimal balance between computational efficiency and detection accuracy.

To address the diverse needs of inline and offline defect detection, multiple machine-learning networks are employed, each optimised for a specific purpose.

Anomaly detection

The first approach is anomaly detection [4], which identifies regions on the wafer that deviate from the expected norm; more specifically, unsupervised anomaly detection using an autoencoder network. Autoencoders learn a compact representation of the input data. Anomalies can be detected by comparing the reconstruction error (difference between input and output) of the autoencoder. When the model encounters data that significantly deviates from the learned patterns, the reconstruction error tends to be high, indicating an anomaly.

This method generates heatmaps with probability scores for each pixel, highlighting potential defects. Anomaly detection is highly efficient, processing a 20-megapixel image in approximately 320 ms (on an Nvidia RTX A4500 GPU). One of its key advantages is its minimal reliance on annotated data. These networks can deliver meaningful results even without annotations (unsupervised), making them particularly useful in early development stages or when labelled datasets are limited. Additionally, they are

tolerant of small annotation errors, further reducing the burden on human annotators. The output includes heatmaps and overview maps that highlight potential defect hotspots, providing a foundation for further analysis.

Object detection

The second approach is object detection, which classifies and localises defects with near-pixel-perfect accuracy. Object-detection networks require more extensive and annotated datasets, as they are trained to recognise specific defect types. These networks process a 20-megapixel image in approximately 250 ms. By focusing on precise defect localisation, object detection provides detailed insights into the nature and position of defects, making it ideal for scenarios where accuracy is critical.

Classification per patch

Another method is classification per patch, which analyses small patches of the image, typically 64x64 or 96x96 pixels, and assigns a single classification label to each patch. While this approach offers limited spatial accuracy, it is extremely fast, processing a 20-megapixel image in just 100 ms. This makes it suitable for quickly identifying defects in scenarios where precise localisation is not required.

Combined network

Finally, a combined network approach integrates patch classification with anomaly detection. This hybrid method first uses patch classification to quickly and roughly identify potential defects. Anomaly detection is then applied only to the flagged patches, providing pixel-perfect results for the most critical areas. This combined approach balances speed and precision, ensuring efficient and accurate defect detection.




Outputs of the different algorithms.

(a) The anomaly detector produces a heatmap.

(b) The object detector produces accurate location and classification information for each defect.

(c) The combined network roughly identifies potential defects, after which the anomaly detector is run for all flagged areas.



Filtered matches of the fiducial template on a scanned image (left), and the template overlaid on the scan (right).

The choice of network depends on the specific requirements of the inspection process, such as the need for speed, accuracy or detailed localisation; see Figure 6.

The success of these algorithms depends on the quality and quantity of the training data. Anomaly-detection networks require relatively little annotated data, as they focus on learning the characteristics of defect-free regions.

In contrast, object-detection networks demand larger, well-annotated datasets to achieve high accuracy. To streamline the annotation process, tools like NoblAI (see the text box) prioritise images with the highest anomaly scores, ensuring that annotators focus on the most critical data. NoblAI can also suggest bounding boxes for defects, allowing annotators to refine these suggestions rather than starting from scratch, significantly reducing the time and effort required.

Wafer alignment

Accurate wafer alignment is a critical step in the inspection process, ensuring that images captured during scanning are precisely linked to the wafer's design. While rough alignment is typically known beforehand, fine alignment requires a more sophisticated approach to achieve the necessary precision. This process begins by capturing images of fiducials – specific alignment markers on the wafer – and using them to align the images to the design.

To localise these alignment markers, the system employs a state-of-the-art feature-matching neural network. The vector description of an alignment marker is added to the test recipe and a template of the marker is rendered. The system then scans the wafer to locate the marker. The marker template and scan image are used as input image pair for the neural network, calculating a set of corresponding keypoints (matches) between the two images. These keypoints represent the locations where parts of the alignment marker are detected; see Figure 7.

Similar to the defect-detection algorithms described earlier, the high-resolution input images used in the alignment

process are too large to process directly. Therefore, the same patching approach is applied here. The neural network is run on each patch, generating matches for the alignment marker. These matches are then combined and filtered using a RANSAC [5] algorithm to determine the final position of the marker. This approach ensures that the alignment process is both robust and accurate.

The localisation algorithm achieves a position estimate with an accuracy of ± 1 pixel, corresponding to a physical accuracy of 0.125 to 0.500 μm , depending on the camera resolution used. By combining the detected location of multiple fiducials spread across the wafer, the transformation between each image and the design can be calculated, allowing the system to overlay detected defects on the design file for further analysis.

Defect-inspection results

The added value of the defect-inspection tooling depends heavily on the performance. In close collaboration with photonic foundries, such as LioniX and a recent Swiss customer, the following three use cases have been defined, for which the algorithm performance is being validated (Figure 8).

Inspection on component level (use case 1)

The first use case relates to the capability of the software to detect process variations. Manufacturers of high-mix, low-volume photonic assemblies spend significant time testing their products electrically and optically. Algorithm-validation activities are planned to perform inspections on component level. The objective of this track is to find a correlation between functional behaviour and component assessment with machine learning. When found, this will allow photonic foundries to reduce effort in active testing, which saves time and money from critical resources.

Critical defect detection (use case 2)

The second use case focuses on the detection of critical defects. Defects on light-tunnelling structures are much more function-critical than a particle on a non-functional

area of the wafer. The performance of the algorithms is evaluated for defect maps generated with the anomaly detector when filtered with design information.

The software outputs defect maps for each processing step that can be overlayed with the design file (available to the manufacturer only). This allows quality and production engineers to assess critical defects only. By analysing the series of maps for consecutive processing steps, an understanding of the origin of the defect in the more than 50 processing steps can be created.

Categorise defects (use case 3)

The third and final use case classifies the detected defects into a customer-specific defect category. Photonic foundries have an extensive set of labels for different defect types linking to the different processing steps. Training an object detector for such a set requires sufficient examples for each defect and substrate type. This is a practical challenge for smaller foundries. Therefore, the annotation time that customers put in should be time-boxed to balance benefits and cost. Object-detection networks are rolled out starting with course defect categories and making them more and more granular over time as the annotated dataset grows. This allows for gradual improvement and granularity of the categorisation.

Conclusion

The development of surface-inspection solutions for photonic wafers is tailored to meet the specific needs of different stages in the product lifecycle, from early development to high-volume production. By offering both cost-effective offline solutions and high-end inline systems, the inspection process is adaptable, scalable and efficient.


During the early phases of product development, a flexible offline scanning solution is employed. This approach is cost-effective and supports visual inspection by providing detailed scans of produced parts. The algorithms used in this phase can identify at least the most critical defects, ensuring that early-stage issues can be addressed promptly. Additionally, scanning a variety of wafers during this phase helps build an annotated dataset, laying the groundwork for more advanced machine-learning models. This offline solution has already garnered enthusiasm from lead customers, highlighting its value in supporting early-stage development.

As production scales up, the focus shifts to a high-end, custom inline scanning solution tailored for high-volume manufacturing. This system performs real-time inspection of produced parts, applying validated inspection algorithms to ensure consistent quality. The inline solution is designed to integrate seamlessly into production lines, providing a robust hardware and software platform.

By addressing the unique requirements of each phase, these surface-inspection solutions ensure that photonic wafers are inspected with precision and efficiency, supporting both innovation and scalability in the photonics industry.

REFERENCES

- [1] "MEKOPP – Metrology Equipment for critical scale up of PIC Production", www.hightechnl.nl/mekopp/projecten/mekopp (accessed 7-4-2025).
- [2] "NXTGEN Hightech domain Semiconductors", High Tech NL Semiconductors, online, www.hightechnl.nl/nxtgen-hightech/projecten/nxtgen-hightech-domain-semiconductors (accessed 7-4-2025).
- [3] R. Kelder, "PIC testing on wafer level", *Mikroniek*, vol. 65 (1), pp. 19-25, 2025.
- [4] V. Kumar, "Anomaly Detection Techniques: A Comprehensive Guide with Supervised and Unsupervised Learning", *Medium*, 2023, online www.medium.com/@venujkvenk/anomaly-detection-techniques-a-comprehensive-guide-with-supervised-and-unsupervised-learning-67671cdc9680 (accessed 27-3-2025).
- [5] M.A. Fischler, and R.C. Bolles, "Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography", *Comm. ACM*, vol. 24 (6), pp. 381-395, 1981.



Use cases for validation of algorithm performance.

- (a) Inspect on component level:
Train an anomaly detector on a specific component to monitor process quality and/or component-related defects.
- (b) Detect critical defects from design info:
Train an anomaly-detection network on dies and AIOs in dies of a wafer and filter the defect map in a post-processing step using design information.
- (c) Categorise defects:
Train an object-detection network to classify defects into the right category.