Convolutional Neural Network Classification for Femtosecond Laser Texturing

Cynthia AlLabaky and Roland Bejjani *

Department of Mechanical Engineering, Lebanese American University, Byblos, Lebanon
Email: labakycynthia@gmail.com (C.A.); rolandbejj@gmail.com (R.B.)
*Corresponding author

***Abstract*—Given the substantial progress in advanced manufacturing, reaching the micro-scale has become necessary in several industries, such as the sensor, semiconductor, and aerospace sectors. This requires effective methods to detect, classify, and determine the topological parameters of micro-scale features on a workpiece. In this paper, microstructures were imprinted using the femtosecond laser machine. To accurately detect and classify these microstructures, Artificial Intelligence (AI) techniques were employed. This study integrated Convolutional Neural Networks (CNNs), specifically GoogleNet and ResNet-50, to achieve this goal. Various microstructures, also known as dimples, were simulated and used as proof of concept to train the CNNs and determine appropriate input parameters, which were subsequently used for training on the experimental dataset. With varying initial learning rates and mini-batch sizes, GoogleNet and ResNet-50 achieved validation accuracy typically ranging from 80% to 100% across simulated textures, with higher consistency observed in round and oval geometries. In the experimental dataset, accuracies generally ranged between about 80% and 90% for GoogleNet, while ResNet-50 exhibited wider variations, approximately between 50% and 95% depending on processing conditions. Both networks demonstrated strong performance in distinguishing textures produced under different marking speeds and layer conditions, effectively identifying low and high-gradient patterns. The results, thus, indicate that CNN-based approaches can reliably classify femtosecond laser-induced micro-textures, with ResNet-50 providing more stable performance across varying conditions and GoogleNet offering faster yet slightly more variable predictions. This multi-stage approach provides a comprehensive assessment of CNN-based micro-texture analysis and introduces an effective framework for quality control in manufacturing.**

***Keywords*—femtosecond laser, convolutional neural networks, GoogleNet, ResNet-50**

I. Introduction

Laser ablation is recognized for its high flexibility and precision in generating surface textures. When compared to lasers with longer pulses, femtosecond lasers offer higher precision and minimal thermal damage, making them ideal for micro-scale surface texturing [1, 2].

Laser Surface Texturing (LST) creates arrayed micro-scale features, known as dimples, which enhance tribological properties such as reduced friction and improved wear resistance [3]. Dimples also improve hydrodynamic effects, increasing bearing capacity [4].

Micro-textures function as cavities and oil reservoirs to optimize engine efficiency by generating eddy currents that reduce friction. Fully textured circular dimples can reduce friction by up to 85%. In the case of elliptical dimples, adjusting their orientation provides a means to control the friction coefficient on steel surfaces [5].

Surface properties can further be tuned by controlling geometry, wavelength, and pitch, affecting micro- and nano-scale characteristics such as hydrophobicity [6].

LST has been applied with various patterns (crater arrays, ellipses, octagons) to improve surface preparation for welding and tribological applications, with morphological and elemental analysis confirming suitability for industrial application [7]. On tool rake faces, microchips and textured surfaces improved tribological behavior by reducing average shear strength, tool-chip contact length, and diffusion wear [8]. Femtosecond laser texturing is also used in solar cells to create shapes like cones and hemispheres that optimize light and absorption and minimize charge recombination [9]. In biomedical applications, surface patterns are created on stainless steel to control cell adhesion and migration, depending on the geometry and orientation of these patterns [10].

At the micro-scale, femtosecond laser texturing produces surface patterns highly sensitive to subtle geometric variations, edge definition, and depth gradient, which are often missed by conventional inspection methods due to irregularities and noise. In contrast, deep learning has achieved state-of-the-art performance in surface defect detection, with nearly 49% of studies focusing on image-based classification tasks [11].

To effectively detect micro-textures produced by femtosecond lasers, transitioning to a "Smart Factory" is essential to modernize traditional quality control methods and integrate AI. This shift enables improved performance, higher efficiency, and enhanced process longevity through predictive maintenance, real-time monitoring, and adaptive inference [12, 13].

Manuscript received November 21, 2025; revised December 26, 2025; accepted April 2, 2026; published June 15, 2026.

In semiconductor manufacturing, where processes are influenced by time-variant behavior and interdependent parameters, CNNs, ResNet, and LSTM architectures have been applied for wafer defect detection. Among them, ResNet achieved the highest accuracy of around 97% due to its ability to handle complex surface features more effectively [14]. CNNs have also been applied to metal surface inspection, machining quality assessment, and textured material evaluation, consistently achieving high classification performance, often exceeding 90% under controlled conditions [15].

Similarly, deep learning networks, such as VGGNet, ResNet, DenseNet, and GoogleNet, have been used to classify electric wire defects. DenseNet achieved the best performance, followed by GoogleNet. Moreover, applying Region of Interest (ROI) extraction through image processing improved the performance of all pre-trained models [16]. ResNet-50 has been successfully employed in industrial quality control for welding defect detection and surface defect classification. Its residual learning structure enables stable training and effective handling of complex surface textures, reaching a classification accuracy of 96% [17].

Several studies have compared deep learning architectures for industrial surface inspection using transfer learning, which consistently outperformed self-built networks. Multi-modal inspection strategies further improved defect detection without added computational complexity. Among these, ResNet-50 achieved the highest classification accuracy for real-time industrial quality control systems [18].

Similarly, models such as VGG16, AlexNet, GoogleNet, and ResNet have been used to detect defects in laser-directed energy deposition. Applying data augmentation improved performance, yielding a maximum accuracy of 94.7% for VGG-16 [19].

In another experiment, sandblasting defects in investment casting were detected using AlexNet, VGG-16, GoogleNet, and ResNet-34. Image pixel size and mini batch size during training were varied, resulting in optimized computational speed and flexibility [20].

In literature, traditional image processing methods have been compared with deep learning approaches. Although image processing can perform adequately under low variability or clearly defined defect conditions, deep learning methods are generally more effective at capturing complex patterns and maintaining high detection accuracy [21]. In femtosecond laser texturing, subtle geometric variations and irregular edge profiles are prevalent, making CNNs more effective for defect detection in this context.

CNN-based approaches have previously been used for surface texture analysis and defect detection; however, their application to femtosecond laser-textured micro-scale surfaces remains limited. This study addresses this gap by applying deep learning models to detect and classify defects in micro-textures with feature sizes down to 50 μm, where subtle variations in edge profiles, marking speeds, number of layers, and jump delays create complex features that are difficult to identify using conventional methods. Furthermore, the selected models are designed to be efficiently deployed on conventional commercial PCs, enabling practical and cost-effective implementation without the need for high-end computational resources.

This paper presents the application of deep learning for defect detection in femtosecond laser-textured surfaces. The background section reviews GoogleNet's and ResNet's architectures, followed by a methodology section describing the textures' simulations and laser processing for experimental design. The results and discussion section evaluates model performance, including the impact of data augmentation and comparisons between the simulated and experimental training results. Finally, conclusions are drawn regarding the effectiveness of the proposed approach which provides a foundation for future industrial implementation of femtosecond laser texturing processes.

II. Background

A. Architecture

CNNs are categorized under deep learning algorithms. They are known for their capacity to classify images based on the dominating features and adaptively learn spatial features from the input image. The main layers are:

Convolutional Layers: A kernel with adjustable weights moves across the image, performing element-by-element matrix multiplication between the filter and pixel values. The computed values for each feature are then summed and stored in a feature map. If the values are high, the feature is considered prominent, whereas lower values indicate that the feature is either weak or absent.

Rectified Linear Unit (ReLU): This function enhances the network's ability to learn complex patterns by setting all negative values in the feature map to zero. By doing so, it ensures that the network focuses on the most significant features while preventing overfitting.

Pooling Layer: Responsible for downscaling the feature map, this layer maintains essential features while decreasing computational complexity. It plays a crucial role in retaining essential details and preventing overfitting.

SoftMax Layer: This layer assigns probabilities to different categories based on the input image, with the category having the highest probability selected as the final classification.

B. GoogleNet

GoogleNet, a deep convolutional neural network, was first introduced by Google in 2014. It is best known for its inception module, which processes features at multiple scales by combining convolutional layers with different kernel sizes (1×1, 3×3, 5×5) and pooling layers in parallel. With a total of 22 layers, GoogleNet improves computational efficiency through 1×1 convolutions, which optimize performance while reducing the number of parameters. When compared to AlexNet and custom CNNs, GoogleNet has demonstrated greater accuracy, achieving 0.95 for grayscale images and 0.96 for

RGB [22]. Fig. 1 shows its architecture, known for these inception layers.

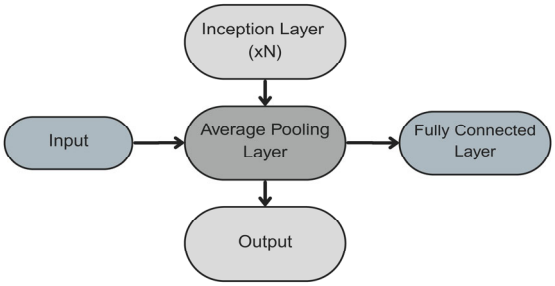

Fig. 1. GoogleNet architecture.

The following equation describes the multi-branch architecture, which processes multiple filter sizes in parallel and concatenates their outputs:

$$Y = Concat(f_{1\times1}(X), f_{3\times3}\left(f_{1\times1}(X)\right), (f_{5\times5}(f_{1\times1}(X), f_{1\times1}(Pool_{3\times3}(X)))$$

where X is the input feature map; $f_{k\times k}$ denotes a convolution with $k\times k$ filters; *Pool* 3×3 is a 3×3 max pooling; *Concat* represents depth-size concatenation; and Y is the output feature map of the inception module.

C. ResNet-50

ResNet is a different type of CNN, which is known for solving the problem of vanishing gradients that have been prevalent in other networks. Gradients are values used to update weights during training. To modify such weights, the error is propagated backward from the output layer to the preceding layers. When gradients become minimal, earlier layers fail to receive sufficient weight updates, causing them to learn very slowly or not at all. This issue is particularly critical in deep networks with many layers, as gradients tend to diminish progressively during backpropagation. To address this problem, ResNet introduces a variable x that bypasses the weight layers, enabling more effective gradient flow [23].

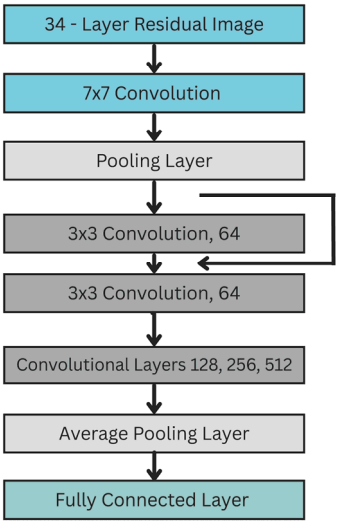

Fig. 2. ResNet architecture.

In training, ResNet-50 was chosen because it has a moderate number of layers (50 layers). This decision was based on the relationship between training errors and test errors. Training errors occur within the training data set, whereas test errors indicate the AI's performance on unseen data after training. Ideally, both errors should be similar; otherwise, overfitting is occurring. As the number of layers increases, errors tend to rise; however, deeper networks allow for better feature extraction and more thorough image detection. ResNet-50 strikes a balance by minimizing errors while still achieving effective feature detection. Fig. 2 shows the ResNet architecture, which portrays the bypassed layer.

This equation illustrates the residual connection, which facilitates gradient flow during training, enabling effective training of deeper networks:

$$Y = F(X, \{W_i\}) + X$$

where X denotes the input to the residual block, which is subsequently added to the final output through the skip connection; $F(X, \{W_i\})$ is the residual function typically composed of 2 or 3 convolutional layers; and Y is the output of the residual block.

III. Methodology

Using pre-trained networks in MATLAB, the performance of GoogleNet and ResNet was evaluated by categorizing and identifying various imprinted micro-features commonly encountered in manufacturing applications. All experiments and neural network training were performed using MATLAB R2022b with an Intel Core i7-10510U CPU (4 cores, 8 threads, 1.8 GHz base frequency), 20 GB RAM, and Windows 11 Pro. Extensive femtosecond laser experiments were first conducted by varying laser input parameters such as marking speeds, number of layers, and laser jump delays. In these experiments, recurring errors were observed, including deformation of round features into oval shapes, square features into rectangles, and edge disruptions. To better understand and control these dominant errors before applying artificial intelligence, a simulation framework was introduced to replicate similar geometric distortions and edge irregularities in a controlled manner. For example, experimentally observed edge irregularities such as slight elongation of circular dimples, geometric deformation of square features, and gradual fading at boundaries were reproduced in the simulations by introducing corresponding geometric distortions and gradient variations at the edges to mimic edge fading. This allowed systematic variations of the input datasets, which is difficult to achieve using solely experimental data. The simulated dataset was therefore used as a proof-of-concept and as a parameter screening step before applying the selected training settings to the experimental dataset.

These micro-textures were simulated by using drawing software to generate synthetic images and apply gradient-based attribution methods. In contrast to experimental images obtained from femtosecond laser

imprinting, which often contain noise, artifacts, and surface irregularities, simulated images provide idealized, noise-free representations of the target geometries. This enabled a clearer evaluation of the networks' ability to recognize and differentiate shapes under ideal conditions. Based on the simulation results, the optimal input parameters were then selected for testing on the experimental dataset. Fig. 3 shows the different shapes that were trained, consisting of macro shapes, which include circular versus oval and square versus rectangular shapes, and gradient shapes.

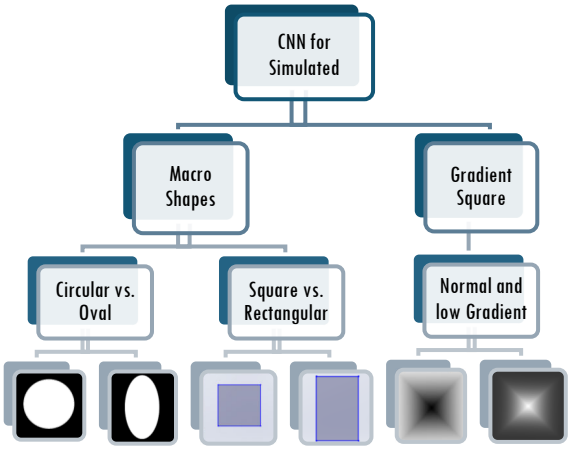

Fig. 3. Simulated micro-texture shapes.

The second phase consisted of optimizing femtosecond laser parameters in a micro-texturing process to get accurate textures with the required dimensions. Illumination conditions were controlled to ensure consistent lighting and reduce imaging-related variability in the experimental dataset. The imaging system used a USB3.0 industrial camera with a 6.3 MP Sony IMX178 sensor, color CMOS, and 60 fps, connected via USB and mounted on a tripod plate. A C-mount 25 mm lens with manual iris F2.8 was used, along with four 1 mm and four 5 mm extension rings to achieve a working distance of 30 mm. Illumination was provided by a white LED bar light (41×16 mm, 1.3 W) positioned to enhance surface feature visibility. The models were trained and tested under controlled imaging conditions; however, in industrial deployment, they may be sensitive to changes in illumination. The experimental conditions should be carefully controlled to minimize variability; otherwise, corrective measures such as system calibration or data augmentation with varied lighting can be employed to maintain robust performance.

Laser parameters such as power and frequency, for example, have a significant effect on the dimensions and depth of the micro-dimple [24]. Other parameters such as laser scanning speed and laser scan times can influence the width and the depth of the micro-texture in the case of linear grooves [25]. The geometrical characteristics of laser-induced micro-textures can affect their performance in their intended application, such as friction reduction [26, 27], necessitating optimizing laser parameters to generate micro-textures with the desired geometrical characteristics.

The laser micro-texturing process was done using a femtosecond laser (HTJ-10F Femtosecond-Hitec) with the following properties listed in Table I:

TABLE I. FEMTOSECOND LASER PARAMETERS

Feature	Value
Maximum Laser Power	10.3
Wavelength (nm)	1033
Repetition Rate (MHz)	1
Single Pulse Energy (μJ)	10.3
Beam Quality	1.15/1.09
Pulse Width (fs)	285
Spot Roundness (%)	95
Spot Diameter (μm)	13 (approx.)

The micro-textures were generated on a 100×100 mm aluminum sheet sample. Several important laser parameters significantly affect the quality and accuracy of the engraved micro-textures. A notable parameter is the laser jump delay, measured in μs, which has received limited attention in the literature. In laser marking machines, Galvano mirrors are used to quickly steer the laser beam across the workpiece area. The jump delay is the duration given to the Galvano mirrors to rotate so that the laser can "jump" from one marking position to another, avoiding Galvano mirror hysteresis. A relatively low jump delay results in the laser marking while jumping, causing residual traces outside the desired marking area. Consequently, low jump delay affects the geometrical borders of the micro-textures, making it necessary to optimize this parameter to obtain accurate edges [28].

Other critical parameters are the marking speed and the number of layers, both of which influence micro-texture depth. The marking speed, measured in mm/s, corresponds to the speed of the Galvano mirrors. A low marking speed improves edge accuracy but can increase depth beyond desired values and may cause heat damage, whereas a high marking speed does not allow sufficient time to engrave the intended depth, producing a non-uniform height gradient. Similarly, the number of layers indicates how many times the laser passes over the 3D model of the micro-texture. A low number of layers produces shallow micro-textures with depth gradients, while a high number of layers results in deeper engraving. Optimizing these parameters is therefore essential to achieve geometrically and structurally accurate micro-textures.

Based on these considerations, the following parameters were selected for texturing. A jump delay of 100 μs was assigned to defective dimples, while 500 μs was used for acceptable dimples, capturing edge transitions accurately. With respect to marking speed and the number of layers, defective dimples were processed with 2 layers at 8000 mm/s, and ideal dimples with 20 layers at 50 mm/s, enhancing gradient differentiation. Round and square dimples were textured with varied parameters to produce different edge qualities and depth gradients, as shown in Fig. 4. Fig. 4(a) and (b) depict micro-dimples at jump delays of 500 μs and 100 μs, respectively, illustrating edge defects, while Fig. 4(c) and (d) show micro-dimples at 20

layers with 50 mm/s and 2 layers with 8000 mm/s, respectively, demonstrating depth gradient defects.

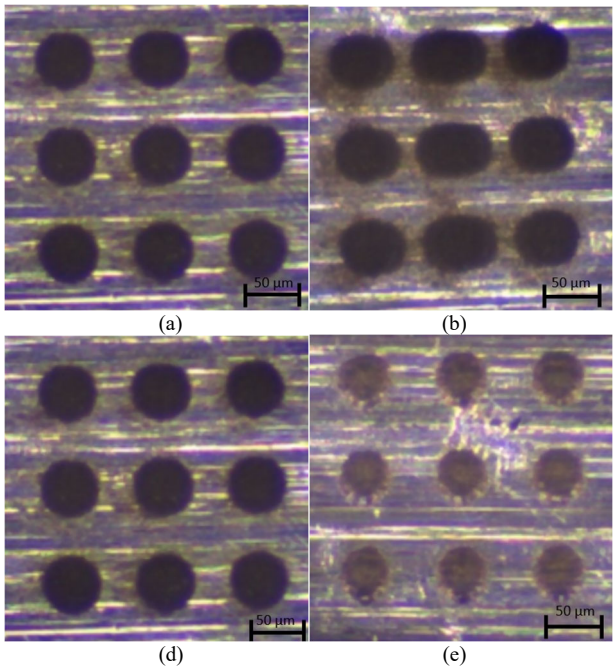

Fig. 4. Generated circular micro-textures at a (a) jump delay of 500 us; (b) jump delay of 100 us; (c) 20 layers with a scanning speed of 50 mm/s; (d) 2 layers with a scanning speed of 8000 mm/s.

The input dataset for each category consisted of 100 images. Although image augmentation was applied to increase the dataset to 540 images through rotations, translations, scaling, and shearing, only marginal changes in performance were observed; therefore, the original dataset of 100 images was retained for training.

During training, 15% of the images were used for validation, while 10% were set aside for testing to generate the confusion matrix.

To reduce potential sampling bias associated with the limited dataset size, the validation dataset was randomized in MATLAB for each training round. This ensured that different image subsets were used during validation across the training runs. The reported performance metrics therefore correspond to the average results obtained across these runs, providing a more robust evaluation of the model performance.

During training, multiple options directly impact the validation process and overall accuracy. The selected solver is Stochastic Gradient Descent with Momentum (SGDM), which updates the model's parameters using a gradient loss function applied to a mini batch, representing the subset selected in each epoch for training. Compared to standard SGD, SGDM includes a momentum term that accelerates training while reducing oscillations. The Initial Learning Rate determines how much the model adjusts after each iteration based on the estimated error and corresponding weight updates. Validation frequency defines how often the training is evaluated against the validation set, while the max epoch indicates the number of complete passes the CNN makes through the entire training dataset. Finally, the mini-batch size specifies the number of images processed in a single iteration.

After testing on the simulated images, these training options were selected and varied in the subsequent experiments to compare the performance of ResNet and GoogleNet. The primary objective of training was to evaluate the networks under different initial learning rates and mini-batch sizes, with input options determined based on the values outlined in Table II. Six rounds of experiments were conducted, each varying a single parameter at a time to better understand the underlying relationships and their impact on model performance.

TABLE II. SOLVER INPUT PARAMETERS

Rounds	Solver	Initial Learning Rate	Validation Frequency	Max Epochs	Mini Batch Size
1	SGDM	0.0001	5	8	11
2		0.0001			8
3		0.0001			4
4		0.001			11
5		0.001			8
6		0.001			4

The number of maximum epochs was chosen based on preliminary tests to balance convergence and training efficiency. Early stopping was implemented by monitoring the validation loss, with training terminated when no further reduction was observed over two consecutive epochs. Since most networks converged around the sixth epoch, training was extended to the eighth epoch to ensure stability and convergence in all cases.

Accuracy indicators used for evaluation include validation accuracy, which was obtained directly from MATLAB, as well as precision, recall, and F1-Scores, which were calculated using the following equations:

$$Precision = \frac{TP}{TP + FP}$$

$$Recall = \frac{TP}{TP + FN}$$

$$F1 - Scores = \frac{2 \times Precision \times Recall}{Precision + Recall}$$

Image acquisition and dataset preparation ensured a consistent number of micro-textures per image (a 4×4 array) to avoid bias during training. Image preprocessing was handled directly by CNN architectures. Input images were resized to the required network dimensions (224 × 224 pixels for both GoogleNet and ResNet-50) and normalized according to the standard procedures of the

pre-trained MATLAB architectures, including pixel scaling and channel-wise normalization.

IV. Results and Discussion

A. Simulated Micro-Textures

1) Training for round and oval micro-features detection

Round and oval dimples, widely used as industrial micro-textures, were categorized by GoogleNet and ResNet according to the training input parameters. The results are demonstrated in Fig. 5. Both GoogleNet and ResNet-50 demonstrated high validation accuracy across all rounds, with values ranging from approximately 92% to 100%.

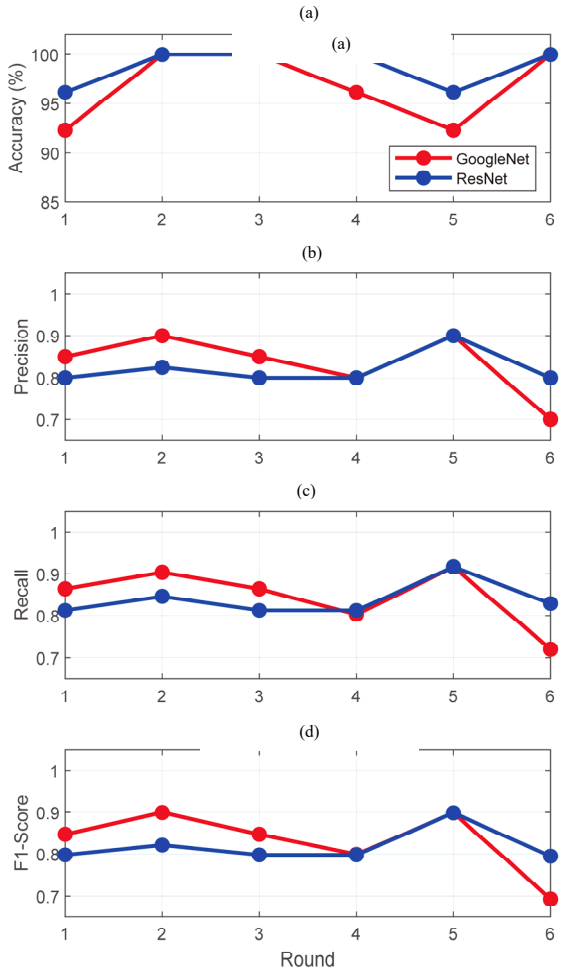

Fig. 5. Simulated round & oval training results. (a) validation accuracy; (b) precision; (c) recall; and (d) F1-Scores.

During training, ResNet consistently outperformed GoogleNet, suggesting greater robustness in classification. Precision and F1-Scores showed a similar trend: ResNet achieved more consistent values (0.8–0.9) compared with GoogleNet (0.7–0.9), indicating better overall performance. For recall, both networks effectively captured round and oval shapes, with ResNet demonstrating greater stability and consistency.

Regarding the effect of input parameters, both models exhibited stable performance at a lower learning rate of 0.0001, whereas a higher rate of 0.001 introduced inconsistencies, particularly in GoogleNet. This behavior is consistent with the observation that smaller learning rates generally require longer training but promote stability. Similarly, reducing the mini-batch size maintained ResNet's performance while slightly decreasing stability in GoogleNet. In the latter, smaller batch sizes likely increased gradient variance and the risk of overfitting, resulting in lower precision and accuracy.

Finally, training times showed that GoogleNet was consistently faster than ResNet across all rounds, completing each session in 5–8 min, whereas ResNet required 13–20 min, making GoogleNet more suitable for time-sensitive training while ResNet prioritized evaluation accuracy.

2) Image augmentation

Image augmentation was applied to evaluate its effect on classification accuracy. The input dataset was augmented using MATLAB, increasing the number of training images from 90 to 540 and the validation images from 10 to 60. The augmentation process included rotation, translation, scaling, and shearing to expand and diversify the dataset. This training was done on the first three rounds of GoogleNet to check the percentage difference with a bigger input batch size. The percentage change is the difference in validation accuracy between a normal and augmented batch size. The results can be observed in Table III. Performance variation was minimal, with percentage differences ranging from 1.9% to about 5% across the evaluated iterations. The training time, on the other hand, is amplified by around 12 folds. Therefore, all subsequent analyses used 100 images per class, since this dataset size was sufficient to capture the dominant geometric features and ensure stable training. This approach served as proof of concept, forming the basis for the integration of higher-specification workstations and the implementation of real-time monitoring under industrial conditions.

TABLE III. Augmented Sample-Training Results

Network	Validation Accuracy (%)	Precision	Recall	F1-Score	Percentage Change (%)	Training Time (min)
GoogleNet	100	0.892	0.897	0.893	5.31	59
	98.77	0.883	0.885	0.883	1.86	57
	99.38	0.867	0.87	0.866	2.12	89

3) Training for simulated square and rectangular micro-features detection

Other commonly used micro-feature shapes in modern industries are squares and rectangles. Detecting and classifying these shapes is crucial for advancing Industry 4.0 applications.

The results of the batches trained can be observed in Fig. 6. For validation accuracy, GoogleNet ranged from

80.77% to 100%, reaching the highest peak, while ResNet ranged from 84.62% to 96.15%, demonstrating more stable and consistent performance. Across the remaining metrics, GoogleNet achieved higher peak values but exhibited greater variability, whereas ResNet, despite slightly lower peak results, maintained greater consistency. At a lower initial learning rate of 0.0001, ResNet's performance plateaued, while GoogleNet reached higher peak accuracies at a learning rate of 0.001.

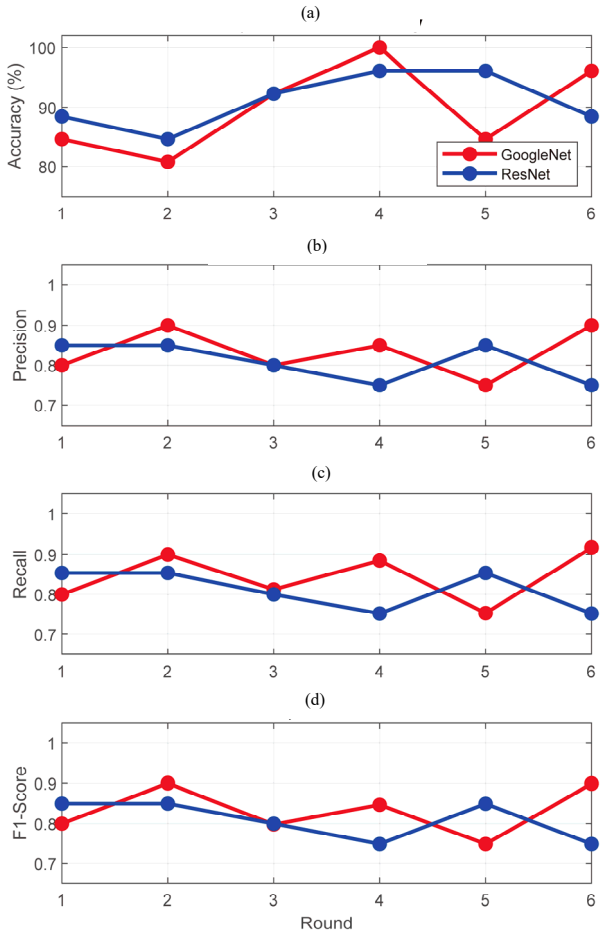

Fig. 6. Simulated square and rectangle training results. (a) Validation accuracy; (b) Precision; (c) F1-Scores; and (d) Recall.

4) *Gradient shape detection*

Once macro-shape detection was evaluated, the networks' performance was further analyzed for gradient-based feature detection. Gradient shapes were used to represent variations in edge definition, as light reflection depends on edge sharpness and contrast, influencing how the network perceives and classifies the dimples. A change in gradient shape is represented as a change in the number of layers and marking speed with the femtosecond laser samples. The two chosen colors were the following: one having a white core and the other having a black core, as seen in Fig. 7.

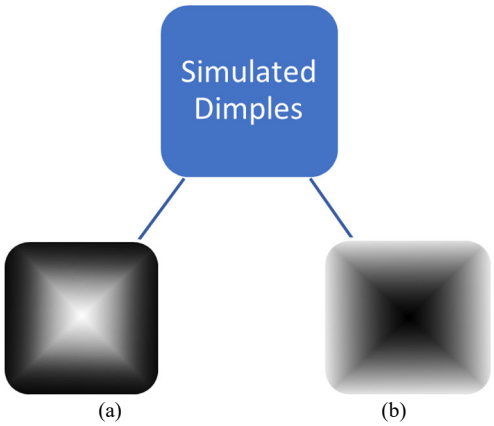

Fig. 7. Simulated gradient dimples. (a) white core; (b) black core.

Both GoogleNet and ResNet were able to reach 100% accuracy, with precision, recall, and F1-Scores also equal to 1, across all training runs with varying initial learning rates and mini-batch sizes. This is due to the binary classification, which is considered straightforward in detecting high-contrast features (white or black core). The gradient shapes were considered easier to detect when compared to the previous complex features.

B. *Femtosecond Laser Textures*

In the experimental datasets, both GoogleNet and ResNet-50 were processed using the same procedure. The initial learning rate and mini-batch sizes were altered to test the accuracy under varied conditions.

The round and square dimples were designed in SolidWorks and imprinted on metal sheets with variations in jump delay, number of layers, and marking speeds. The convolutional networks were trained on various shapes, including round dimples (Fig. 8) and square dimples (Fig. 9). Subsequently, a selection process was applied to fine-tune the training inputs and identify both ideal and defective dimples.

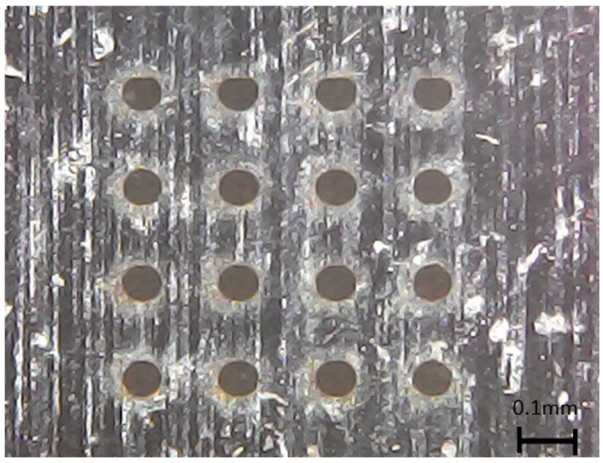

Fig. 8. Round dimples array.

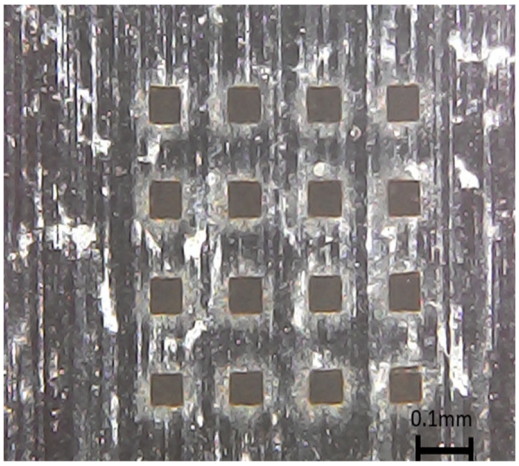

Fig. 9. Square dimples array.

1) *Training roundness for different jump delays*

Different jump delays directly affected the shape of the micro-textures as observed in Fig. 10. Fig. 10(a) represents a jump delay of 500, yielding defined edges, while Fig. 10(b) represents a jump delay of 100, yielding defective shapes.

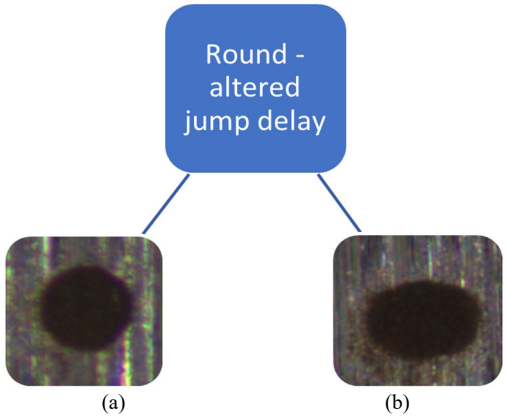

Fig. 10. Round dimples with altered jump delay. (a) ideal dimple and (b) defective dimple.

The training results are portrayed in Fig. 11 with respect to the validation accuracy, precision, recall, and F1-Scores.

Starting with the validation accuracy, ResNet's performance was low with an Initial Learning Rate of 0.0001 but improved steadily across training rounds, peaking at Round 5. Despite the lower learning rate, ResNet demonstrated strong adaptability to varying training parameters. In contrast, GoogleNet maintained consistent accuracy across all parameter changes, which may be advantageous if predictability and stability are preferred over high variance.

Precision, recall, and F1-Scores achieved similar results across the testing parameters for both networks. ResNet showed fluctuations, achieving higher values in certain rounds, while GoogleNet remained relatively stable throughout.

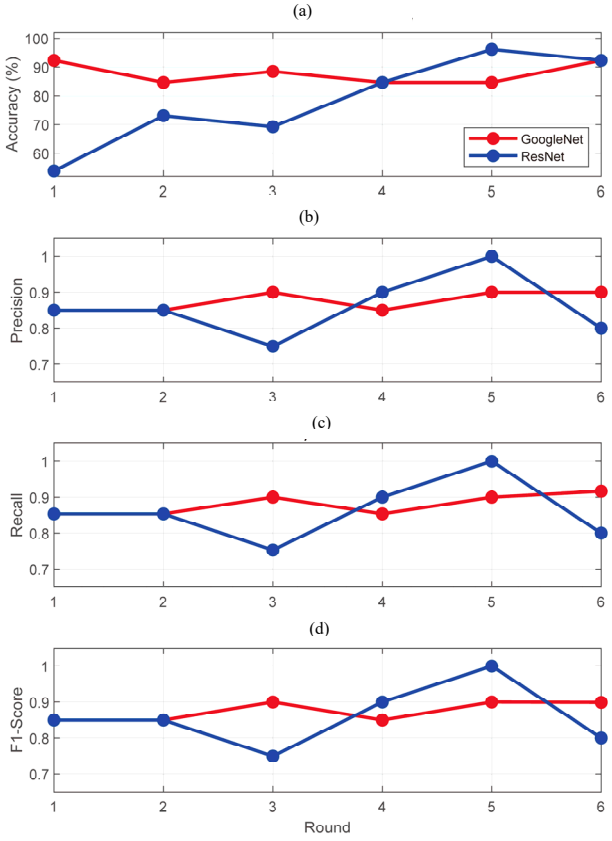

Fig. 11. Round dimples with altered jump delay training results: (a) validation accuracy; (b) precision; (c) recall; and (d) F1-Scores.

2) *Training square dimples for different jump delays*

Similarly, the different jump delays affected the edges of square dimples, as observed in Fig. 12.

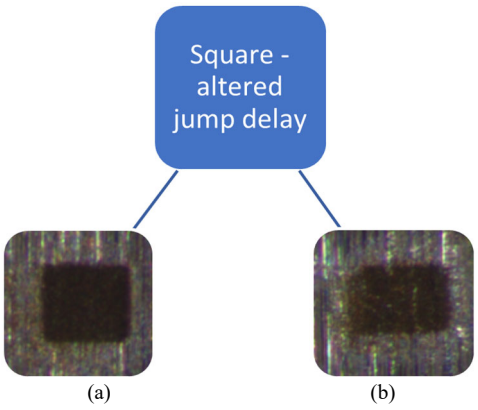

Fig. 12. Square dimples with altered jump delay. (a) ideal dimple; and (b) defective dimple.

The training results are depicted in Fig. 13:

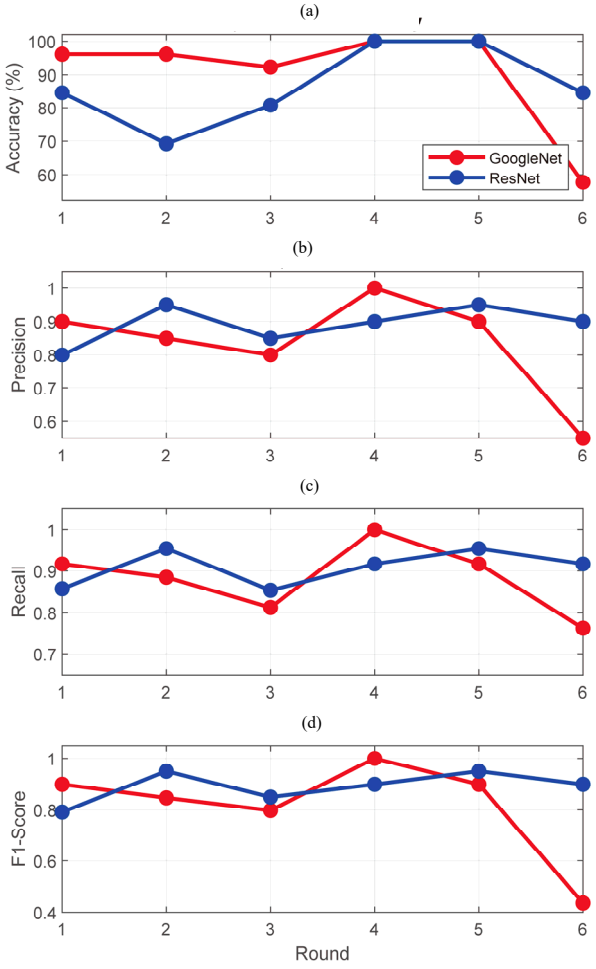

Fig. 13. Square dimples with altered jump delay training results (a) validation accuracy; (b) precision; (c) recall; and (d) F1-Scores.

When trained on square dimples, GoogleNet exhibited strong learning and maintained high accuracy, except in the final round, where a low mini-batch size of 4 caused some instability. ResNet-50 performed better on oval dimples compared to square dimples. While its validation accuracy fluctuated, overall performance during testing, measured by precision, recall, and F1-Scores, remained stable.

3) Training square dimples for number of layers and marking speeds

To test the performance of GoogleNet and ResNet for different gradients, the layers and marking speeds were varied on the femtosecond laser machine. Fig. 14 shows the difference between them. In Fig. 14(a), 20 layers and a speed of 50 mm/s were used, whereas in the low-gradient dimple shown in Fig. 14(b), only 2 layers and a higher marking speed of 8000 mm/s were employed to imprint these micro-textures.

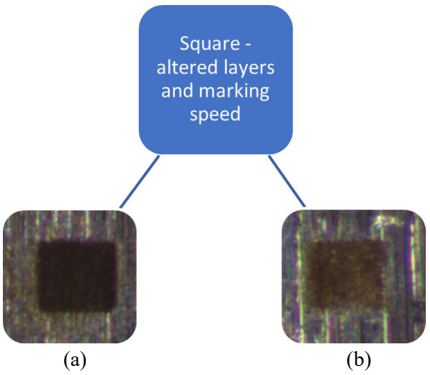

Fig. 14. Square dimples with altered layers and marking speed. (a) ideal dimple; and (b) defective dimple.

The training results are tabulated in Table IV.

TABLE IV. GRADIENT SQUARE DIMPLES TRAINING RESULTS

Round	GoogleNet		ResNet-50	
	Validation Accuracy (%)	Precision, Recall, & F1-Scores	Validation Accuracy (%)	Precision, Recall, & F1-Scores
1	95.15	1	100.00	1
2	96.15	1	96.15	1
3	96.15	1	96.15	1
4	100.00	1	100.00	1
5	96.15	1	100.00	0.95
6	92.31	1	100.00	1

Similar to the results of the simulated dimples, both networks were able to accurately differentiate low and high gradient images. Although ResNet achieved a slightly higher overall validation accuracy, it was more time-consuming and consumed more CPU resources. This is an unnecessary trade-off, given that GoogleNet delivered strong and reliable results.

In simulated micro-features, ResNet outperformed GoogleNet through more consistent and robust classification. While GoogleNet offered faster training (5–8 mins vs 11–24 mins for ResNet), ResNet delivered superior validation accuracy, precision, and F1-Scores, especially with round/oval shapes. ResNet maintained stable performance across different learning rates and mini-batch sizes, whereas GoogleNet showed inconsistency at higher learning rates (0.001). The simulations enabled the identification of effective training parameters, which were subsequently validated through experiments, confirming their suitability for accurate classification. ResNet's architecture prioritizes accuracy at the expense of training time, making it the preferred choice for simulated dimple classification where precision matters more than speed.

The evaluation on experimental femtosecond laser-imprinted images revealed that the networks' performance differed across texture types and experimental settings. With round dimples at different jump delays, ResNet showed significant improvement

over training rounds despite initial low accuracy (around 53.3%), while GoogleNet maintained consistency but with lower peak performance. When training square dimples, GoogleNet exhibited exceptional learning and high accuracy except when using smaller mini-batch sizes. Both networks achieved near-perfect results (96–100% accuracy) in gradient detection across both simulated and experimental datasets. GoogleNet requires fewer computational resources than ResNet, making it more efficient when the two networks deliver comparable performance. Indeed, average training times across the different texture types indicate that ResNet required approximately 2.5–3 times longer than GoogleNet.

This difference in computational efficiency could be further explained by the respective parameter counts and Floating Point Operations (FLOPs) of the two architectures. GoogleNet contains 5.9 million parameters and requires around 1.8 billion FLOPs per forward pass, whereas ResNet-50 has 23.5 million parameters and requires approximately 3 billion FLOPs.

A statistical comparison between simulated and experimental datasets is presented in Table V using the mean validation accuracy and the corresponding percentage error. Square textures showed small deviations, with percentage errors of approximately 0.8% for GoogleNet and 4.9% for ResNet, indicating strong agreement between simulated and experimental results. Round textures exhibited larger deviations (9.3% for GoogleNet and 20.8% for ResNet), which can be attributed to the higher sensitivity of circular geometries to small variations in edge distortion and marking conditions. Despite these variations, the training parameters identified from the simulations enabled both networks to maintain robust classification performance when applied to the experimental dataset. Gradient textures showed negligible deviation, with both datasets achieving approximately 100% accuracy.

TABLE V. PERCENTAGE ERROR OF MEAN VALIDATION ACCURACY FOR SIMULATED AND EXPERIMENTAL DATASETS

Network	Case	Simulated Mean (%)	Experimental Mean (%)	Error (%)
GoogleNet	Round Textures	96.8	87.8	9.3
	Square Textures	89.7	90.4	0.8
ResNet-50	Round Textures	98.7	78.2	20.8
	Square Textures	91.0	86.5	4.9
Both	Gradient	~100	~100	~0

The collective findings from the simulated and experimental datasets demonstrate that CNN-based models can reliably classify femtosecond laser textures, validating their suitability for automated inspection of micro-features. Given the difficulty of detecting these features with conventional inspection methods, CNNs provided an effective approach for identifying errors in shape, gradient, and edge profiles. Building on the training procedures described in this study, future work will focus on real-time monitoring, enabling the detection of surface irregularities during the process and allowing immediate corrective action.

This study demonstrates the application of CNN-based models to femtosecond laser-textured surfaces, extending the use of deep learning to a domain with unique micro-scale features and processing variations. While previous studies have applied CNNs to surface defect detection in wafer [14], welding [17], and casting [20] processes, the present work builds on these approaches by addressing challenges specific to laser-textured surfaces and validating model performance under experimental conditions. By providing a framework capable of classifying complex micro-textures, this study highlights potential applications in real-time monitoring and inline quality control, illustrating how CNNs can support automated inspection workflows in industrial laser processing.

V. CONCLUSION

The novelty of this work lies in integrating femtosecond laser micro-texturing with convolutional neural networks for the detection and classification of micro-scale surface textures. The proposed framework enabled robust identification of key femtosecond laser-induced textures and demonstrated how simulated data can enhance experimental training by providing controlled insight into parameter selection. This approach established a foundation for future industrial quality control applications, where real-time monitoring of micro-features may improve manufacturing precision and operational efficiency.

Different neural networks, namely GoogleNet and ResNet, were trained to analyze their performance. Textures were both simulated and imprinted at the micro-scale using a femtosecond laser, with variations in marking speed, number of layers, and jump delay. Training parameters were then adjusted to evaluate the accuracy and behavior of each network, highlighting the role of AI in advanced manufacturing techniques and quality control in high-tech industries.

Based on the training of simulated and experimental micro-textures, the networks demonstrated the following results:

- Both GoogleNet and ResNet-50 demonstrated relatively high validation accuracy across most tests.
- ResNet-50 exhibited more consistent and stable performance across different micro-texture shapes (round, oval, square, and gradient), particularly in terms of precision, recall, and F1-Score.
- GoogleNet occasionally reached higher peak accuracy but showed more variability, especially under higher learning rates or smaller mini-batch sizes.
- The smaller size and lower computational demand of GoogleNet contributes to its reduced training and inference times, while ResNet-50's deeper architecture

enables potentially higher feature representation under complex conditions at the cost of increased computational load.

- Data augmentation increased training images six-fold but yielded only marginal accuracy changes (~1.9% to 5.3%), while increasing training time by a factor of approximately 12.
- Binary classification of gradient features was highly successful for both networks, achieving 100% accuracy due to the high-contrast nature of the task.
- When applied to experimental femtosecond laser images, ResNet-50 initially underperformed but improved with tuning, while GoogleNet maintained a more consistent performance across all parameter variations (jump delay, layers, and marking speeds).

The results indicate that ResNet-50 exhibits greater robustness and accuracy under complex conditions, whereas GoogleNet offers higher efficiency and reliability for time-sensitive applications. Future work will focus on optimizing a selected network architecture, expanding the dataset to include a wider range of surface patterns, investigating advanced feature extraction techniques, and exploring models suitable for real-time inspection. In addition, domain adaptation strategies will be explored to improve generalization across experimental variations, and mitigation approaches will be applied to enhance model reliability.

These networks support the smooth transition to Industry 4.0, particularly in industries involving close-contact components such as automotive, sensor, and semiconductor industries. To ensure increased production efficiency, enhanced performance, and longevity of a mechanical process, the integration of AI for quality control and real-time monitoring is paramount. The CNNs will finally be able to categorize different textures to provide adaptable solutions for manufacturing companies.

Conflict of Interest

The authors declare no conflict of interest.

Author Contributions

CA conducted the experimental work related to femtosecond laser texturing, performed the AI modeling and data analysis, and wrote the original manuscript draft. RB supervised the research activities, contributed to the methodology and interpretation of results, and reviewed and revised the manuscript. All authors approved the final version of the manuscript.

References

[1] A. Ancona, S. Doring, C. Jauregui *et al.,* "Femtosecond and picosecond laser drilling of metals at high repetition rates and average powers," *Opt. Lett.,* vol. 34, no. 21, pp. 3304–3306, 2009. https://doi.org/10.1364/OL.34.003304

[2] J. Zhang, D. Yang, A. Rosenkranz *et al.,* "Laser surface texturing of stainless steel—effect of pulse duration on texture's morphology and frictional response," *Adv. Eng. Mater.*, vol. 21, no. 3, 1801016, 2019. https://doi.org/10.1002/adem.201801016

[3] I. Etsion, *Handbook of Lubrication and Tribology, Vol. II: Theory and Design*, 2nd ed. Boca Raton: CRC Press, 2012, pp. 1–16.

[4] Z. Wu, H. Bao, Y. Xing *et al.,* "Tribological characteristics and advanced processing methods of textured surfaces: A review," *Int. J. Adv. Manuf. Technol.,* vol. 114. no. 5, pp. 1241–1277, 2021. https://doi.org/10.1007/s00170-021-06954-2

[5] A. Ancona, G. Carbone, M. De Filippis, "Femtosecond laser full and partial texturing of steel surfaces to reduce friction in lubricated contact," *Adv. Opt. Technol.,* vol. 3, no. 5–6, pp. 539–547, 2014. https://doi.org/10.1515/aot-2014-0045

[6] T. Aizawa, T. Inohara, K. Wasa, "Femtosecond laser micro-/nano-texturing of stainless steels for surface property control," *Micromachines*, vol. 10, no. 8, 512, 2019. https://doi.org/10.3390/mi10080512

[7] E. R. Moldovan, C. D. Cimpoesu, J. L. Ocaña *et al.,* "Morphological analysis of laser surface texturing effect on AISI 430 stainless steel," *Materials,* vol. 15, no. 13, 4580, 2022. https://doi.org/10.3390/ma15134580

[8] J. Zhang, H. Yang, S. Chen *et al.,* "Study on the influence of micro-textures on wear mechanism of cemented carbide tools," *Int. J. Adv. Manuf. Technol.*, vol. 108, pp. 1701–1712, 2020. https://doi.org/10.1007/s00170-020-05530-4

[9] A. Sharma, D. Marla, S. Joshi *et al.,* "A study of femtosecond laser processed microtextures on silicon wafers to enhance optical absorption," *Laser Manuf. Mater. Process.*, vol. 9. no. 3, pp. 277–291, 2022. https://doi.org/10.1007/s40516-022-00176-4

[10] M. Calderon, M. Manso-Silván, A. Rodríguez *et al.,* "Surface micro- and nano-texturing of stainless steel by femtosecond laser for the control of cell migration," *Sci. Rep.,* vol. 6, no. 1, 36296, 2016. https://doi.org/10.1038/srep36296

[11] E. Cumbajin, N. Rodrigues, P. Costa *et al.,* "A systematic review on deep learning with CNNs applied to surface defect detection," *J. Imaging*, vol. 9, no. 10, 193, 2023. https://doi.org/10.3390/jimaging9100193

[12] P. Osterrieder, L. Budde, and T. Friedli, "The smart factory as a key construct of Industry 4.0: A systematic literature review," *Int. J. Prod. Econ.*, vol. 221, 107476, 2020. https://doi.org/10.1016/j.ijpe.2019.08.011

[13] T. Wang, Y. Chen, M. Qiao *et al.,* "A fast and robust convolutional neural network-based defect detection model in product quality control," *Int. J. Adv. Manuf. Technol.*, vol. 94, no. 9, pp. 3465–3471, 2018. https://doi.org/10.1007/s00170-017-0882-0

[14] P. Tchatchoua, G. Graton, M. Ouladsine *et al.,* "Application of 1D ResNet for multivariate fault detection on semiconductor manufacturing equipment," *Sensors*, vol. 23, no. 22, 9099, 2023. https://doi.org/10.3390/s23229099

[15] J. Khanam, M. Hussain, R. Hill *et al.,* "A comprehensive review of convolutional neural networks for defect detection in industrial applications," *IEEE Access*, vol. 12. pp. 94250–94295, 2024. https://doi.org/10.1109/ACCESS.2024.3425166

[16] H. Nguyen, G. Yu, N. Shin *et al.,* "Defective product classification system for smart factory based on deep learning," *Electronics,* vol. 10, no. 7, 826, 2021. https://doi.org/10.3390/electronics10070826

[17] G. A. Elhendawy and Y. El-Taybany, "Machine vision-assisted welding defect detection system with convolutional neural networks," *Int. J. Precis. Eng. Manuf.*, vol. 26, no. 12, pp. 3185–3194, 2025. https://doi.org/10.1007/s12541-025-01281-y

[18] A. Semitela, M. Pereira, A. Completo *et al.,* "Improving industrial quality control: A transfer learning approach to surface defect detection," *Sensors*, vol. 25, no. 2, 527, 2025. https://doi.org/10.3390/s25020527

[19] D. Patil, A. Nigam, S. Mohapatra *et al.,* "A deep learning approach to classify and detect defects in the components manufactured by laser directed energy deposition process," *Machines*, vol. 11, 854, 2023. https://doi.org/10.3390/machines11090854

[20] J. Kuo, J. Wu, P. Huang *et al.,* "Inspection of sandblasting defect in investment castings by deep convolutional neural network," *Int. J. Adv. Manuf. Technol.*, vol. 120, no. 3, pp. 2457–2468, 2022. https://doi.org/10.1007/s00170-022-08841-w

[21] R. Tyrone, "Benchmarking classical and deep learning approaches for defect detection in high-resolution wafer inspection imaging," *QIT Press Int. J. Comput. Sci.*, vol. 3, no. 1, pp. 7–12, 2023. https://qitpress.com/articles/QITP-IJCS_3_01_002

[22] C. Szegedy, W. Liu, Y. Jia *et al.,* "Going deeper with convolutions," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.,* 2015, pp. 1–9. https://doi.org/10.1109/CVPR.2015.7298594

[23] K. He, X. Zhang, S. Ren *et al.,* "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.,* 2016, pp. 770–778. https://doi.org/10.1109/CVPR.2016.90
[24] N. Deshmukh, A. Rajurkar, O. Kolekar *et al.,* "Thermal modeling of laser surface micro-texturing: Investigation on effects of laser parameters on dimple-texture dimensions and aspect ratio," *Mater. Today Proc.,* vol. 46, pp. 8374–8380, 2021. https://doi.org/10.1016/j.matpr.2021.03.420
[25] H. Wei, C. L. Huang, L. D. Hanlian *et al.,* "Modeling and optimizing femtosecond laser process parameters for high-efficient and near damage-free micromachining of single-crystal GaN substrate," *Mater. Sci. Semicond. Process*., vol. 153, 107123, 2023. https://doi.org/10.1016/j.mssp.2022.107123
[26] Z. Shen, F. C. Wang, R. X. Zhiguo *et al.,* "Numerical simulation of lubrication performance on chevron textured surface under hydrodynamic lubrication," *Tribol. Int.*, vol. 154, 106704. https://doi.org/10.1016/j.triboint.2020.106704
[27] M.-S. Suh, Y.-H. Chae, S.-S. Kim *et al.,* "Effect of geometrical parameters in micro-grooved crosshatch pattern under lubricated sliding friction," *Tribol. Int.,* vol 43, no. 8, pp. 1508–1517, 2010. https://doi.org/10.1016/j.triboint.2010.02.012
[28] H. Jigang, Q. Qin, W. Jie *et al.,* "Two-dimensional laser galvanometer scanning technology for additive manufacturing," *Int. J. Mech. Mechatron. Eng.,* vol. 6. no. 5, pp. 332–336, 2018.