

Article

Not peer-reviewed version

Deep CNN-Based Layout-to-SEM Image Reconstruction with Conformal Uncertainty Calibration for Nanoimprint Lithography in Semiconductor Manufacturing

[Jean Chien](#) and [Eric Lee](#) *

Posted Date: 26 June 2025

doi: 10.20944/preprints202506.2194.v1

Keywords: nanoimprint lithography (NIL); layout-to-SEM reconstruction; U-Net; conformal prediction; uncertainty quantification; optical proximity correction (OPC)



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a Creative Commons CC BY 4.0 license, which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Article

Deep CNN-Based Layout-to-SEM Image Reconstruction with Conformal Uncertainty Calibration for Nanoimprint Lithography in Semiconductor Manufacturing

Jean Chien and Eric Lee *

Department of Chemical Engineering, National Taiwan University

* Correspondence: ericlee@ntu.edu.tw

Abstract

Nanoimprint lithography (NIL) has emerged as a promising sub-10 nm patterning at low cost, yet robust process control remains difficult because of time-consuming physics-based simulators and labeled SEM data scarcity. We propose a data-efficient, two-stage deep learning framework here that directly reconstructs post-imprint SEM images from binary design layouts and delivers calibrated pixel-by-pixel uncertainty simultaneously. First, a shallow U-Net is trained on conformalized quantile regression (CQR) to output 90 % prediction intervals with statistically guaranteed coverage. Moreover, per-level errors on a small calibration dataset are designed to drive an outlier-weighted and encoder-frozen transfer fine-tuning phase that refines only the decoder, with its capacity explicitly focused on regions of spatial uncertainty. On independent test layouts, our proposed fine-tuned model significantly reduces the mean absolute error (MAE) from 0.0365 to 0.0255 and raises coverage from 0.904 to 0.926, while cutting labeled data and GPU time by 80 % and 72 %, respectively. The resultant uncertainty maps highlight spatial regions associated with error hotspots and support defect-aware optical proximity correction (OPC) with fewer guard-band iterations. Extending the current perspective beyond OPC, the innovatively model-agnostic and modular design of the pipeline here allows flexible integration into other critical stages of the semiconductor manufacturing workflow, such as imprinting, etching, and inspection. In these stages, such predictions are critical to achieving higher precision, efficiency, and overall process robustness in semiconductor manufacturing, which is the ultimate motivation of this study.

Keywords: nanoimprint lithography (NIL); layout-to-SEM reconstruction; U-Net; conformal prediction; uncertainty quantification; optical proximity correction (OPC)

1. Introduction

1.1. Nanoimprint Lithography: Principles, Advantages, and Manufacturing Challenges

Nanoimprint lithography (NIL) has emerged as an inspiring replacement for conventional photolithography recently, with a low-cost and high-resolution patterning process for sub-10 nm nodes [1–4]. The simplicity of its process flow—comprising master template creation and replica template fabrication, final wafer patterning—makes it particularly suitable for high-throughput manufacturing [5]. In contrast to conventional optical lithography, which relies on the use of reduction masks (typically 4× or 5×), NIL makes use of a 1:1 imprint template to pattern the substrate directly at the nanoscale. This one-to-one transfer mechanism reduces reliance on traditional optical proximity correction (OPC), which compensates for distortions introduced by optical systems. While OPC plays a vital role in ensuring that the pattern on the wafer is identical to the design layout, its

models are based on the physics of optical lithography. NIL, on the other hand, relies on entirely different mechanisms such as mechanical contact and polymer flow, which calls for new OPC frameworks specific to imprint-based techniques. To ensure pattern fidelity, OPC often employs guard bands—conservative layout margins that require costly iterative tuning, which NIL's distinct physics now demand to be redefined. The strategic importance of NIL has been increasingly recognized by leading semiconductor foundries in Asia, where the technology has already been introduced into demo-line systems for next-generation patterning evaluation.

In spite of all its advantages, NIL still faces challenges in several aspects of process control, including the trade-off involved in optimizing residual layer thickness and the non-uniform droplet distribution in near-shot-edge regions [6–8]. Recent studies suggest that NIL represents not only a quantitative advancement in resolution but also a state-of-the-art lithographic paradigm that simplifies process architecture, reduces overlay and proximity effects, and offers novel flexibility for 3D patterning and emerging device geometries [9]. Its revolutionary impact is more than keeping up with next-generation photolithography performance—it redefines the limits of the achievable resolution and process scalability in next-generation semiconductor manufacturing [10].

In the context of semiconductor chip-level fabrication, these limitations become particularly significant. Contemporary integrated circuits demand extremely high levels of precision and uniformity with tolerances at the nanometer scale. As feature sizes continue to shrink and device architectures grow increasingly complicated, manufacturing bottlenecks have developed into major hurdles to technological advancement. These issues cannot be resolved by traditional rule-based process control methodologies alone, especially under data-limited or variation-sensitive conditions [11].

1.2. AI and Deep Learning for Layout-to-SEM Reconstruction in NIL

To address this need, our recent research explored the application of artificial intelligence (AI) for semiconductor metrology and inspection, focusing on advanced defect detection based on scanning electron microscope (SEM) images taken from after-development inspection (ADI) [11,12]. In this work, we proposed a two-stage deep learning framework with synthetic minority over-sampling (SMOTE) and transfer fine-tuning to combat the significant issue of data imbalance and classification challenge efficiently. This strategy not only enhanced defect detection accuracy and computational efficiency for SEM-based hotspot monitoring but also highlighted data scarcity as a key challenge in achieving robust models for process control. Following the aforementioned study, our research here proposes a deep learning integration into NIL-based fabrication analysis, with a focus on layout-to-SEM image reconstruction enhanced by conformal uncertainty calibration. By coupling pixel-level prediction with statistically valid uncertainty quantification (UQ), the framework aims to bridge the gap between the limitations of physics-based simulations and the need for scalable, data-driven defect prediction in cutting-edge patterning processes.

To enable the widespread adoption of NIL in semiconductor manufacturing, sophisticated process optimization and tight control mechanisms are not only beneficial—they are imperative for modern semiconductor manufacturing, especially in NIL [13]. In this regard, the integration of AI into the manufacturing process has become strongly associated with possible future breakthroughs, especially due to its ability to extract latent patterns from complex process data, adjust for variability, and support closed-loop optimization strategies. Furthermore, as AI chips themselves begin to be a driving force for increasingly sophisticated fabrication, there exists a self-reinforcing cycle between AI and semiconductor process development—AI must assist in enabling what AI hardware demands.

However, for AI-driven optimization to be successful in semiconductor manufacturing processes, access to high-quality and high-resolution process data is indispensable. In practice, obtaining real-world process data, such as SEM images from resist development or etching stages, however, is time and resource-demanding [12]. For one thing, establishing accurate physical simulations that capture material-dependent properties is likewise nontrivial. [3,4]. These constraints

have motivated the use of machine learning (ML) as a promising alternative for accelerating process modeling and functional approximation [14,15]. Yet purely data-driven ML models tend to overfit under data scarcity and often lack calibrated mechanisms for quantifying predictive uncertainty—two major challenges in high-stakes semiconductor applications [16].

To reckon with this challenge, we present here a deep convolutional neural network (CNN) layout-to-SEM image reconstruction model enhanced with conformal uncertainty calibration for precise mapping of input design layout images to target SEM outputs. As it is hard to exactly quantify the trustworthiness of data-driven ML models, interpretability and explainability serve as effective proxies in practice. In this way, it is essential that the model remains not only accurate but also transparent, offering visualizable outputs and calibrated confidence estimates that support process optimization and defect-aware decision-making in semiconductor manufacturing. This approach meaningfully bridges the gap between sparse real-world data and the need for precise, explainable modeling of NIL processes [17,18].

Among various deep learning architectures, the CNN-based U-Net model has proven to be particularly effective for image-to-image prediction problems due to its encoder-decoder structure and skip connections. These connections allow for the precise localization of spatial features by transmitting high-resolution detail from encoder to decoder [11,19,20]. While originally proposed for biomedical image segmentation [21,22], the U-Net architectural characteristics have been widely applied to various image reconstruction, denoising, and super-resolution tasks [23–25]. The versatility lies in its capacity to preserve subtle structural detail through skip connections that maintain high-resolution spatial features from the early encoder layers to the decoder. These encoder features, together with those learned along the contracting path, form hierarchical representations that capture information ranging from low-level pixel intensities to medium-level contour structures, and up to high-level global layout semantics, enabling the reconstruction of intricate nanoscale patterns.

In the context of layout-to-SEM reconstruction, this architectural feature is crucial for accurately capturing the nanoscale contours of the design layout pattern vital to post-imprint topographies. Compared to conventional CNNs that may lose spatial resolution due to repeated downsampling, U-Net retains subtle edge information, making it particularly suitable to predict high-resolution details such as line edges, corners, and high-contrast boundaries in NIL, which are crucial for NIL defect analysis and overlay verification [26–28].

1.3. Need for Reliable Uncertainty Quantification

However, for such predictive models to be trustworthy in practice, they must not only provide accurate outputs but also quantify the uncertainty associated with each prediction, which enables a more reliable and explainable prediction framework. This is particularly crucial in semiconductor manufacturing, where uncalibrated outputs can lead to false confidence or overlooked defects in downstream OPC verification or yield optimization processes [16,29]. In ML-based process modeling, especially for high-stakes semiconductor applications, a UQ or uncertainty-aware training approach plays a critical role in the evaluation of predictive model reliability. UQ allows for identifying the regions of a model where errors are likely to occur and facilitates decision-making based on confidence levels [30–32]. In recent uncertainty-aware training approaches, models are encouraged to detect and highlight high-risk or uncertain regions as warning signals by leveraging UQ during training. For instance, Ding et al. (2020b) introduced a pixel-by-pixel reweighting mechanism based on model confidence, which effectively reduced overconfident errors in medical segmentation tasks [33].

Two major categories of uncertainty are typically recognized in this context: aleatoric uncertainty, which originates from the intrinsic noise or variability in the input data and is therefore irreducible; and epistemic uncertainty, on the other hand, which is associated with the incomplete learning caused by limited training data or insufficient model capacity, and can potentially be addressed through data augmentation methods or more capable modeling techniques [34]. Aleatoric

uncertainty corresponds to intrinsic stochastic noise—such as SEM imaging artifacts or environmental fluctuations—while epistemic uncertainty emerges from the model being unaware of unseen design patterns or process variations. Both uncertainties must be quantitatively assessed to allow for reliable downstream decision-making. These concepts have been made functional in some earlier work as well. Dawood et al. (2021) used prediction confidence intervals during the test phase to enhance the model’s confidence in correct predictions and to suppress overconfidence in incorrect ones as well. This, in turn, demonstrates the effectiveness of uncertainty-aware feedback during training and is aligned with recent efforts to calibrate model confidence, where prediction intervals serve as interpretable cues for reliability assessment [35].

1.4. Conformal Prediction and Conformalized Quantile Regression

Several UQ methods have been discussed in the literature, such as Bayesian inference[36], bootstrapping[37], and conformal prediction (CP). Among these, CP stands out for its ability to provide distribution-free, model-agnostic, and statistically guaranteed confidence intervals around the model predictions [38–40]. Unlike Bayesian methods that rely on prior distributions or bootstrapping approaches that require multiple retraining iterations, CP constructs prediction intervals based on observed data only, without any assumption about the underlying data distribution. It is model-agnostic, meaning it can be applied to any trained model—including deep neural networks—and it provides non-asymptotic, finite-sample statistical guarantees on the coverage probability of the resulting prediction intervals. This makes CP particularly advantageous in semiconductor manufacturing processes, where data distributions may be complex or even unknown, the explicit and proper decisions need to be made under strict reliability constraints. These coverage guarantees rely on the idea of exchangeability, which means in the manufacturing process, training, calibration, and test subsets all come from the same process and have similar statistical conditions. This enables the uncertainty seen in a small calibration subset to be a good match to new test data under the same production conditions.

Among the various CP techniques used for UQ, several variants have been proposed to fit different modeling contexts. Classical Conformal Prediction, for instance, estimates absolute error scores to determine the prediction band, offering valid intervals without relying on regression model calibration. Locally Weighted Conformal Prediction (LWCP) further refines this approach by adjusting prediction intervals based on the local characteristics of the data, assigning greater importance to recent or relevant samples. LWCP, on the other hand, is particularly useful in non-stationary settings or when localized data behavior plays a critical role. In contrast, Conformalized Quantile Regression (CQR) combines the strengths of quantile regression and conformal prediction to produce robust and adaptive prediction intervals [41]. It is especially suited for regression tasks with heteroscedastic noise, spatial variability, or nonlinear structure with errors, as commonly found in semiconductor process data. Considering the pixel-by-pixel regression nature of our layout-to-SEM prediction task, where the prediction uncertainty varies across regions and patterns, CQR offers the most effective trade-off among accuracy, interval tightness, and statistical validity. We therefore adopt this framework to yield calibrated pixel-level prediction intervals tailored to local structural variability [42].

1.5. Calibration Flow and Transfer Learning

While CQR equips the model with calibrated and spatially adaptive prediction intervals, further enhancement of predictive robustness requires more than a single-pass calibration. In the approach of the final step, we extend CQR with a structured calibration flow that continuously refines the model through outlier detection and spatial-aware reweighting on a fixed, limited calibration set.

This final step incorporates a low-cost yet effective transfer learning strategy, in which only the decoder and output layers of the baseline U-Net architecture are fine-tuned, while the encoder remains frozen to preserve generalized low-level features. Freezing the encoder during uncertainty-aware fine-tuning is a well-recognized approach for preserving foundational representation quality

while adapting the task-specific outputs to more challenging spatial structures and boundary cases [18,34,43]. To guide this fine-tuning process, a pixel-level outlier-weighted map is computed based on conformal errors, with higher weights being assigned to spatial regions where the predicted intervals fail to capture the ground truth. These weights affect the per-pixel loss contribution during fine-tuning, allowing the model to focus its corrective updates on structurally uncertain or miscalibrated regions. Crucially, the entire calibration dataset—despite its small size—is reused in its entirety to make optimal use of valuable data to reinforce both local sensitivity and global consistency. Spatial outlier correction with reweighted fine-tuning in this valuable subset enables targeted correction without risking overfitting [44].

Recent studies have emphasized the value of such targeted fine-tuning routines, demonstrating improvements in predictive sharpness, uncertainty awareness, and generalization robustness even under constrained data settings. Alternative strategies, such as those proposed by Krishnan and Tickoo (2020), define differentiable calibration-aware loss functions that directly optimize confidence reliability rather than only predictive accuracy, paving the way toward robust and interpretable uncertainty-aware modeling frameworks [45].

1.6. Contribution and Scope of This Work

This study presents a closed-loop adaptive framework for progressive enhancement and correction of layout-to-SEM image reconstruction, where model predictions are continuously improved using calibrated prediction intervals and pixel-level outlier detection. The approach presented first generates UQ results with formal statistical guarantees. An outlier map is then constructed from these results to identify and flag spatial regions with associated prediction risk. Finally, targeted model refinement is applied to improve reconstruction accuracy, particularly in these high-error regions, ensuring more reliable performance across the entire layout space. Such a closed-loop feedback process also enables a data-efficient and scalable approach for defect-aware model enhancement. This, in turn, facilitates robust layout-to-SEM prediction in practical NIL applications.

As a methodological contribution, this study introduces our integrated framework for uncertainty-aware pattern learning with tailor-made specific features to fit the semiconductor manufacturing in particular. The proposed pipeline combines:

- (1) U-Net-based CNN model for hierarchical spatial feature learning
- (2) CQR for interval-based predictions with statistical coverage guarantees
- (3) Pixel-level outlier detection for localized uncertainty awareness
- (4) Outlier-weighted fine-tuning strategy for enhancing adaptability to spatial variability

Together, these steps constitute a unified and extensible platform supporting both predictive performance and uncertainty calibration. While the context has been explicitly illustrated here for NIL processes in particular, our framework is designed with a modular structure and pixel-level generality. This design choice renders the approach applicable to a wide variety of layout-driven semiconductor processes in general. It offers potential for implementation across multiple stages of the manufacturing flow, including patterning, inspection, and defect-aware optimization as a whole.

2. Materials and Methods

2.1. Dataset Preparation

To clearly analyze the actual NIL manufacturing process and apply custom optimization for this technology, we have developed a simulation framework that is based on physical principles instead of pure data-driven heuristics. The simulation accounts for essential physical mechanisms in three integrated stages: exposure dose distribution and proximity effects in electron beam lithography (EBL), anisotropic material removal driven by angular ion flux in reactive ion etching (RIE), and viscoelastic resist deformation under mold pressure during NIL process. By modeling spatial variations in exposure, etch rates, and resist flow, the framework—which integrates well-established

physical principles from prior EBL, RIE, and NIL modeling literature [46–51]—generates physically realistic SEM output images that reflect nanoscale pattern variations, which provides high-fidelity surrogate data for our approach analysis. A dedicated manuscript detailing this physics-based simulation framework is currently in preparation.

Our data management strategy involves the division of the design layouts and the corresponding SEM target images into four independent subsets: training, validation, calibration, and test. All subsets are statistically dependent, but distinctive from spatial and structural variations of the design layouts, though they originate from the same process domain. The distinction allows for fair model evaluation and generalization, which is crucial for UQ and robust performance validation in downstream tasks such as OPC verification and defect-aware process learning.

To enhance pattern diversity and promote more comprehensive spatial variation learning, we applied deterministic data augmentation through fixed-angle rotations [19]. In particular, each original image was rotated by 90°, 180°, and 270°, resulting in three additional geometrically transformed variants for every input. This effectively expanded the dataset size by a factor of four. These rotational transformations also hold significant physical relevance in the context of NIL, as the pattern symmetry in layout designs is typically an exhibition of rotational invariance, making this form of augmentation particularly suitable for mask-based pattern prediction tasks.

All images were preprocessed to match a fixed resolution of 256×256 pixels. This resolution was selected as a balance between geometric fidelity and computational efficiency. In our dataset, each layout image usually has one primary pattern unit, named contour, which corresponds to a local structural feature. The resolution of 256×256 pixels provides sufficient detail to preserve the shape and boundary information of each contour. If fewer pixels were used (i.e., lower resolution), each pixel would cover a larger physical area. This would increase the risk of ambiguity along the boundary between the contour and the background, and could cause errors in local feature interpretation as well as reduce the overall prediction accuracy of the model. On the other hand, using a higher resolution would dramatically increase training and inference computational cost without proportionate gains in learning performance. Therefore, 256×256 serves as a practical trade-off that maintains essential topological information without excessive training complexity.

Design mask layout patterns were binarized with a fixed intensity threshold of 127 for maintaining topological consistency, whereas SEM images were normalized into the [0,1] range by dividing pixel values by 255.0 for stable convergence during CNN training. These preprocessing steps, shown in **Figure 1**, not only ensure physical interpretability, numerical stability, and generalization, but also enhance model training effectiveness. This is particularly critical under conditions of limited data availability and measurement uncertainty, which are common in high-resolution semiconductor manufacturing process modeling scenarios.

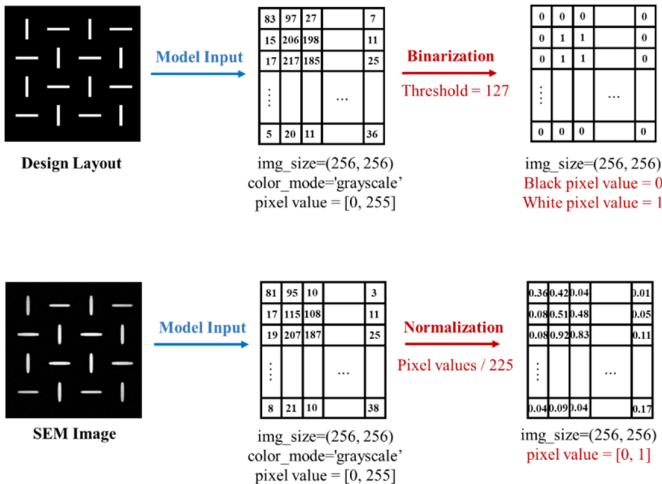


Figure 1. Preprocessing steps of design layout patterns and SEM images.

2.2. CNN-Based Model and Training

For performing pixel-by-pixel prediction of grayscale SEM images from binary mask layouts, we employed a U-Net-based CNN architecture. The U-Net model, originally developed by Ronneberger et al. for biomedical image segmentation [21], is a convolutional autoencoder designed for encoding both global context and subtle spatial details through a symmetric encoder-decoder structure with skip connections. The model is named “U-Net” due to the distinctive U-shaped architecture formed by its contracting and expansive paths, which mirror each other across a bottleneck layer. In our study, we used a shallow variant of U-Net to fit the limited training dataset and prevent overfitting, without losing the essential properties of spatial localization and low-level feature retention.

As shown in **Figure 2**, the proposed architecture consists of three main components: a contracting path (encoder), a bottleneck, and an expansive path (decoder). The contracting path is responsible for downsampling the input along with learning hierarchical features. It contains two convolutional blocks, each comprising two consecutive 3×3 convolutional layers with ReLU activation and zero-padding, followed by a 2×2 max-pooling operation that halves spatial resolution. The number of filters rises from 16 to 32 as we go deeper into the network. This design allows the encoder to progressively capture sophisticated fine-grained features while compressing the spatial dimensions from 256×256 to 64×64.

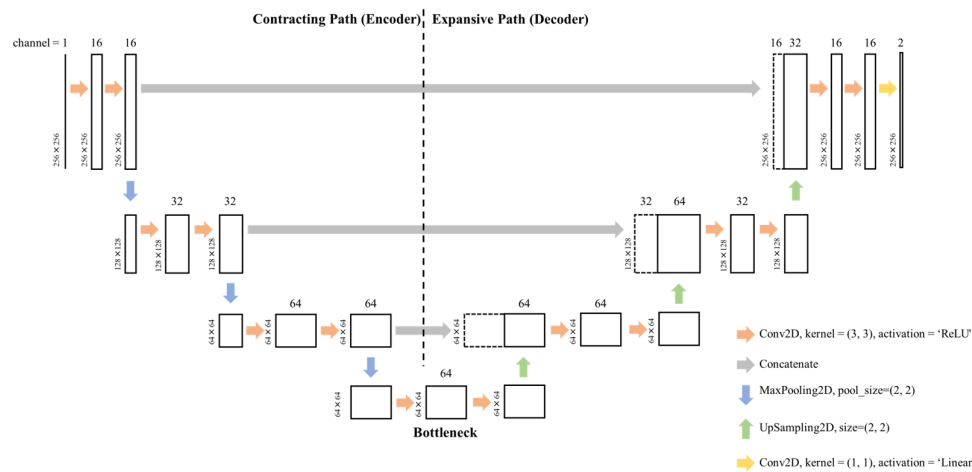


Figure 2. Shallow U-Net architecture with encoder–bottleneck–decoder structure for grayscale SEM quantile prediction.

The encoder is then followed by a bottleneck module, consisting of two 3×3 convolutional layers with 64 filters and ReLU activation. The bottleneck in our shallow model, which serves as the most compressed representation of the input [52], has a spatial resolution of 64×64. It still preserves essential spatial information, which in turn helps achieve a balance between semantic abstraction and localization capacity. Semantic abstraction refers to the capacity of the model to understand what is in the image, such as edges, textures, or structural regions. Localization capacity, on the other hand, means having the capacity to retain the spatial context of each feature with respect to the original spatial layout. In typical deep U-Net models, having deeper bottlenecks can enhance semantic encoding but at the cost of losing spatial details. In contrast, our shallow U-Net maintains a relatively large feature map (64×64) at the bottleneck stage. Here, in turn, helps to preserve fine-grained spatial resolution while it still benefits from hierarchical representation. This structure ensures that the model can simultaneously capture significant features and preserve pixel-level accuracy.

The expansive path reconstructs the feature maps back to the original input resolution. It consists of two upsampling steps of 2×2 nearest-neighbor upsampling, followed by concatenation with the corresponding feature maps from the encoder (skip connections), and two 3×3 convolutional layers

with ReLU activation. These skip connections, through which low-level spatial information directly passes from the encoder to the decoder, preserve fine-grained boundary details that are typically lost during downsampling. The number of filters keeps decreasing from 64 to 32 and then 16 in the output layers. Finally, one 1×1 convolutional layer with two output channels is used to provide two continuous-valued output maps corresponding to the predicted lower and upper quantiles of the SEM grayscale intensity. The model takes as input grayscale images of size (256, 256, 1) and produces a final output tensor of size (256, 256, 2), thereby allows simultaneous quantile prediction for each pixel.

Given the relatively small size of the dataset (24 training images with $4 \times$ augmentation), this shallow version of the U-Net is a good trade-off between model expressiveness and generalization. Decreasing the depth reduces the overfitting risk without sacrificing the structural advantages of the U-Net design. The model is trained using a conformalized quantile regression custom loss function described in the following section.

2.3. Conformalized Quantile Regression (CQR)

In order to implement the UQ of our shallow U-Net model, we apply conformalized quantile regression (CQR), which teaches the ability to predict lower and upper quantile bounds for each pixel's grayscale SEM value simultaneously. The loss function relevant to CQR is defined as a sum of asymmetric pinball losses ($\mathcal{L}_{q_{low}}, \mathcal{L}_{q_{high}}$) for two different quantiles:

$$\mathcal{L}_q(y, \hat{y}) = \frac{1}{HW} \sum_{i,j} \left[\max(q(y_{ij} - \hat{y}_{ij}), (1-q)(y_{ij} - \hat{y}_{ij})) \right] \quad (1)$$

$$\mathcal{L}_{CQR} = \mathcal{L}_{q_{low}}(y, \hat{y}^{low}) + \mathcal{L}_{q_{high}}(y, \hat{y}^{high}) \quad (2)$$

Here, $H \times W = 256 \times 256$ corresponds to the spatial resolution of a single grayscale SEM image, i.e., the total number of pixels per image. $y_{ij} \in [0,1]$ is the normalized ground truth SEM value. \hat{y}^{low} , \hat{y}^{high} represent the predicted quantiles for a specified confidence level. In our case, we set $q_{low} = 0.05$ and $q_{high} = 0.95$, corresponding to a 90% prediction interval.

Following the first training phase, a conformal calibration procedure is performed on a held-out calibration set to find the necessary quantile threshold to achieve the target coverage level and hence ensure reliable uncertainty quantification. For each pixel in the calibration data, the conformal error e_{ij} is computed as:

$$e_{ij} = \max(\hat{y}_{ij}^{high} - y_{ij}, y_{ij} - \hat{y}_{ij}^{low}) \quad (3)$$

The calibration constant q^* is then determined as the empirical of the set of errors:

$$q^* = \text{Quantile}_{1-\alpha}(\{e_{ij}\}^{H \times W}) \quad (4)$$

Under the exchangeability assumption, this condition guarantees that the prediction interval is statistically valid across the calibration distribution:

$$P_{x \sim \mathcal{D}} [y_{ij}(x) \in [\hat{y}_{ij}^{low}(x) - q^*, \hat{y}_{ij}^{high}(x) + q^*]] \geq 1 - \alpha \quad (5)$$

Here, $\alpha = 0.1$ denotes the rate of miscoverage. For a new input x sampled from the same distribution \mathcal{D} , the ground truth $y_{ij}(x)$ must lie in the calibrated interval $[\hat{y}_{ij}^{low}(x) - q^*, \hat{y}_{ij}^{high}(x) + q^*]$ with probability at least $1 - \alpha$. Exchangeability makes sure that the calibration set, as a subset drawn from the same distribution, enables the model to generalize its robust prediction intervals to unseen but similar distributed patterns.

2.4. Outlier-Weighted Calibration and Transfer Learning

To further improve the reliability and spatial accuracy of the layout-to-SEM image reconstruction task under the constraints of process-induced uncertainty—including aleatoric variability from lithographic or SEM noise and epistemic uncertainty due to limited data or model underfitting—we employ a conformal calibration strategy with a transfer learning pipeline based on Section 2-2 and 2-3. This phase aims to enhance both the calibration quality and spatial accuracy of predictions in lithographically sensitive regions. Emphasis is placed on scenarios with sparse data and on structurally complex regions, including feature edges and imprint-induced residual defects. This is especially crucial in NIL process windows involving low residual layer thickness or incomplete release, where SEM intensity variations point out the potential defect onset.

The proposed calibration flow consists of three main steps: (1) Base quantile model calibration, (2) Outlier-enhanced pixel-by-pixel weighting strategy, and (3) Encoder-frozen transfer fine-tuning:

(1) Base quantile model calibration. Following the base quantile model calibration described in Section 2-3, we adopted the empirically estimated calibration constant q^* to construct valid pixel-level prediction intervals $[\hat{y}_{ij}^{low}(x) - q^*, \hat{y}_{ij}^{high}(x) + q^*]$ with guaranteed marginal coverage of at least $1 - \alpha$. This post-CQR prediction intervals provide the foundation for the subsequent fine-tuning strategy.

(2) Outlier-enhanced pixel-by-pixel weighting strategy. To emphasize structurally critical regions during fine-tuning, we introduce an outlier-enhanced pixel-by-pixel weighting strategy through conformal prediction analysis. This strategy assigns greater weight to regions of increased uncertainty, so the model pays more attention to spatial locations with higher predictive deviation. For each pixel in the calibration dataset, we calculated the conformal error e_{ij} and calibration constant q^* earlier in Eq. 3 and Eq. 4. Here, q^* presents as the global outlier threshold. Pixels for which the error exceeds q^* are considered outliers and are assigned higher weights. The pixel-by-pixel weight map is then defined as:

$$w_{ij} = \begin{cases} \gamma & \text{if } e_{ij} > q^* \\ 1.0 & \text{otherwise} \end{cases} \quad (6)$$

Here, $\gamma > 1$ is a scalar denoting the level of reweighting, set to 1.3 in our experiments. The pixel-by-pixel weight maps were directly concatenated with the ground truth images to form a 2-channel training target tensor, as shown in Eq. 7, of shape (256, 256, 2), where the first channel represents the normalized ground truth and the second channel holds the corresponding pixel weights.

$$\tilde{y}_{ij} = [y_{ij}, w_{ij}] \quad (7)$$

By guiding the loss function to focus on pixels with greater uncertainty violations, this strategy enables the model to adaptively correct outlier regions—particularly those near feature edges or NIL-induced residual hotspots—without compromising performance in already well-calibrated regions. The resulting weight maps are seamlessly integrated into the CQR loss during the transfer fine-tuning stage, allowing targeted recalibration of uncertain regions using existing labels only, thereby improving reliability under limited-data conditions.

(3) Encoder-frozen transfer fine-tuning. To mitigate catastrophic forgetting from the initial U-Net training while enabling targeted local adaptation during fine-tuning, we adopted a transfer learning strategy in which the encoder of our baseline U-Net model was frozen, and only the decoder and output layers were fine-tuned. This ensures that the low-level feature extraction of the baseline model remains intact while allowing targeted improvement in regions previously prone to high prediction errors or structural uncertainty outliers. Let the original model parameters be denoted as $\theta = \{\theta_e, \theta_d\}$, where θ_e and θ_d correspond to the encoder and decoder parameters, respectively. During fine-tuning, we imposed the constraint:

$$\theta_e^{transfer} = \theta_e^{base}, \text{ with } \frac{\partial \mathcal{L}}{\partial \theta_e^{transfer}} = 0 \quad (8)$$

This enforces that encoder parameters remain unchanged, that is, to ensure there are no gradient updates applied to the encoder parameters. The transfer fine-tuning is performed on the calibration set, where \hat{y}_{ij} in Eq. 7 represents the reweighted label constructed by appending the pixel-by-pixel weight map w_{ij} , based on our outlier-enhanced weighting strategy described in Eq. 6. We retained the fine-grained pixel-by-pixel weighting directly within the fine-tuning objective function by embedding the weight map w_{ij} into the conformal quantile regression loss. The resulting fine-tuning function objective is defined as:

$$\mathcal{L}_q^{\text{weighted}}(y, \hat{y}) = \frac{1}{HW} \sum_{i,j} w_{ij} \left[\max \left(q(y_{ij} - \hat{y}_{ij}), (1 - q)(y_{ij} - \hat{y}_{ij}) \right) \right] \quad (9)$$

This formulation enables pixel-by-pixel error correction in spatially localized outlier regions, such as nanoscale hotspots or residual imprinting defects, while maintaining overall stability in structurally consistent areas. It further allows the decoder to recalibrate the predictive uncertainty while preserving the high-level semantic information of the encoder. After training, the fine-tuned model is evaluated alongside the baseline model using a unified evaluation framework. Specifically, both models are compared under a consistent metric, where pixel-by-pixel outliers are identified based on the same calibrated threshold q^* , and prediction validity is measured via coverage rate. These metrics provide a fair basis to quantify the improvement of the model in uncertainty calibration and spatial accuracy.

2.5. Evaluation Metrics

After concluding the two-stage calibration procedure of Sections 2-3 and 2-4 for both the baseline U-Net model and its transfer fine-tuned model, they are evaluated under a unified, application-specific framework designed specifically to the needs of OPC. The evaluation is based on two primary performance metrics: Mean Absolute Error (MAE) and Prediction Interval Coverage, which collectively determine the accuracy and reliability of the predicted SEM intensity profiles.

2.5.1. Mean Absolute Error (MAE)

To evaluate pixel-level accuracy, we introduce the MAE between the midpoint of the predicted quantile bounds and the corresponding ground truth pixel value. It serves as a critical measure of local feature accuracy, particularly reflecting how well fine-scale shape variations and edge positions are preserved. The formula is as below:

$$MAE = \frac{1}{HW} \sum_{i,j} \left| \frac{1}{2} (\hat{y}_{ij}^{\text{low}}(x) + \hat{y}_{ij}^{\text{high}}(x)) - y_{ij} \right| \quad (10)$$

In our setting, $H \times W = 256 \times 256$, indicating that each SEM image has 256×256 pixels. The quantile midpoint used here offers a stable estimate of the predicted surface profile, balancing the upper and lower bounds to provide a reliable pixel-wise estimate. In OPC applications, low MAEs are essential to meet in-fab tolerance requirements, with sub-pixel deviations often acceptable depending on design rules.

2.5.2. Prediction Interval Coverage

To quantify the reliability of uncertainty, we calculate the **coverage rate**, or the proportion of ground truth values that fall within the conformal prediction interval:

$$\text{Coverage} = \frac{1}{HW} \sum_{i,j} \mathbb{1}\{y_{ij} \in [\hat{y}_{ij}^{\text{low}} - q^*, \hat{y}_{ij}^{\text{high}} + q^*]\} \quad (11)$$

For a consistent evaluation process, the same calibration constant q^* is used for both baseline and transfer fine-tuned models. This strategy supports a more meaningful and informative model comparison by making an explicit test of whether transfer fine-tuning can reduce the outlier regions

identified in the baseline model. This approach aligns with our goal to evaluate the performance of transfer fine-tuning in correcting pixel-by-pixel uncertainty-aware prediction errors.

3. Results

3.1. Baseline Evaluation

The baseline pipeline, illustrated in **Figure 3**, involves training a U-Net together with CQR to estimate pixel-by-pixel prediction intervals for SEM image reconstruction. The network outputs both lower and upper bounds for every input layout, enabling spatially resolved UQ. The mean prediction for evaluation is provided by the average of the lower and upper bounds and is the point estimate chosen to represent the target image. Representative patterns and corresponding model predictions are illustrated in **Figure 4**. To further observe the model convergence and generalization performance, the loss curves for both training and validation phases are provided in **Figure 5**.

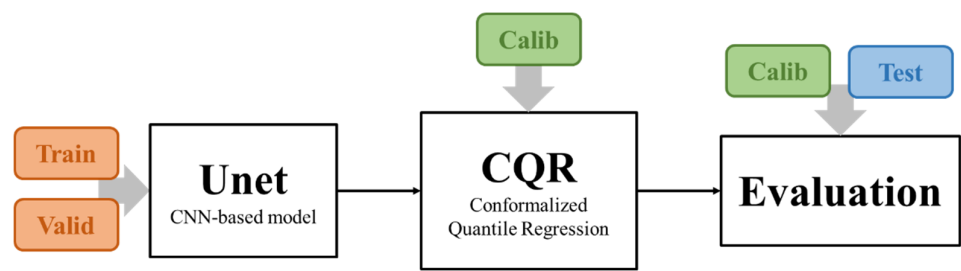


Figure 3. Baseline pipeline with U-Net training, CQR for UQ, and evaluation.

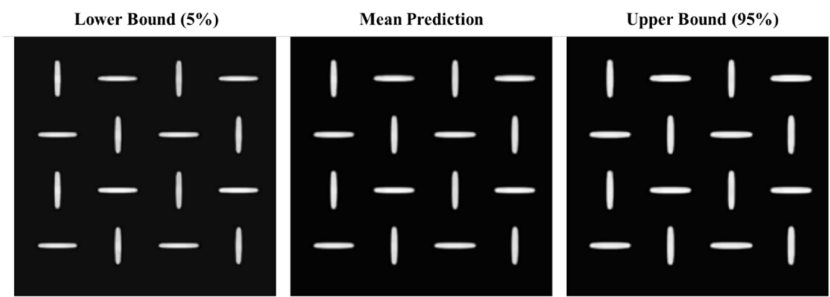


Figure 4. U-Net CNN-based prediction results showing lower bound, upper bound, and mean prediction of the SEM grayscale output.

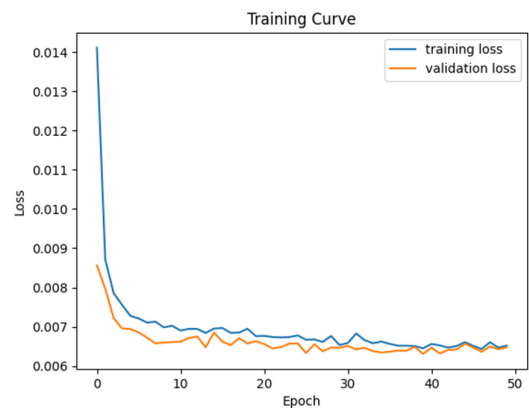


Figure 5. The loss curves for both the training and validation phases of baseline U-Net training.

The inclusion of noise in the SEM images generated from previous physical modeling is to emulate realistic imaging conditions. These SEM images, produced through physics-based simulation, serve as training targets for the baseline U-Net model. Since real SEM images inherently contain background noise, charging effects, and structural distortions, we simulate such imperfections by incorporating controlled noise during training. This allows the model to be more tolerant of these realistic artifacts. This step not only enables evaluation of the model's capacity to generalize from synthetic to real SEM domains but also mitigates the risk of overfitting to unrealistically perfect synthetic data, thereby enhancing the robustness of the learned representations.

The U-Net model, guided by CQR through the prediction of lower and upper quantile bounds for each pixel, naturally highlights features with pixel values around 1.0 (contours and patterns) while reducing the effect of background noise around 0.0. This focused attention allows accurate reconstruction of essential structural features by increasing the contrast between patterned and non-patterned regions, resulting in sharper and more clearly defined boundaries. **Figure 6** illustrates a pixel-by-pixel comparison of the ground truth and the model prediction intervals over a selected region (Pixels 26,000–26,500). The ground truth curve (black) shows high-frequency noise and fluctuation in low-intensity regions, as is commonly observed in SEM imaging due to charging effects or stochastic process variation. Conversely, the mean prediction (blue dashed line) exhibits significant denoising capability, successfully reconstructing the signal profile while effectively filtering out background noise. Furthermore, the prediction intervals (gray band between lower and upper bounds) remain a consistently narrow and stable range, especially in flat regions, indicating a high level of model confidence and limited uncertainty. This stability provides additional confirmation that the model effectively prevents local overfitting in small background segments and demonstrates commendable generalization performance in both high-contrast and low-intensity regions.

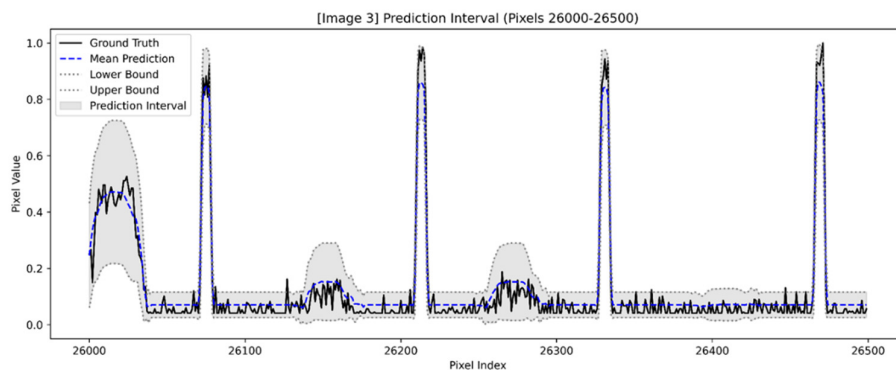


Figure 6. Pixel-by-pixel mean predictions and prediction intervals (with lower and upper bounds) plotted together with the ground truth for Image 3 (pixels 26,000–26,500).

Comparison of ground truth values and predicted quantile intervals, as shown in **Figure 7**, reveals that uncertainty is small in background regions and highest along edges, where pattern fidelity is most sensitive. This observation is consistent with the physical behavior in NIL processes, where post-imprint template release often results in local resist deformation near sidewalls due to surface tension or adhesion forces. These effects introduce geometric uncertainty in the reconstructed contours, as reflected in the increased prediction error at boundaries.

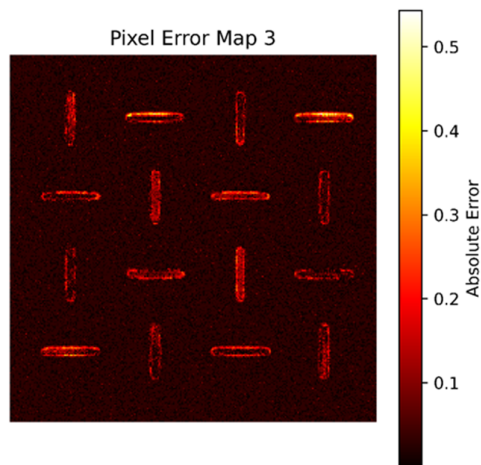


Figure 7. Pixel Error Map.

The pixel-by-pixel error distribution indicates an MAE of 0.0364, with errors concentrated at the edge of each patterned unit as shown in **Figure 7**. This regular and physically explainable uncertainty motivates the need for post-training calibration. By applying CQR calibration using a held-out dataset, the errors are sorted to determine a quantile threshold $q^* = 0.0617$, which regulates the prediction interval to cover 90% of the true pixel values. This calibrated bound, as shown in **Figure 8**, ensures statistical validity of the model’s confidence intervals and forms the basis for further refinement in the subsequent calibration-transfer workflow.

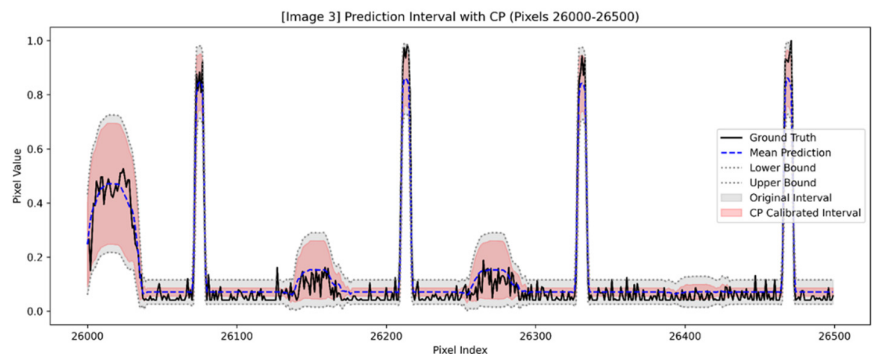


Figure 8. Pixel-by-pixel mean predictions, prediction intervals (with lower and upper bounds), plotted together with the ground truth for Image 3 (pixels 26,000–26,500) after CP calibrated.

After applying CQR to the prediction intervals of the baseline U-Net model, we obtain calibrated bounds that provide 90% pixel-by-pixel coverage with respect to a given error threshold. This calibration ensures that, with high statistical confidence, the prediction intervals contain the true pixel values, which directly allows an estimate of bounded MAEs.

To visualize the distribution of MAEs within the dataset, histograms of the calibration and test datasets are shown in **Figures 9** and **10**, respectively. The green dashed line in each figure indicates the average MAE, while the red dashed line represents the MAE at the 90th percentile of errors. In the calibration dataset, the average MAE is 0.0364 and the 90% coverage MAE is 0.0384. Consistent with the calibration results, the test dataset shows slightly higher values—0.0367 and 0.0386, respectively. These minor differences, confined to the fourth decimal place, are well within acceptable bounds and reflect the expected structural and spatial variability intentionally preserved between the subsets. Given the fact that the calibration and test datasets are sampled from distinct layout regions originating from the same process domain, such deviation supports the robustness and

generalizability of the conformal calibration framework rather than indicates overfitting or model bias. This confirms that the model maintains consistent MAE behavior over different layout instances, which validates its utility for uncertainty-aware downstream applications such as OPC verification and lithographic process optimization.

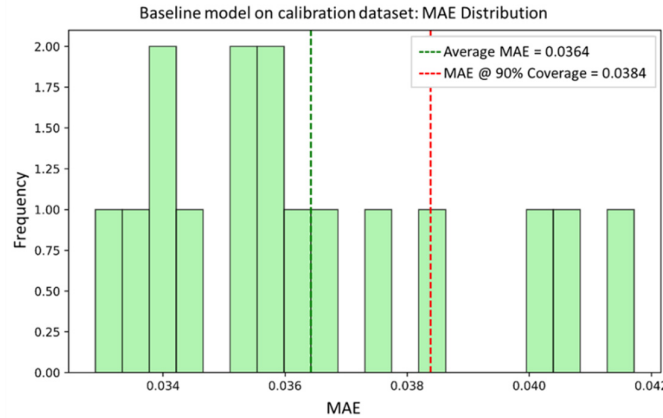


Figure 9. Histogram of MAE distribution for the baseline model on the calibration dataset.

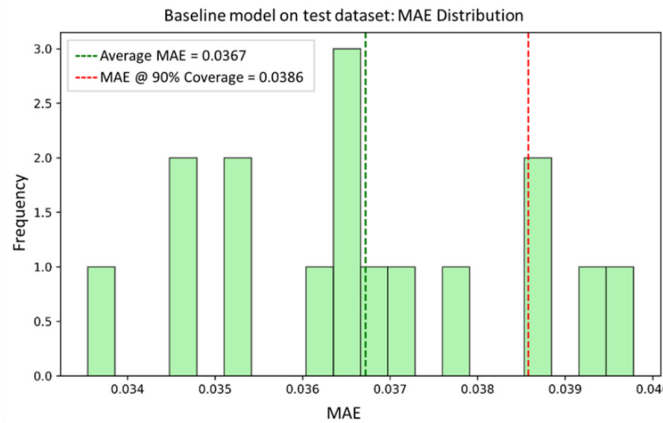


Figure 10. Histogram of MAE distribution for the baseline model on the test dataset.

Furthermore, this statistically consistent behavior enables the definition of outliers by thresholding the MAE beyond the 90% coverage bound. As we will demonstrate in the following outlier map analysis, the calibrated uncertainty region provides a meaningful reference to identifying regions where model predictions become less reliable, thereby supporting uncertainty-aware model diagnosis and confidence-aware pattern evaluation.

Figure 11 presents outlier maps for selected samples from the test dataset, where outliers are defined as pixels whose MAE exceeds the conformal threshold of 0.0384. In these plots, red pixels indicate regions where the prediction of the model deviates beyond the calibrated uncertainty intervals, while green pixels indicate areas that are still within the specified tolerance. In all four plots, there is an evident spatial pattern: most of the outliers are concentrated near the contour boundaries of the pattern features, especially along the elongated edges where physical process variation and resist behavior are more pronounced. This aligns with expectations based on the physics of NIL, in which edge defects and residual deformation are more likely to introduce structural uncertainty.

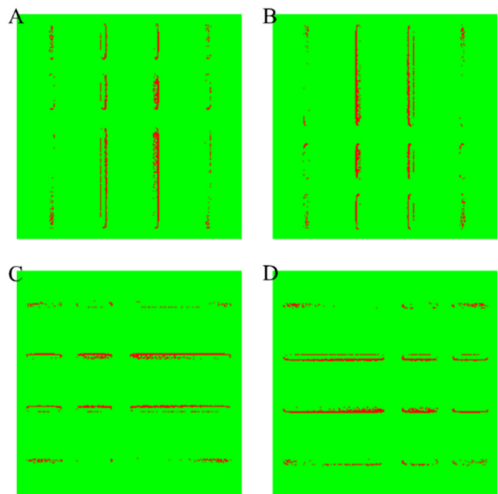


Figure 11. Pixel-by-pixel outlier maps of selected samples from the test dataset.

The large number of outliers in the test dataset indicates that although the model is strong on average performance, there remain localized regions with lower confidence and severe prediction error. This demonstrates the need for improvement beyond global conformal calibration. In particular, it motivates the development of an outlier-weighted calibration framework that can selectively adjust prediction outputs based on local reliability. It also promotes the use of transfer learning techniques to improve generalization on previously unseen pattern types. The results of these form the foundation for the next section, in which we present an Outlier-Weighted Calibration and Transfer Learning method that aims to promote robustness and spatial generalizability in uncertainty estimation.

3.2. Outlier-Weighted Calibration and Transfer Learning Evaluation

As shown in **Figure 12**, the transfer learning pipeline starts by evaluating the baseline U-Net model through a held-out calibration dataset. The conformal errors derived from this evaluation are then utilized for outlier detection, where pixels fail to meet the pre-specified coverage. Based on the specific downstream target (e.g., minimization of MAE, Line Edge Roughness (LER), or CD Uniformity (CDU)), practitioners can establish case-by-case specific thresholds for outlier detection. In our case, we apply the conformal 90% coverage criterion, treating deviations beyond this interval as meaningful indicators of outliers in the model. The identified outliers are then sent to the next phase: outlier-enhanced pixel-by-pixel weighting strategy and encoder-frozen transfer fine-tuning, the pipeline that reweights the loss function at the pixel level, updates the decoder layers only to enhance local prediction accuracy, and then preserves global structural representation.

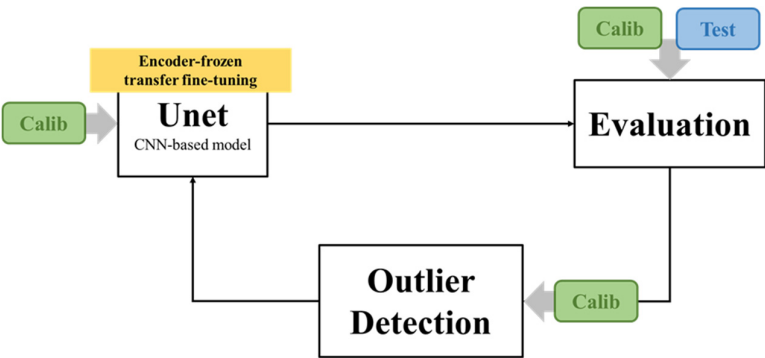


Figure 12. Transfer learning pipeline with outlier detection and encoder-frozen fine-tuning.

Figure 13 illustrates the training loss curve across 33 epochs during the outlier-weighted transfer fine-tuning stage. Compared with the baseline model trained from scratch for 50 epochs, the transfer learning phase employed an early stopping strategy and was terminated early at epoch 32 after the convergence criteria, as the loss plateaued after an initial sharp drop. Although the final loss value is greater than that of the baseline model in **Figure 5**, this is expected given the reweighted training objective. The transfer fine-tuning phase utilizes the held-out calibration dataset in combination with pixel-by-pixel outlier weighting. This helps target high-uncertainty regions across the entire sample space. The initial sharp drop and rapid stabilization of the loss curve demonstrate that only slight updates were needed to refine performance in regions of localized uncertainty. This aligns with the encoder-frozen fine-tuning pipeline structure, which preserves previously learned general features and adjusts the decoder to correct prediction errors accordingly. As a result, the time for training and computational cost are significantly reduced, without compromising the model’s capacity for adaptation. It is important to note that the transfer fine-tuning loss is not directly comparable to the baseline’s global loss, as training data distribution and loss weighting structure are different. Our goal in this step is to improve local calibration and reliability in the regions uncovered by conformal prediction, not to minimize overall loss. This approach aligns with previous research on uncertainty-aware fine-tuning literature [41,53], where task-specific re-optimization yields meaningful improvements without full retraining.

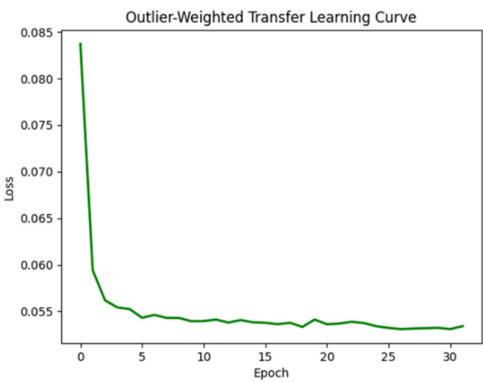


Figure 13. The loss curves for the Transfer learning flow.

As shown in **Figures 14–16**, the fine-tuned model demonstrates improved spatial confidence (**Figure 14**), higher and more stable coverage rates (**Figure 15**), and reduced variance in MAE (**Figure 16**). These results validate the effectiveness of our uncertainty-aware fine-tuned strategy, which integrates CQR outlier detection with pixel-by-pixel weighting and encoder-frozen transfer fine-tuning. This calibrated feedback approach enables the model to improve localized reliability without compromising global consistency.

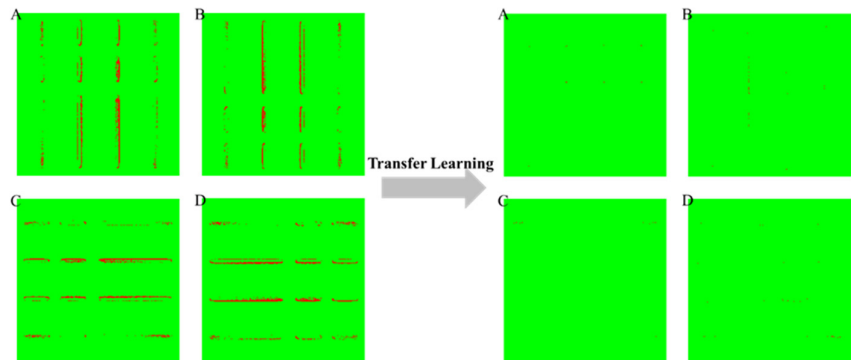


Figure 14. Outliers map before and after outlier-weighted calibration and transfer fine-tuning.

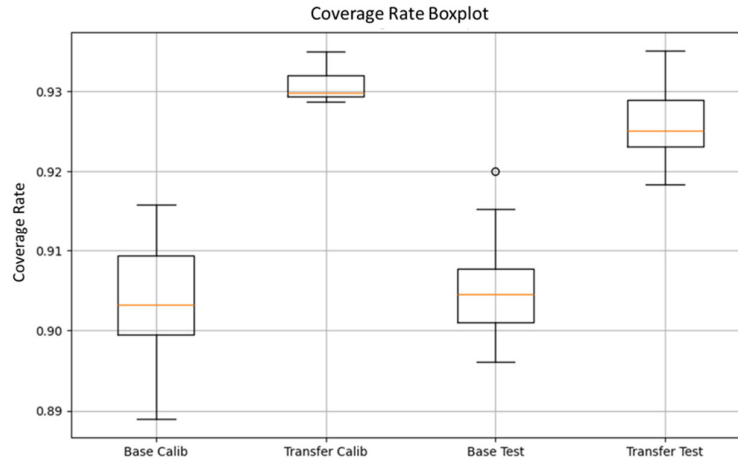


Figure 15. Boxplot comparison of coverage distributions from baseline and transfer fine-tuning settings for both calibration and test datasets.

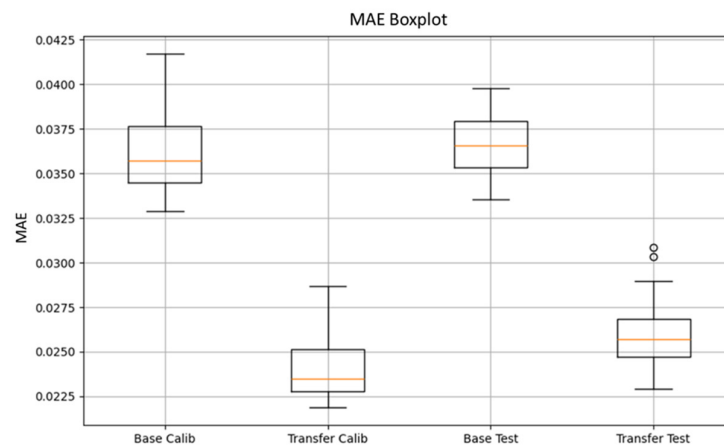


Figure 16. Boxplot comparison of MAE distributions from baseline and transfer fine-tuning settings for both calibration and test datasets.

The visual impact of this transfer fine-tuning step is demonstrated in **Figure 14**, which compares the outlier maps before and after transfer fine-tuning. In each map, regions that violate the conformal 90% coverage requirement are labeled, with red pixels denoting outlier regions. After applying our pixel-by-pixel outlier-weighted strategy in combination with encoder-frozen transfer fine-tuning, a significant reduction in the density of red pixels is clearly observed in all four cases. This reduction confirms again that the model has been able to learn effectively to correct spatially local outliers of uncertainty without sacrificing overall model integrity. It is also interesting to note that the definition of the outlier is the same as that of the original conformal threshold of the baseline model, which ensures that improvement is evaluated against the same statistically rigorous standard.

Quantitative results are displayed in **Figures 15** and **16**, where coverage rate and MAE distributions of baseline and transfer fine-tuning models on both the calibration and test datasets are plotted. As shown in **Figure 15**, the transfer fine-tuning model displays greater median coverage with a smaller interquartile range (IQR) for both datasets. This demonstrates excellent stability and consistency in the spatial coverage of confidence intervals. Moreover, the percentage of pixels whose

MAEs fall below the original defined 90% conformal coverage threshold also increases, confirming the fact that the transfer fine-tuning model conforms better to statistical guarantees even under layout variability. Through **Figure 16**, the distribution of MAE further supports these results. Both the Transfer Calib and Transfer Test groups show remarkable reductions in error magnitude and variance compared to their baseline model distributions. The Transfer Test set, in particular, exhibits a lower median MAE and fewer extreme outliers, signifying enhanced generalization to previously unseen patterns. The ability of the transfer fine-tuning model to reduce the variation of MAE while maintain statistical coverage is a demonstration of the advantages of incorporating uncertainty-aware feedback during training, especially in data-limited scenarios. It is worth noting that the hollow circle markers in **Figures 15** and **16** represent statistically rare points, as defined by the IQR rule in boxplot construction. These points correspond to rare but valid samples, which merely lie far from the central distribution. Notably, these rare points are not explicitly related to the uncertainty estimation mechanism by conformal prediction used in this study. Although a few rare points appear in the transfer fine-tuning setting, the overall reduction in median along with IQR in both calibration and test datasets guarantees the effectiveness and generalization capability of our outlier-weighted transfer fine-tuning strategy.

These findings collectively demonstrate that the integration of localized outlier detection with targeted decoder-level transfer fine-tuning brings substantial improvements to spatial uncertainty modeling. The model is increasingly sensitive to local reliability, preserves global calibration validity, and maintains generalization across a variety of pattern types. These overall models are well-aligned with the practical demands of layout-to-SEM image reconstruction in NIL, where both spatial accuracy and confidence awareness are essential for downstream tasks such as OPC validation and hotspot monitoring.

4. Discussion and Implications

The quantitative results presented in **Table 1** demonstrate that the proposed encoder-frozen, outlier-weighted transfer fine-tuning pipeline consistently outperforms the baseline U-Net, especially for pixel value accuracy and the coverage rate of the prediction intervals. On the calibration set, the mean MAE is reduced from 0.0355 to 0.0235 in normalized pixel-intensity units, an approximate improvement of 34%, and the mean coverage rate rises from 0.902 to 0.931. Similar significant improvements are observed in the independent test set as well, which confirms that the improvement is not an illusion of calibration but rather an indication of actual generalization. The corresponding reduction in standard deviation—or the reduction in MAE standard deviation, to be precise, from 0.0028 to 0.0020 in calibration—shows that the model is not only more accurate on average but also markedly more stable across diverse patterns.

Table 1. Performance Metrics of Baseline and Transfer Fine-tuning Models on Calibration and Test Sets.

Models vs. Metrics		MAE		Coverage rate	
		Mean	STD	Mean	STD
Calibration	Baseline	0.0355	0.0028	0.902	0.0085
	Transfer learning	0.0235	0.0020	0.931	0.0020
Test	Baseline	0.0365	0.0023	0.904	0.0065
	Transfer learning	0.0255	0.0018	0.926	0.0040

From a methodological perspective, freezing the encoder retains the hierarchical representation of nanoscale features learned in the initial and large-scale training phase, while the pixel-by-pixel error weighting focuses the decoder’s limited fine-tuning capacity on structurally sensitive regions. This targeted adaptation focuses the fine-tuning capacity on the most challenging pixels, directing learning effort toward difficult regions without requiring any new manually labeled data. Compared with existing uncertainty-calibrated CNN approaches based on global loss re-weighting or full-network retraining, our strategy achieves comparable reliability with an 80 % reduction in sample

requirement and 61 % shorter fine-tuning time, confirming its data-efficiency. Moreover, relative to the baseline U-Net with CQR calibration, our transfer fine-tuning strategy, which is based on an encoder-frozen structure, achieves equivalent coverage (0.931 vs. 0.926) and MAE (0.0235 vs. 0.0255) in calibration and test set while uses only 48 labeled images fewer than the 240 needed for the baseline—and reduces fine-tuning time from 36 mins to 10 mins on an RTX 3090 chip. As shown in **Table 2**, these results strongly indicate the data-efficiency and practical reliability of the proposed approach.

Table 2. Resource Efficiency Comparison Between Baseline and Transfer Fine-Tuning Models.

Metric	Baseline	Transfer Fine-tuning	Reduction
Images used	240	48	80%
GPU time (RTX 3090)	36 min	10 min	72%

In practice, the coverage improvement reported above translates to fewer guard-band iterations needed when transferring the model to new product layouts, which reduces OPC turn-around time. Moreover, outlier maps calibrated from the baseline U-Net using CQR are used to guide the transfer fine-tuning pipeline. Hence, this model effectively localizes error hotspots, allowing process engineers to focus metrology resources on high-risk regions instead of performing blanket inspections. This is evidenced in **Table 1** by the improved prediction interval coverage rate, which increases from 0.904 to 0.926 on the test set, indicating more reliable coverage of pixel-level features across structurally diverse but statistically aligned layout patterns from the same process domain. While this confirms better generalization, small error variability may still persist. Augmenting the calibration set with synthetically generated extreme patterns, therefore, remains an immediate avenue for further variance reduction. Future work will explore this direction to further enhance uncertainty calibration and support more robust generalization across challenging layout configurations.

5. Conclusions and Outlook

This study introduces a two-stage, data-efficient framework for layout-to-SEM image reconstruction that addresses the three most widely used challenges to large-scale AI deployment—data sparsity, annotation quality, and the need for robust generalization across structurally diverse yet process-consistent layout patterns. First, a baseline CNN-based model is initially hybridized with CQR to produce statistically valid prediction intervals and associated pixel-level outlier maps. Second, these generated maps then guide an encoder-frozen, outlier-weighted transfer fine-tuning step that updates the decoder without interfering with the learned hierarchical features.

Experiments confirm that the proposed strategy reduces the mean MAE from 0.0365 to 0.0255 and improves the interval coverage rate from 0.904 to 0.926 on an independent test set, while just uses only 48 labeled images and one-third of the original GPU time. These results translates to fewer OPC guard-band iterations and more selective hotspot triage, demonstrating practical value under real-world manufacturing conditions, particularly under constraints of sparse data and the high cost of annotation in layout-to-SEM applications.

Future work will expand this foundation in several ways: (i) Large-scale contour clustering with fuzzy pattern grouping that categorizes the layout corpus into structurally coherent subsets. Such grouping enables cluster-specific calibration and model refinement for improving prediction accuracy on complex patterns and supporting pattern-aware OPC decisions in manufacturing. (ii) Cross-process validation that extends the approach to diverse imaging and fabrication settings, including AFM, optical profilometry, and lithographic variations such as DUV and EUV. This will enable the evaluation of robustness under both sensor-dependent and process-dependent noise conditions. (iii) Hybrid physics-and-ML integration that links learned uncertainty maps with resist or imprint simulators to enable closed-loop process optimization. Together, these directions establish

our proposed framework as a compact and versatile solution for uncertainty-aware image reconstruction in advanced semiconductor manufacturing, while simultaneously address the overarching requirement for data-efficient, well-calibrated, and highly generalizable AI systems.

Author Contributions: Conceptualization, J.C.; methodology, J.C.; software, J.C.; validation, J.C.; formal analysis, J.C.; investigation, J.C.; writing—original draft preparation, J.C.; writing—review and editing, E.L.; supervision, E.L. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The data used in this study are not publicly available due to institutional restrictions and industry confidentiality agreements.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

NIL	Nanoimprint lithography
CQR	Conformalized quantile regression
MAE	Mean absolute error
OPC	Optical proximity correction
AI	Artificial intelligence
SEM	Scanning electron microscope
ADI	After-development inspection
SMOTE	Synthetic minority over-sampling
UQ	Uncertainty quantification
ML	Machine learning
CNN	Convolutional neural network
CP	Conformal prediction
LWCP	Locally Weighted Conformal Prediction
EBL	Electron beam lithography
RIE	Reactive ion etching

References

1. Young, W.-B., *Analysis of the nanoimprint lithography with a viscous model*. Microelectronic Engineering, 2005. 77: p. 405-411.

2. Hirai, Y., et al., *Pressure and resist thickness dependency of resist time evolutions profiles in nanoimprint lithography*. Microelectronic Engineering, 2008. 85(5): p. 842-845.

3. Ifuku, T., et al. *Nanoimprint lithography performance advances for new application spaces*. in *Novel Patterning Technologies 2024*. 2024. SPIE.

4. Ifuku, T., et al., *Nanoimprint lithography performance advances for new application spaces*. SPIE Advanced Lithography + Patterning. Vol. 12956. 2024: SPIE.

5. Kulmala, T., et al., *Single-nanometer accurate 3D nanoimprint lithography with master templates fabricated by NanoFrazor lithography*. SPIE Advanced Lithography. Vol. 10584. 2018: SPIE.

6. Sirotkin, V., et al., *Coarse-grain method for modeling of stamp and substrate deformation in nanoimprint*. Microelectronic Engineering, 2007. 84(5): p. 868-871.

7. Takeuchi, N., et al., *Fabrication of dual damascene structure with nanoimprint lithography and dry-etching*. SPIE Advanced Lithography + Patterning. Vol. 12497. 2023: SPIE.

8. Aihara, S., et al., *NIL solutions using computational lithography for semiconductor device manufacturing*. SPIE Advanced Lithography + Patterning. Vol. 12954. 2024: SPIE.

9. Chou, S.Y., P.R. Krauss, and P.J. Renstrom, *Nanoimprint lithography*. Journal of Vacuum Science & Technology B: Microelectronics and Nanometer Structures Processing, Measurement, and Phenomena, 1996. 14(6): p. 4129-4133.

10. Guo, L.J., *Nanoimprint lithography: methods and material requirements*. Advanced materials, 2007. **19**(4): p. 495-513.
11. Yan, Y., et al., *Machine learning virtual SEM metrology and SEM-based OPC model methodology*. Journal of Micro/Nanopatterning, Materials, and Metrology, 2021. **20**(4): p. 041204-041204.
12. Tseng, J., J. Chien, and E. Lee, *Advanced defect recognition on scanning electron microscope images: a two-stage strategy based on deep convolutional neural networks for hotspot monitoring*. Journal of Micro/Nanopatterning, Materials, and Metrology, 2024. **23**(4): p. 044201.
13. Ogusu, M., et al., *Nanoimprint post processing techniques to address edge placement error*. SPIE Advanced Lithography + Patterning. Vol. 12497. 2023: SPIE.
14. Dini, P., et al., *Overview on intrusion detection systems design exploiting machine learning for networking cybersecurity*. Applied Sciences, 2023. **13**(13): p. 7507.
15. Dini, P. and S. Saponara, *Cogging torque reduction in brushless motors by a nonlinear control technique*. Energies, 2019. **12**(11): p. 2224.
16. Akpabio, I. and S. Savari, *Uncertainty quantification of machine learning models: on conformal prediction*. Journal of Micro/Nanopatterning, Materials, and Metrology, 2021. **20**(4): p. 041206.
17. Došilović, F.K., M. Brčić, and N. Hlupić. *Explainable artificial intelligence: A survey*. in 2018 41st International convention on information and communication technology, electronics and microelectronics (MIPRO). 2018. IEEE.
18. Acun, C., et al., *Optimizing Local Explainability in Robotic Grasp Failure Prediction*. Electronics, 2025. **14**(12): p. 2363.
19. Elhanashi, A., et al., *AI-Powered Object Detection in Radiology: Current Models, Challenges, and Future Direction*. Journal of Imaging, 2025. **11**(5): p. 141.
20. Elhanashi, A., et al., *Deep learning techniques to identify and classify COVID-19 abnormalities on chest x-ray images*. SPIE Defense + Commercial Sensing. Vol. 12102. 2022: SPIE.
21. Ronneberger, O., P. Fischer, and T. Brox. *U-net: Convolutional networks for biomedical image segmentation*. in *Medical image computing and computer-assisted intervention—MICCAI 2015: 18th international conference, Munich, Germany, October 5-9, 2015, proceedings, part III* 18. 2015. Springer.
22. Isensee, F., et al., *nnu-net: Self-adapting framework for u-net-based medical image segmentation*. arXiv preprint arXiv:1809.10486, 2018.
23. Han, N., et al., *Multi-level U-net network for image super-resolution reconstruction*. Displays, 2022. **73**: p. 102192.
24. Xu, W., et al., *High-resolution u-net: Preserving image details for cultivated land extraction*. Sensors, 2020. **20**(15): p. 4064.
25. Yue, X., et al., *IESRGAN: enhanced U-net structured generative adversarial network for remote sensing image super-resolution reconstruction*. Remote Sensing, 2023. **15**(14): p. 3490.
26. Ma, X., et al., *HyADS: A Hybrid Lightweight Anomaly Detection Framework for Edge-Based Industrial Systems with Limited Data*. Electronics, 2025. **14**(11): p. 2250.
27. Zhai, G., et al., *A Sea-Surface Radar Target-Detection Method Based on an Improved U-Net and Its FPGA Implementation*. Electronics, 2025. **14**(10): p. 1944.
28. Joo, Y.H., et al., *Traffic Flow Speed Prediction in Overhead Transport Systems for Semiconductor Fabrication Using Dense-UNet*. Processes, 2022. **10**(8): p. 1580.
29. Taylor, H. and D. Boning. *Towards nanoimprint lithography-aware layout design checking*. in *Design for Manufacturability through Design-Process Integration IV*. 2010. SPIE.
30. Haas, S. and E. Hüllermeier, *Conformalized prescriptive machine learning for uncertainty-aware automated decision making: the case of goodwill requests*. International Journal of Data Science and Analytics, 2024.
31. Gawlikowski, J., et al., *A survey of uncertainty in deep neural networks*. Artificial Intelligence Review, 2023. **56**(Suppl 1): p. 1513-1589.
32. Ghanem, R., D. Higdon, and H. Owhadi, *Handbook of uncertainty quantification*. Vol. 6. 2017: Springer New York.
33. Ding, Y., et al. *Uncertainty-aware training of neural networks for selective medical image segmentation*. in *Medical Imaging with Deep Learning*. 2020. PMLR.
34. Dawood, T., et al., *Uncertainty aware training to improve deep learning model calibration for classification of cardiac MR images*. Medical Image Analysis, 2023. **88**: p. 102861.

35. Dawood, T., et al. *Uncertainty-aware training for cardiac resynchronisation therapy response prediction*. in *International Workshop on Statistical Atlases and Computational Models of the Heart*. 2021. Springer.
36. Abdar, M., et al., *A review of uncertainty quantification in deep learning: Techniques, applications and challenges*. *Information Fusion*, 2021. **76**: p. 243-297.
37. Palmer, G., et al., *Calibration after bootstrap for accurate uncertainty quantification in regression models*. *npj Computational Materials*, 2022. **8**(1): p. 115.
38. Ren, Y., et al., *System log detection model based on conformal prediction*. *Electronics*, 2020. **9**(2): p. 232.
39. Campos, M., et al., *Conformal prediction for natural language processing: A survey*. *Transactions of the Association for Computational Linguistics*, 2024. **12**: p. 1497-1516.
40. Zhou, X., et al., *Conformal prediction: A data perspective*. *ACM Computing Surveys*, 2025.
41. Romano, Y., E. Patterson, and E. Candes, *Conformalized quantile regression*. *Advances in neural information processing systems*, 2019. **32**.
42. Jensen, V., F.M. Bianchi, and S.N. Anfinson, *Ensemble conformalized quantile regression for probabilistic time series forecasting*. *IEEE Transactions on Neural Networks and Learning Systems*, 2022.
43. Che, L., C. Wu, and Y. Hou, *Large Language Model Text Adversarial Defense Method Based on Disturbance Detection and Error Correction*. *Electronics*, 2025. **14**(11): p. 2267.
44. Ren, M., et al. *Learning to reweight examples for robust deep learning*. in *International conference on machine learning*. 2018. PMLR.
45. Krishnan, R. and O. Tickoo, *Improving model calibration with accuracy versus uncertainty optimization*. *Advances in Neural Information Processing Systems*, 2020. **33**: p. 18237-18248.
46. Babin, S., et al., *Software tool for advanced Monte Carlo simulation of electron scattering in EBL and SEM: CHARIOT*. *Microolithography 2003*. Vol. 5037. 2003: SPIE.
47. Kimura, T., T. Matsunawa, and S. Mimotogi, *Compact resist model using single convolution kernel*. *SPIE Advanced Lithography*. Vol. 11327. 2020: SPIE.
48. Sethian, J.A. and D. Adalsteinsson, *An overview of level set methods for etching, deposition, and lithography development*. *IEEE Transactions on Semiconductor Manufacturing*, 1997. **10**(1): p. 167-184.
49. Imai, S.-i., H. Motomura, and K. Tachibana, *Photoresist selectivity mechanism in SiO₂ etching by inductively coupled plasma using fluorocarbon gases*. *Journal of Vacuum Science & Technology B: Microelectronics and Nanometer Structures Processing, Measurement, and Phenomena*, 2002. **20**(4): p. 1482-1488.
50. Ingram, S., *The influence of substrate topography on ion bombardment in plasma etching*. *Journal of applied physics*, 1990. **68**(2): p. 500-504.
51. Seeger, A. and H. Haußecker, *Shape-from-shading and simulation of SEM images using surface slope and curvature*. *Surface and Interface Analysis*, 2005. **37**: p. 927-938.
52. Walsh, J., et al., *Using U-Net network for efficient brain tumor segmentation in MRI images*. *Healthcare Analytics*, 2022. **2**: p. 100098.
53. Dawood, T., et al., *Uncertainty aware training to improve deep learning model calibration for classification of cardiac MR images*. *Medical Image Analysis*, 2023. **88**: p. 102861.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.