lubricants

MDPI

Article

U-Net Segmentation with Bayesian-Optimized Weight Voting for Worn Surface Analysis of a PEEK-Based Tribological Composite

Yuxiao Zhao [1] and Leyu Lin [1,2,*]

[1] Chair of Composite Engineering (CCe), Rheinland-Pfälzische Technische Universität (RPTU) Kaiserslautern-Landau, 67663 Kaiserslautern, Germany; yuxiao.zhao@mv.rptu.de
[2] Research Center OPTIMAS, Rheinland-Pfälzische Technische Universität Kaiserslautern-Landau, 67663 Kaiserslautern, Germany
* Correspondence: leyu.lin@mv.rptu.de

Abstract

This study presents a U-Net-based automatic segmentation framework for quantitative analysis of surface morphology in a PEEK-based composite following tribological testing. Controlled Pin-on-Disc tests were conducted to characterize tribological performance, worn surfaces were captured by laser scanning microscopy to acquire optical images and height maps, and the model produced pixel-level segmentation masks distinguishing different regions, enabling high-throughput, objective analysis of worn surface morphology. Sixty-three manually annotated image sets—with labels for fiber, third-body patch, and matrix regions—formed the training corpus. A 70-layer U-Net architecture with four-channel input was developed and rigorously evaluated using five-fold cross-validation. To enhance performance on the challenging patch and fiber classes, the top five model instances were ensembled through Bayesian-optimized weighted voting, achieving significant improvements in class-specific F1 metrics. Segmentation outputs on unseen data confirmed the method's robustness and generalizability across complex surface topographies. This approach establishes a scalable, accurate tool for automated morphological analysis, with potential extensions to real-time monitoring and other composite systems.

Keywords: Artificial Neural Network (ANN); U-Net; tribology; PEEK-based composite

check for updates

Received: 7 June 2025
Revised: 19 July 2025
Accepted: 22 July 2025
Published: 24 July 2025

Citation: Zhao, Y.; Lin, L. U-Net Segmentation with Bayesian-Optimized Weight Voting for Worn Surface Analysis of a PEEK-Based Tribological Composite. *Lubricants* **2025**, *13*, 324. https://doi.org/10.3390/lubricants13080324

Copyright: © 2025 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (https://creativecommons.org/licenses/by/4.0/).

1. Introduction

Polymer materials are extensively valued in mechanical engineering for their favorable performance in high-friction, high-wear environments, finding applications in aerospace, automotive braking systems, and beyond [1–7]. To withstand substantial tribological stress, polymer matrices are often reinforced with fibers—such as carbon or glass—and supplemented with functional fillers, including solid lubricants (e.g., graphite) and submicron- or nanoparticles, which synergistically enhance wear resistance and reduce friction [8–11]. Early studies highlighted the benefits of combining fibers with solid lubricants for improved load-bearing and friction-reduction performance [12–15], while more recent work has demonstrated that the inclusion of nanoparticles further promotes tribological synergies within polymer-based tribocomposites [2,16–19]. These nanoparticles are believed to alleviate surface damage during sliding by providing a "rolling" or "filling" effect (Figure 1), as evidenced by micro-scratches observed on carbon fiber surfaces, suggesting the release and redeposition of particulate from the polymer matrix under sliding conditions [20]. Consequently, worn surfaces exhibit plastically deformed layers whose mechanical properties

differ markedly from those of the bulk matrix, leading to altered wear and friction behaviors during steady-state sliding against metallic counterparts [21]. In carbon fiber reinforced composites, wear performance is governed by fiber damage modes—including fracture, detachment, pulverization, and thinning—with fiber thinning correlating with enhanced wear resistance, whereas fiber fracture and pull-out lead to inferior performance [22,23]. The interfacial bonding strength between the fibers and matrix critically influences these damage modes: stronger bonding enables fibers to better bear applied loads, reducing matrix plowing and interfacial shear forces, and thus enhancing both friction reduction and anti-wear properties [24–26]. During repetitive sliding, wear debris and matrix material compact at the interface, forming plastically deformed layers in which portions of fibers may protrude or become embedded within polymer patches. Such patches, reported in various systems (e.g., polyamide 66 (PA 66) [17], polyphenylensulfid (PPS) [10], phenolic resin [27], and polyetheretherketone (PEEK) [15]), encapsulate fibers within wear particles relocated or even sintered onto the surface. Drawing on the terminologies of Godet [28], Denape [29], and Krejčí et al. [30], we define these structures as "third-body patches", characterized by accumulations of submicron- or nanoparticles on the leading side of fibers that create a "jumping pit" effect (Figure 1), dispersing stress concentrations at fiber ends and thereby mitigating fiber damage and lowering wear rates.

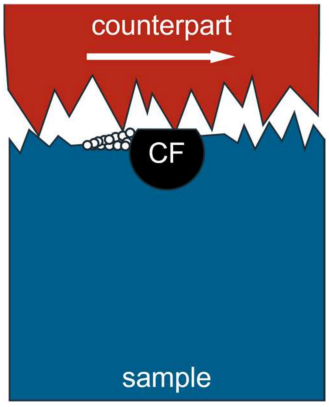


Figure 1. Schematic representation of jumping pit effects of submicron particles in composite material during sliding process. "CF" stands for carbon fiber and white balls accumulated by it are submicron- or nanoparticles. White arrow indicates the sliding direction.

Despite extensive qualitative characterization of third-body patches, most studies have relied on small sample sizes and descriptive observations, lacking robust, high-throughput quantification of patch morphology and distribution. In our prior work, we manually annotated 63 multi-channel microscopy images to quantify fiber, patch, and matrix regions and found that larger third-body patch areas correlated with lower wear rates. Although this manual approach demonstrated the value of morphometric analysis for elucidating tribological mechanisms, it was extremely labor-intensive and inherently subjective, significantly limiting its scalability for large-scale quantitative studies. With the rapid development of deep learning, semantic-segmentation techniques have been widely applied to automate microstructural analysis tasks in materials science. Examples include the successful application of U-Net architectures in the automated phase segmentation of electron backscatter diffraction (EBSD) maps of dual-phase steels [31] and in the semantic segmentation of progressive micro-cracking in polymer composites using Attention U-Net architectures [32]. Nevertheless, precise segmentation of complex tribological worn surfaces remains understudied. To address this research gap, we propose a U-Net-based automated segmentation framework specifically tailored to worn surfaces of PEEK-based composites, enabling high-throughput and objective quantitative morphometric analysis.

Deep learning-based image segmentation has witnessed transformative advances over the past decade, driven largely by the advent of fully convolutional architectures. Early studies applied classical machine learning pipelines such as Support Vector Machine (SVM) to morphological and textural descriptors of micrographs, achieving classification accuracies exceeding 90% [33,34]. Further studies showed that random forest regression on micromagnetic feature sets can reliably predict material properties from microstructural images [35]. More recently, Transformer-based frameworks have combined self-attention with physics-informed feature modules to deliver accurate phase segmentation and real-time defect detection in polymer composites [36,37]. Among these, U-Net has emerged as a paradigmatic model: originally designed for biomedical image segmentation, it demonstrated remarkable success in delineating complex anatomical structures from limited training data. Falk et al. developed an ImageJ plugin (Version 1.1.0) enabling U-Net-based cell counting, detection, and morphometry on small datasets with minimal user intervention [38]. Zhou et al. proposed U-Net++, a nested U-Net variant that further narrows the semantic gap between encoder and decoder features to achieve substantial Intersection-over-Union (IoU) gains over standard U-Net [39]. U-Net's distinctive encoder–decoder topology, with its symmetric skip connections, enables the preservation of fine-grained spatial details while capturing high-level contextual features [40,41]. Beyond medical applications, U-Net and its variants have been effectively adapted to materials science challenges. Warren et al. and Shi et al. both developed U-Net-based segmentation methods for automated grain and grain-boundary analysis in metallic microstructures [42,43]. Martinez et al. and Wang et al. both implemented U-Net-based frameworks to automatically segment metallurgical phases, achieving pixel-wise accuracies above 95% and enabling rapid objective phase quantification [31,44]. Moreover, recent studies have widely applied U-Net and its variants to segmentation of composite microstructure images, achieving accurate and efficient extraction of key features (cracks, fibers, particles, and pores) from limited annotated data, and utilizing the segmentation outputs for quantitative analysis and 3D/2D reconstruction [32,45–47]. The robustness of U-Net to noisy, irregular morphologies and its ability to learn from relatively few annotated examples makes it particularly well suited for the segmentation of worn surface images, where region boundaries are often subtle and visual contrast is low. In addition, advanced segmentation models such as Attention U-Net [41] and DeepLabV3+ [48] have demonstrated improved performance on large-scale benchmarks; however, they also require greater computational power and parameter complexity, which can exacerbate overfitting on small datasets and impose heavy resource demands. The original U-Net architecture therefore offers a well-validated, resource-efficient baseline that converges reliably on limited training data.

In this present study, we apply a U-Net-based convolutional neural network to automate the segmentation of worn-surface micrographs from Pin-on-Disc tests, using combined optical and height data. The model is trained on a manually annotated dataset using cross-validation and enhanced through a multi-network voting scheme with Bayesian-optimized weights. This workflow overcomes the limitations of time-consuming manual labeling and delivers high-precision, pixel-wise segmentation of fibers, third-body patches, and matrix regions. By enabling rapid, large-scale morphometric analysis, our approach provides an objective, scalable tool for investigating tribological surface phenomena. Furthermore, the streamlined pipeline—from image acquisition to segmentation—reduces turnaround time and resource requirements, making routine application feasible in both research and industrial settings. Ultimately, this framework lays the groundwork for more efficient experimental workflows, supports the integration of automated image analysis into tribological studies, and accelerates the critical mechanistic feedback loop in new material development by automating quantitative measurements of fiber and patch morphologies,

which is an essential but time-consuming step. This paves the way for the development of intelligent, data-driven materials design strategies in tribology.

2. Methodology

2.1. Tribological Sample Preparation

A high-performance tribological composite was fabricated using a PEEK matrix (VESTAKEEP 2000 G, Evonik Operations GmbH, Marl, Germany) reinforced with short carbon fibers (Sigrafil C30 S600 APS, SGL Carbon GmbH, Wiesbaden, Germany) and graphite flakes (RGC 39A, Superior Graphite, Sundsvall, Sweden). Submicron titanium dioxide (Kronos 2310, Kronos Titan GmbH, Leverkusen, Germany) and zinc sulfide (Sachtolith HD-S, Venator GmbH, Duisburg, Germany) particles were incorporated as functional fillers. Detailed specifications of each filler are available in our previous work [26]. All components were first compounded using a Leistritz ZSE 18 MAXX twin-screw extruder (Leistritz AG, Nürnberg, Germany) in a two-step masterbatch process to ensure uniform dispersion of nanofillers, with barrel temperatures set between 370 °C and 395 °C [49]. The resulting granules were then injection-molded into 50 × 50 × 4 mm^3 sheets on an ENGEL victory 200/80 SPEX machine (ENGEEL Austria GmbH, Schwertberg, Austria) at a melt temperature of 395 °C. Finally, specimens of 4 × 4 × 4 mm^3 were precisely milled from these sheets and eventually slid against a 100Cr6 bearing steel disc (arithmetical mean roughness R_a = 0.2 ± 0.02 µm, Rockwell hardness ≈ 62 HRC) on a Pin-on-Disc (PoD) setup (Figure 2) in accordance with ASTM G99-17 standards [50]. Tests were conducted at 23 °C and 50% relative humidity under eight distinct pv (pressure × velocity) combinations ranging from 1 to 32 MPa·m/s (cf. Table 1), with a normal load applied via a pneumatic system and sliding speeds adjusted by the disc rotation rate. Two IR-Sensors measured the temperature of sample and counterbody during the test, respectively. For each load condition, three independent specimens were tested until steady-state friction coefficients and wear rates were achieved.

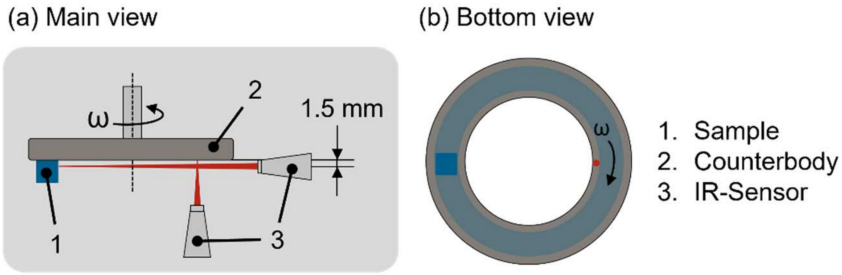


Figure 2. Schematic demonstration of the Pin-on-Disc tribometer in accordance with ASTM G99-17 standards.

Table 1. Selected load combinations of the tribological test.

Test Condition		**Selected pv-Combinations**							
Pressure	MPa	1	1	4	4	4	6	8	8
Velocity	m/s	1	4	1	2	4	4	1	4
pv-product	MPa·m/s	1	4	4	8	16	24	8	32

2.2. Image Acquisition and Annotation

Following Pin-on-Disc testing, worn specimens were examined using a Keyence VK-X1050 laser scanning microscope (Keyence Corporation, Osaka, Japan), with nine strategically distributed fields of view per sample to capture representative surface features (Figure 3). At each location, the instrument generated different complementary image modalities: a high-resolution

optical micrograph, a color-coded height map illustrating topographical variation, a C-Laser differential interference contrast (DIC) image for enhanced edge definition, and a quantitative height dataset in which each pixel encodes an absolute height value. This multimodal acquisition strategy ensures both morphological and height-based information are available, facilitating accurate delineation of fibers, third-body patches, and matrix regions.

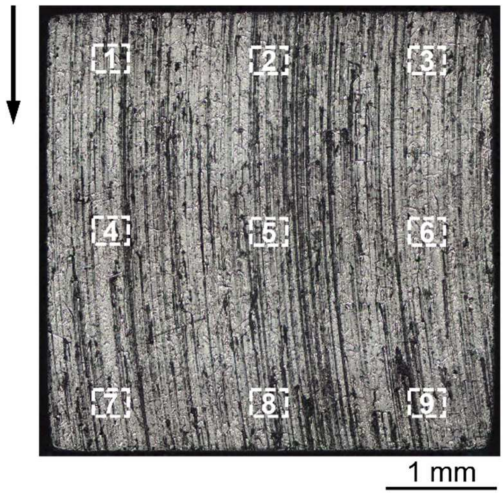


Figure 3. Optical image of worn surface from tested sample with nine LSM sampling positions (white squares marked) and sliding direction (black arrow).

According to the definition of third-body patches (Section 1), carbon fibers protrude slightly above the matrix surface following friction testing, while patches manifest as localized "ramping" regions formed by submicron debris accumulating at the entering side of fibers. These topographical variations are readily apparent in color-coded height heatmaps but are often indistinct in optical micrographs, making heatmaps the primary reference for region delineation. However, reliance on heatmaps alone can lead to annotation errors: for example, deep scratches on fiber surfaces produce color shifts in the heatmap that closely resemble debris-filled patches or matrix inclusions (Figure 4a). In the corresponding optical image (Figure 4b), these scratches remain visually consistent with intact fiber regions and are thus correctly identified. Moreover, because height heatmaps encode elevation via an RGB colormap, they do not capture micro-scale height differences with full fidelity. To address these limitations, we replaced heatmaps with single-channel quantitative height maps wherein each pixel stores an exact height value. This approach not only improves height resolution for precise annotation but also reduces storage requirements compared to three-channel heatmap images. Building on the co-registered optical micrographs and single-channel height images, we manually annotated 63 multi-modal images sets with MATLAB (R2024a), assigning each pixel to one of three classes—Fiber, Patch, and Matrix. The resulting label maps (visualized in Figure 4c) use red, green, and blue to denote Fiber, Patch, and Matrix regions, respectively. To ensure consistent and reproducible annotation, all images were labeled by a single experienced researcher, strictly following the definition of third-body patches provided in Section 1. Specifically, patches were identified by meeting two simultaneous conditions: (1) a clear and localized elevation change ("ramping") at the fiber's entering side, evident in quantitative height maps, and (2) the presence of submicron-scale debris accumulations confirmed by corresponding optical images. In cases where ambiguities arose, annotations were cross-verified and confirmed by a second experienced researcher to maintain consistency. Custom MATLAB script was then applied to process these label maps and calculate key morphometric metrics, including area fractions and fiber segment lengths, which form the quantitative basis for neural network training and performance evaluation.

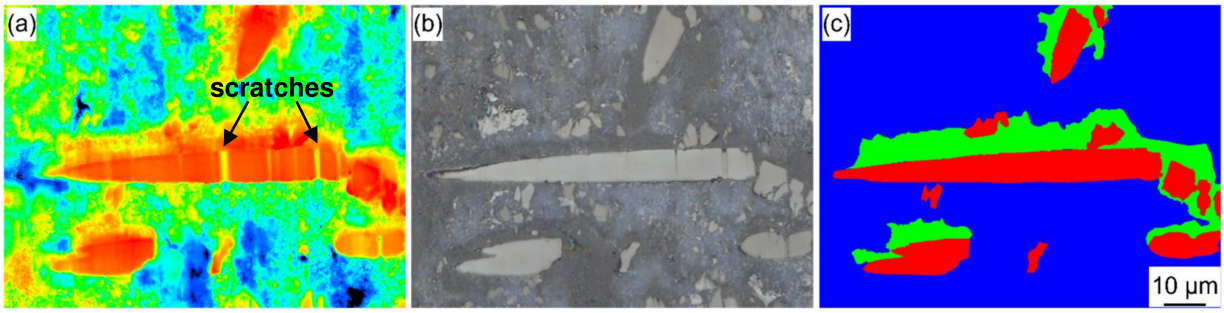


Figure 4. LSM-acquired (**a**) color-coded height map (red stands for higher position, blue for lower position), (**b**) optical micrograph, and (**c**) corresponding visualized manual segmentation mask (fiber: red; patch: green; matrix: blue) of the worn surface.

2.3. U-Net Architecture Design

The U-Net network was initially proposed by Olaf Ronneberger, Philipp Fischer, and Thomas Brox in 2015 to address challenges in biomedical image segmentation, specifically in scenarios involving limited training data [40]. Traditional convolutional neural networks (CNNs) often encounter challenges in accurately segmenting images that require precise localization of features. This difficulty arises primarily due to the pooling operations inherent in CNN architectures, which significantly reduce the spatial resolution of feature maps. As highlighted by Gholamalinezhad et al., pooling layers, while beneficial for reducing computational complexity and controlling overfitting, can lead to the loss of fine-grained spatial information essential for tasks like image segmentation [51]. U-Net resolves this limitation by introducing a symmetric encoder–decoder structure combined with skip connections, thus preserving spatial information while simultaneously capturing hierarchical features.

At its core, U-Net employs a contracting path (encoder) and an expansive path (decoder), which together form a distinctive "U"-shaped architecture (Figure 5). The contracting path consists of sequential convolutional blocks, each comprising two consecutive convolutions followed by Rectified Linear Unit (ReLU) activations (Figure 6a), and a subsequent max-pooling operation. This path progressively reduces the spatial dimensions of the feature maps while increasing their depth, enabling the network to extract meaningful high-level semantic features. Mathematically, the convolution operation performed in each block can be described by the following:

$$Y = f\left(\sum_i X_i \cdot W_i + b\right) \tag{1}$$

where X_i denotes input feature maps, W_i are convolutional kernels, b is the bias term, and $f(\cdot)$ represents the ReLU activation function, defined as follows:

$$ReLU(x) = max(0,\ x) \tag{2}$$

The expansive path reverses this process by employing upsampling operations, typically through transposed convolutions, to recover the spatial resolution of the feature maps. Mathematically, the transposed convolution operation in each upsampling step can be described by (similar to convolution operation) the following:

$$Y = f\left(\sum_i X_i \circledast W\prime_i + b\right) \tag{3}$$

where X_i denotes input feature maps from the lower-resolution layer, $W\prime_i$ are transposed convolutional kernels, b is the bias term, $\circledast$ symbolizes the transposed convolution operator, and $f(\cdot)$ represents the ReLU activation function. By doubling the height and

width of the feature maps at each stage while halving their channel depth, this operation effectively reverses the encoder's downsampling, enabling recovery of the original image dimensions and the fusion of high-resolution spatial details with the learned semantic features. Each upsampling step is followed by concatenation with corresponding feature maps from the contracting path via skip connections. These skip connections directly transfer high-resolution spatial information, mitigating the loss of detailed localization caused by pooling operations. The concatenated feature maps are then subjected to two additional convolutions and ReLU activations, refining the features and facilitating accurate pixel-level segmentation.

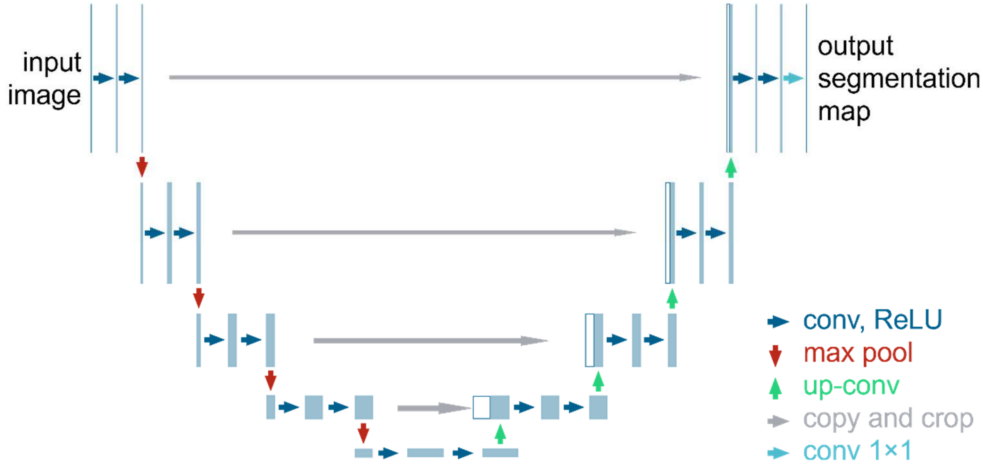


Figure 5. Schematic representation of U-Net architecture. Light blue boxes correspond to feature maps. White boxes denote copied feature maps. Arrows represent different operations [40].

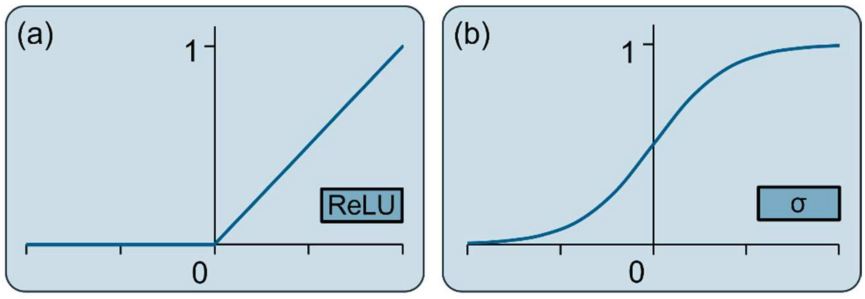


Figure 6. Schematic illustration of (**a**) ReLU function and (**b**) sigmoid function.

The final segmentation map is generated through a 1×1 convolution layer at the end of the expansive path, which maps the high-dimensional feature maps into the desired number of classes. Formally, this operation is described as follows:

$$Z = \sigma(W_{1\times1} * X + b) \tag{4}$$

where Z represents the segmentation probability map, $W_{1\times1}$ represents the convolutional kernel, and $\sigma(\cdot)$ denotes the sigmoid activation function (Figure 6b) for binary segmentation or the softmax function for multi-class segmentation, defined as follows:

$$\sigma(x) = \frac{1}{1 + e^{-x}} \tag{5}$$

The integration of skip connections is a pivotal innovation in U-Net, allowing precise spatial information to persist across the network, thereby significantly enhancing segmentation accuracy, especially along object boundaries. This architectural design enables U-Net to effectively perform complex image segmentation tasks, even with a limited amount

of training data. In the context of analyzing the morphological characteristics of fibers and patches on the worn surface of composite materials, the U-Net architecture offers significant advantages. Its ability to accurately segment complex structures with limited training data aligns well with the challenges presented by such images, where the precise delineation of fiber boundaries is essential for subsequent analysis. The incorporation of skip connections ensures that fine-grained spatial information is preserved, facilitating detailed morphological assessments critical in materials science research.

Building upon the basic U-Net architecture, targeted optimizations tailored specifically for segmenting worn surface morphologies were introduced. The resulting deep neural network accepts four-channel inputs (768 × 1024 pixels), which combine optical micrographs (three channels) and corresponding height-map data (single channel) previously described in Section 2.2. The network produces pixel-wise segmentation maps for three classes: Fiber, Patch, and Matrix. Structurally, the proposed architecture follows the characteristic U-shaped encoder–decoder framework and has been carefully balanced in complexity to suit our workstation's computational capabilities, resulting in approximately 124 million trainable parameters.

The encoder path (contracting path) consists of five successive stages. Each stage (Figure 7a) involves two **convolutional layers**, each applying 3 × 3 convolutional kernels. These convolutional layers capture local spatial patterns by applying learnable filters across feature maps. Each convolutional operation is immediately followed by a **ReLU layer**. After two consecutive convolutional and ReLU activation pairs, each encoder stage includes a 2 × 2 **max pooling layer**, which reduces spatial dimensions by half. Max pooling extracts dominant features within local regions, allowing the network to learn progressively abstract, high-level representations while reducing the computational load. As the spatial dimensions decrease, the number of feature channels doubles at each stage (starting from 64 channels), enhancing the representational capability of the network. The deepest encoder level (Stage 5) incorporates an additional **dropout layer**, randomly zeroing elements in the feature map to regularize the model and prevent overfitting at the network's most abstract representation level, without compromising the shallow spatial details. Connecting the encoder and decoder is a bridge block (Figure 7b), composed of two additional 3 × 3 convolutional layers, each followed by ReLU activations, as well as another dropout layer. This bridge further refines and stabilizes abstract features at the smallest resolution (24 × 32 pixels), ensuring robust semantic information transfer into the decoding stage. The decoder path (expanding path) mirrors the encoder, restoring spatial resolution in five successive stages. Each decoding stage (Figure 7c) begins with a 2 × 2 **transposed convolution layer** (UpConv), which learns an optimal up-sampling operation to double the spatial dimensions. Each transposed convolutional operation is followed by ReLU activation to introduce further nonlinearity. **Depth concatenation layer** concatenates the resulting feature maps along the depth dimension with corresponding encoder features from the contracting path through **skip connection**, reintroducing detailed spatial information (such as fiber boundaries) lost during pooling operations. Following concatenation, two additional 3 × 3 convolutional layers, each with ReLU activations, further refine these combined feature maps, improving the segmentation precision.

Finally, at full resolution (768 × 1024 pixels), the network applies a 1 × 1 convolutional layer, which compresses the depth dimension from 64 to 3 channels, corresponding directly to the segmentation classes. This is followed by a **softmax layer**, which normalizes the output into pixel-wise class probabilities, enabling a clear probabilistic interpretation of the segmentation results. Ultimately, the class probabilities are mapped into discrete class labels by selecting the class with the highest probability at each location. Complete schematic illustration of the custom U-Net architecture is shown in Figure 8.

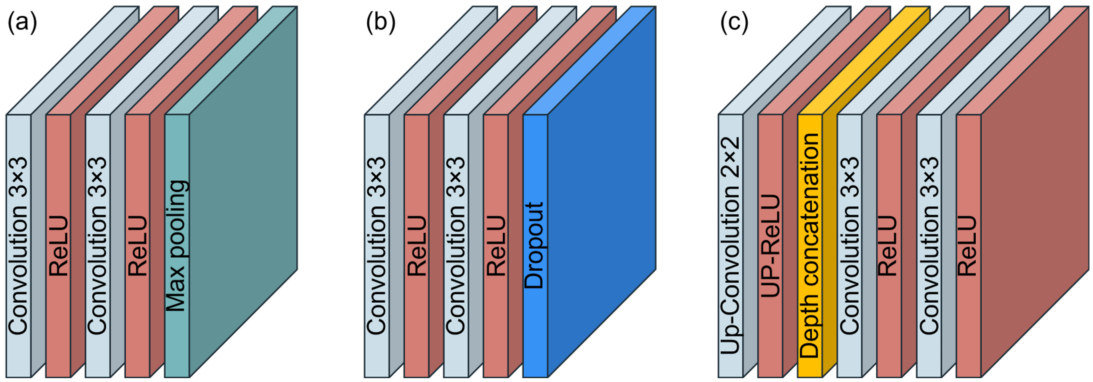


Figure 7. Schematic representation showing the structures of (**a**) encoder stage, (**b**) bridge block, and (**c**) decoder stage.

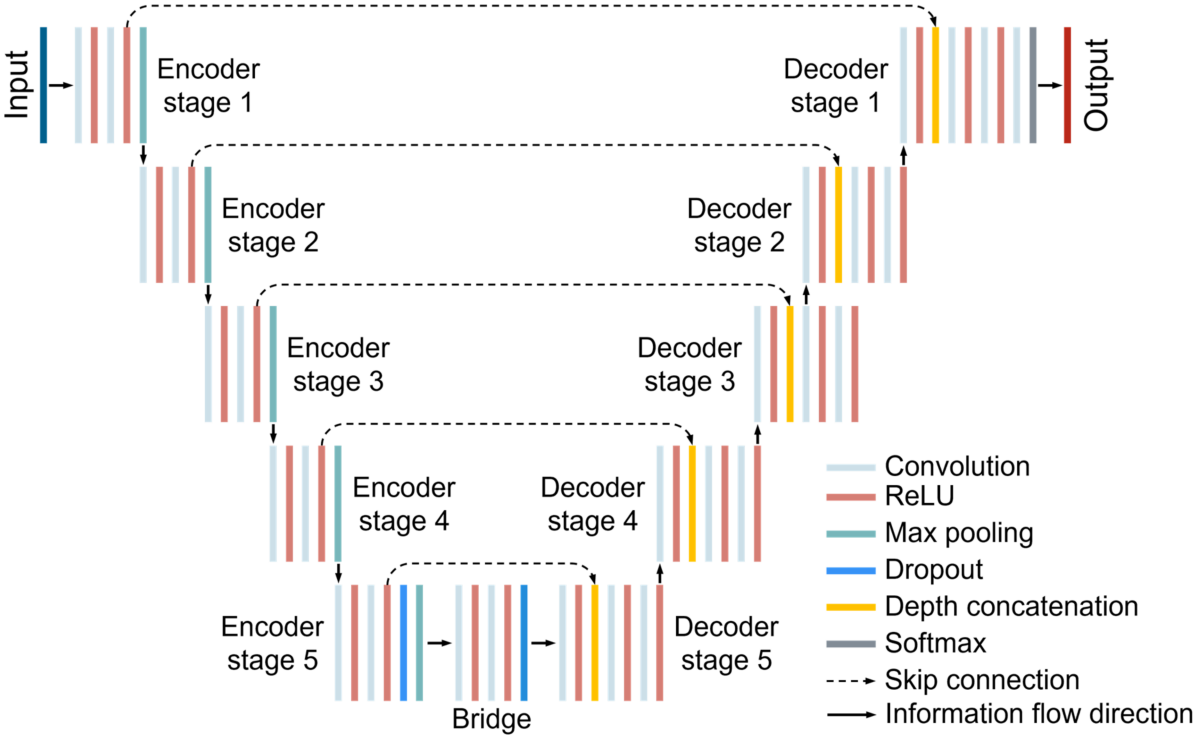


Figure 8. Schematic illustration of specifically optimized U-Net architecture.

Several task-specific design considerations guided the selection of architectural parameters. The choice of five encoder and decoder levels was motivated by the need to adequately capture hierarchical morphological information in the images, ranging from fine-scale textures of fibers to larger-scale spatial distributions of material patches. An initial convolutional filter count of 64 was chosen after careful consideration to balance representational power and computational demand, making the architecture suitable for the processing capabilities of our workstation. Retaining 1024 channels at the bridge level ensures sufficient capacity for learning complex, texture-rich interactions characteristic of tribological surfaces. These carefully designed optimizations and parameter selections ensure that the custom U-Net network effectively captures both global morphological characteristics and detailed structural boundaries of fibers and patches, providing robust segmentation performance tailored specifically to the challenges of composite tribological analysis.

2.4. Training Data Preparation and Network Training

Following the definition and optimization of the U-Net structure described previously, a comprehensive training strategy was established to effectively segment worn surface morphology. The training procedure includes carefully prepared input datasets, a robust training configuration with well-chosen hyperparameters, data augmentation strategies to enhance generalization, and rigorous performance validation via K-fold cross-validation.

The training data set was prepared by combining optical images with corresponding height maps (in micrometers), as introduced in Section 2.2. Each input sample, therefore, comprises four channels—three from the optical data (RGB value ranging from 0 to 255), capturing detailed surface textures, and one additional channel encoding height information, providing crucial depth-related structural insights. All images and corresponding segmentation labels were uniformly resized to 768 $\times$ 1024 pixels to preserve detailed spatial features while ensuring computational feasibility. Labels were verified and cleaned to contain only valid categorical classes (Fiber, Patch, and Matrix) and converted into categorical data types suitable for semantic segmentation tasks. To enhance the model's robustness against variations in real-world data, extensive data augmentation was performed during training. Random horizontal and vertical reflections were applied to simulate variations in fiber orientations. Random translations of up to $\pm$10 pixels in both horizontal and vertical directions, rotations within $\pm$15 degrees, and scaling ranging from 0.8 to 1.2 times the original size were introduced. These augmentation techniques expand the effective size and variability of the training dataset, enabling the network to better generalize to unseen data.

Training was conducted using the Adam optimizer, known for its adaptive learning rate properties and suitability for deep learning segmentation tasks. An initial learning rate of 2.6×10^{-4} was selected based on empirical experimentation, balancing convergence speed with training stability. The model underwent training for a maximum of 300 epochs, and the mini-batch size was set at 4, a choice dictated by the GPU memory constraints of our workstation, ensuring a stable training procedure without memory overflow. To rigorously evaluate and validate the segmentation performance of the network, a K-fold cross-validation strategy with K = 5 folds was employed. This technique randomly divides the dataset into five subsets of approximately equal size. For each iteration, the network was trained on four subsets while the remaining subset served as a validation set. This approach ensures all data contributes to both training and evaluation, providing a thorough and unbiased assessment of model performance and preventing potential bias caused by a single training–validation split. For each fold of the cross validation, the training and validation datasets were explicitly defined, augmented appropriately for training, and left unaltered for validation to accurately represent the model's performance on unseen data. Model training was monitored closely using MATLAB's built-in training progress visualization, providing real-time feedback on metrics such as training loss, validation accuracy, and convergence trends. Additionally, a checkpoint-saving mechanism was implemented to periodically preserve the intermediate states of the network at predefined iterations, safeguarding training progress against interruptions. All training experiments were performed on a workstation equipped with an NVIDIA RTX 4090 GPU (24 GB VRAM, Nvidia, Santa Clara, CA, USA) and a 24-core Intel i9-14900 CPU (128 GB RAM, Intel, Santa Clara, CA, USA). Five-fold cross-validation required approximately one hour per fold ($\approx$5 h total) with minibatches of four, leveraging GPU acceleration and a parallel CPU pool for data augmentation. For multi-modal input images of size 768 $\times$ 1024 pixels (four channels), single-image inference takes on average 1.3 s on the GPU. Leveraging MATLAB's parallel CPU pool, batch inference of ten images completes in approximately 3 s, demonstrating

that the framework can segment hundreds of images per hour without specialized industrial hardware.

Upon completion of each fold's training, a rigorous validation approach was implemented to systematically evaluate the segmentation performance of the trained U-Net models. To quantitatively assess the accuracy and robustness of the model predictions, several metrics were computed using the validation dataset. Specifically, overall accuracy (OA), mean Intersection-over-Union (mIoU), precision, recall, and F1-score metrics were calculated based on the confusion matrix obtained from the validation results of each fold. The confusion matrix, illustrated exemplarily in Figure 9, provides a comprehensive summary of the model's prediction performance across the defined segmentation classes. This matrix presents the distribution of true positives (TP), false positives (FP), false negatives (FN), and true negatives (TN), enabling detailed analysis of classification performance for each class.

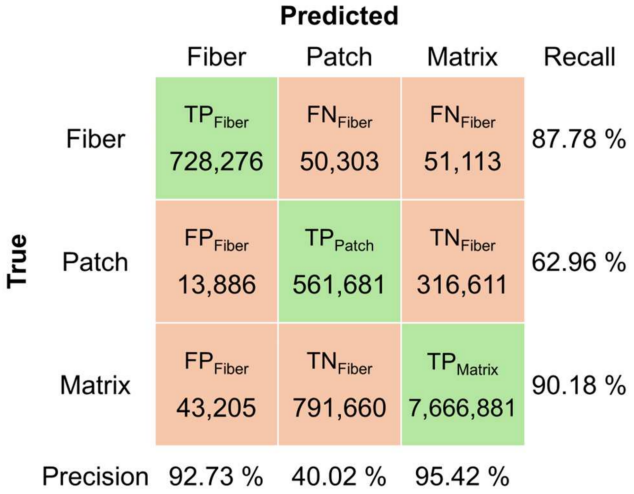


Figure 9. Example of a confusion matrix for a trained model. Green diagonal squares denote each class's true positives, red off-diagonal squares indicate misclassifications and are annotated from the fiber perspective (i.e., showing fiber's false positive, false negatives, and true negatives).

Based on the confusion matrix, precision, recall, and F1-score for each class were calculated as follow:

- Precision represents the proportion of correctly identified instances of a class among all instances predicted as that class. It quantifies the reliability of positive predictions and is computed as follows (i denotes the class index; same applies below):

$$Precision_i = \frac{TP_i}{TP_i + FP_i} \tag{6}$$

- Recall measures the proportion of correctly identified instances of a class out of all actual instances belonging to that class. It indicates the model's sensitivity to correctly detecting positive instances and is defined as follows:

$$Recall_i = \frac{TP_i}{TP_i + FN_i} \tag{7}$$

- F1-score is the harmonic mean of precision and recall, providing a balanced measure between the two metrics. It is especially useful when dealing with imbalanced classes, as it equally emphasizes both false positives and false negatives:

$$F1_i = 2 \times \frac{Precision_i \times Recall_i}{Precision_i + Recall_i} \tag{8}$$

Additionally, two broader measures were computed to offer an overall assessment of model performance across all classes:

- Overall accuracy (OA), defined as the ratio of the correctly classified pixels to the total number of pixels, provides an intuitive measure of general classification effectiveness:

$$OA = \frac{\sum_{i=1}^{n} TP_i}{\sum_{i=1}^{n} (TP_i + FP_i)} \tag{9}$$

- Mean Intersection-over-Union (mIoU) is calculated as the average Intersection-over-Union (IoU) across all segmentation classes. IoU for each class quantifies the overlap between the predicted and ground-truth regions, making it particularly suitable for evaluating segmentation tasks:

$$IoU_i = \frac{TP_i}{TP_i + FP_i + FN_i} \tag{10}$$

$$mIoU = \frac{1}{n} \sum_{i=1}^{n} IoU_i \tag{11}$$

These five metrics (precision, recall, F1-score, OA, and mIoU) together provided a comprehensive and nuanced view of model performance, allowing for in-depth comparisons across the folds. They were essential for subsequent model evaluation and for selecting the best-performing model for practical application. The detailed confusion matrices and associated metrics ensured that network performance was objectively and transparently assessed, providing clear insights into areas of model strength and highlighting potential areas for improvement.

2.5. Multi Nets Weighted Voting

Following the training phase, the performances of each individual model were systematically evaluated using precision, recall, and F1-score metrics, placing particular emphasis on recall and F1-score due to notable class imbalance in the training dataset. Specifically, approximately 80% of the labeled pixels corresponded to the "Matrix" class, while "Fiber" and "Patch" constituted roughly 10% each, resulting in a severe class imbalance. Such imbalance commonly biases neural networks towards dominant classes. Given that the primary objective of the segmentation task was to accurately identify and delineate "Fiber" and "Patch" regions, emphasis was placed on recall and F1-score metrics for these minority classes. Based on these evaluations, five U-Net models exhibiting high recall and balanced F1-score across all three classes were selected for subsequent analysis. Each of these models, trained independently using a randomly partitioned data subset (due to the nature of K-fold cross-validation), inherently captures slightly different aspects of the data distribution and thus provides complementary insights when combined.

To enhance segmentation accuracy, a weighted ensemble approach was adopted, combining predictions from multiple independently trained U-Net models. Ensemble methods, particularly those employing weighted voting schemes, have been widely recognized for improving predictive performance in various image segmentation tasks [52–54]. In this context, each model's contribution to the final prediction was weighted based on its proficiency in identifying specific classes. For instance, certain models exhibited heightened sensitivity in detecting "Fiber" regions, while others were more adept at identifying "Patch" areas. By assigning higher weights to models where they performed best, the ensemble capitalized on the strengths of individual models, leading to more accurate and robust segmentation outcomes.

The weighted voting strategy was implemented as follows: for each pixel, the score that it belongs to a specific class was computed as a weighted sum of predictions from each of the five networks. Mathematically, for each pixel location $(i,\ j)$ and class c, the aggregated score $S_c(i,\ j)$ is given by the following:

$$S_c(i,\ j) = \sum_{k=1}^{5} w_{k,\ c} \cdot F1_{k,\ c} \cdot P_{k,\ c}(i,\ j) \tag{12}$$

where $P_{k,\ c}(i,\ j)$ represents the probability predicted by network k for class c at pixel $(i,\ j)$, and $F1_{k,\ c}$ denotes the baseline F1-score for network k evaluated previously on class c. The parameters $w_{k,\ c}$ are class-specific weights optimized using Bayesian optimization to further enhance the accuracy of predictions. The incorporation of each network's baseline F1-score as part of the weighting formula serves as a mechanism to inherently adjust the influence of networks according to their demonstrated reliability for each class. Thus, networks that historically exhibited superior accuracy for a particular class were effectively granted greater influence in the ensemble prediction for that class, a critical step towards addressing the aforementioned class imbalance and maximizing predictive accuracy.

To determine the optimal weighting parameters for the ensemble model, Bayesian optimization was employed due to its efficiency in optimizing complex, multidimensional, and expensive-to-evaluate function. Bayesian optimization is particularly well-suited for scenarios where the objective function lacks a closed-form expression and is costly to evaluate, such as tuning hyperparameters in machine learning models [55]. It operates by constructing a probabilistic surrogate model, typically a Gaussian Process (GP), to approximate the objective function. In this study, we employed MATLAB's "bayesopt" function with a Gaussian Process surrogate model utilizing an Automatic Relevance Determination (ARD) Matern 5/2 kernel. The acquisition function was set to "expected-improvement-plus" (EI+), which automatically balances exploration–exploitation through adaptive perturbation. We used a single initial point with all weights set to 1.0 to initialize the optimization and performed a total of 60 objective evaluations. Kernel hyperparameters were re-estimated at each iteration via marginal-likelihood maximization. This configuration ensures efficient global search, balancing exploration of the parameter space with the exploitation of known promising regions [56,57]. In the context of this study, the objective was to optimize the weights assigned to each of the five U-Net models in the ensemble, as well as the class-specific weights for "Fiber", "Patch", and "Matrix". The optimization aimed to maximize the ensemble's F1-score, particularly focusing on the "Patch" class due to its critical role in the segmentation task. Bayesian optimization was chosen over traditional methods like grid search or random search because it can achieve better results with fewer evaluations, which is crucial given the computational expense of evaluating each set of weights. This approach has been successfully applied in similar contexts, such as optimizing ensemble weights in image classification tasks. The optimization process began with an initial equal weighting configuration, gradually evolving toward the optimal solution through successive iterations guided by the Bayesian strategy. At each iteration, weighted predictions were evaluated against ground-truth labels to calculate precision, recall, and F1-scores, ensuring comprehensive tracking of ensemble performance improvements. Upon convergence, the final set of optimized weights was determined based on achieving the highest validation performance, measured principally by the F1-score, thereby ensuring maximal segmentation accuracy.

Through the integration of multiple independently trained networks, class-specific weighting schemes informed by each network's baseline performance, and rigorous Bayesian optimization of ensemble weights, the final ensemble model achieved superior segmentation outcomes. This carefully structured approach significantly mitigates the challenges posed by class imbalance, leverages the complementary strengths of individ-

ual networks, and effectively enhances the overall accuracy and reliability of composite material surface segmentation.

3. Results and Discussion

3.1. U-Net Training

Following the methodology detailed in Section 2.3, we selected a set of optimal hyperparameters based on the available dataset size, image resolution, and hardware capabilities of the workstation (cf. Table 2). For example, an initial learning rate of 2.6×10^{-4} was chosen to maintain sufficiently large gradient updates at early training stages while preventing instability from excessively high learning rates. Considering the high resolution of training images and available GPU memory, a minibatch size of four was selected to balance computational efficiency and prevent GPU memory overflow. Additionally, validation was performed every 30 iterations (approximately every two epochs) to monitor potential overfitting without excessively slowing the training process. Furthermore, a five-level encoder depth was adopted to achieve an optimal trade-off between preserving fine fiber textures (shallow-layer features) and capturing the global semantic context of patches (deep-layer features). The initial convolution layer with 64 kernels represented a balanced compromise between model capacity and computational efficiency, sufficient to extract complex fiber features while limiting the overall number of parameters. To maximize training efficiency, parallel CPU processing and GPU acceleration were employed, achieving roughly 3600 iterations per fold within the five-fold cross-validation, with each fold requiring approximately one hour.

Table 2. Selected hyperparameters of five-fold cross validation U-Net training process.

Hyperparameter	Value
Initial learn rate	2.6×10^{-4}
Mini batch size	4
Maximum epoch	300
Validation frequency	30
Encoder depth	5
No. of first encoder filters	64

To mitigate randomness in the training process, a total of 80 networks (16 sets of five-fold) were independently trained. Given that the validation data subsets within K-fold cross-validation were randomly partitioned, these networks can be considered independent from each other. Furthermore, data augmentation strategies were utilized to further enhance generalization and minimize overfitting. As previously described, individual performance metrics were calculated for each network using its respective validation subset. Given the significant dominance (~80%) of the matrix class within tribological surface images, overall accuracy (OA) alone was deemed unsuitable as a performance metric due to the potential bias toward the majority class. Therefore, recall and F1-score were selected as the primary indicators of segmentation performance. As illustrated in Figure 10, the matrix class exhibited exceptionally high average recall and F1-score values of 0.967 and 0.946, respectively, attributed to its large proportional coverage. Fibers, with clearly defined boundaries and height characteristics, also showed strong performance, with average recall and F1-score values of 0.912 and 0.908, respectively. However, the patch regions, characterized by ambiguous boundaries in both optical and height images, posed significant segmentation challenges, even for experienced researchers. Consequently, segmentation accuracy for patches was relatively low and unstable, with average recall and

F1-score values of 0.404 and 0.479, respectively. This level of accuracy was insufficient for our intended precise quantification of patch areas, underscoring the necessity of further methodological enhancements.

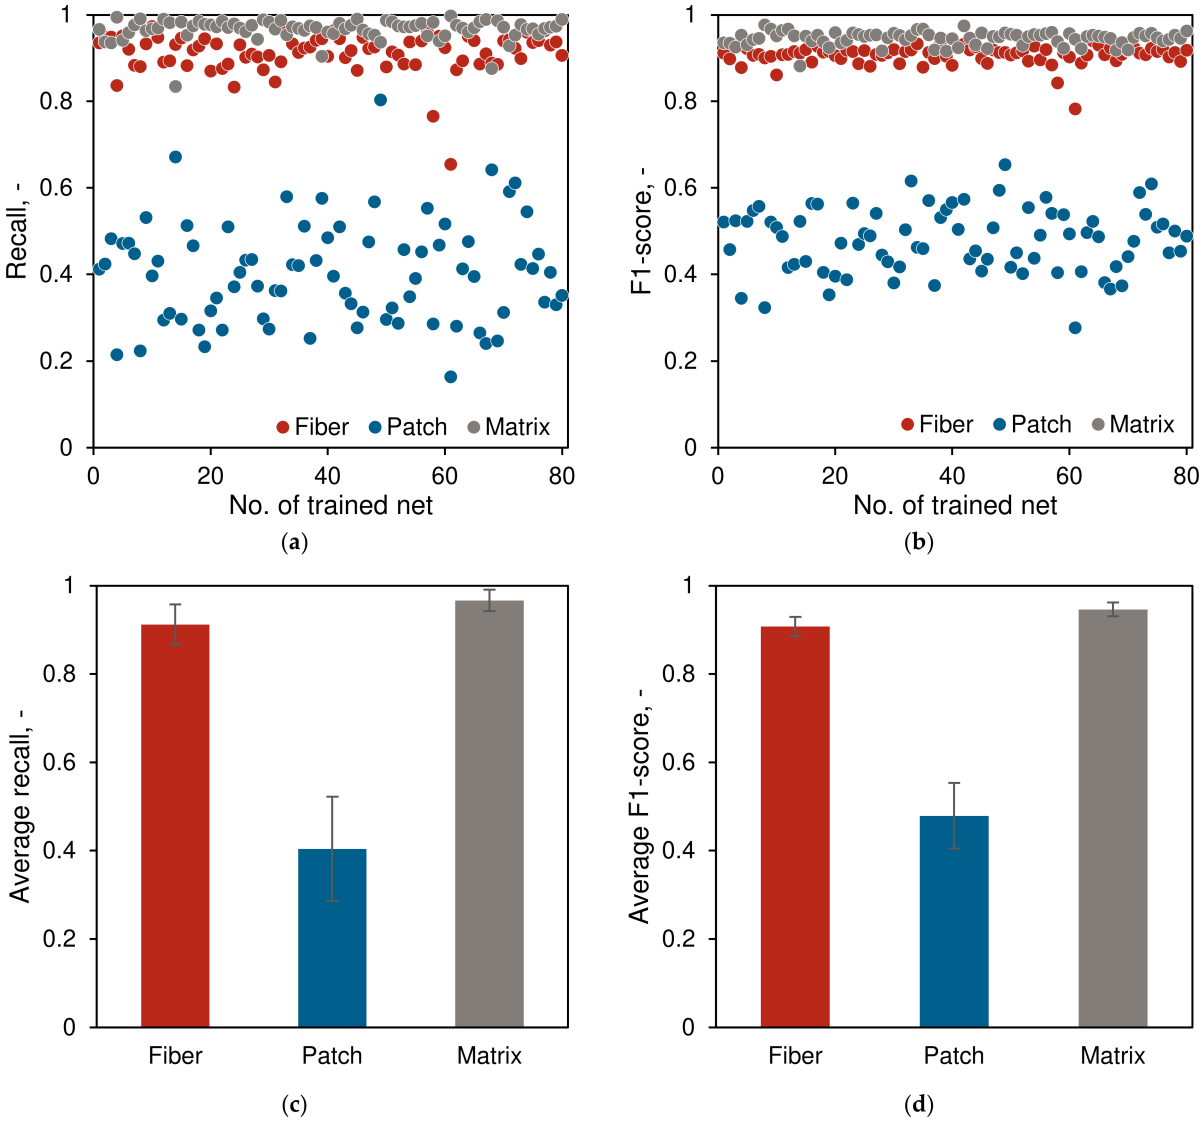


Figure 10. Evaluation metrics for segmentation results: (**a**) recall and (**b**) F1-score values for three labels across 80 trained networks; and (**c**) mean recall and (**d**) mean F1-score values averaged over the three labels.

3.2. Bayesian-Optimized Weighted Voting

As demonstrated in the previous section, the trained individual networks exhibited sufficiently high segmentation accuracy for fibers and matrix regions; however, the prediction accuracy for patch regions remained unsatisfactory. To address this issue and enhance overall segmentation accuracy, a multi-network weighted voting mechanism was introduced, primarily aimed at improving the segmentation accuracy for patch regions while maintaining the performance for fiber and matrix areas. Accordingly, the five networks demonstrating the highest individual patch-region performance (based on their respective F1-scores) were selected for inclusion in the voting ensemble. Following the procedure described in Section 2.3, each of these five networks individually segmented the images from the training dataset. Their pixel-wise predictions were then multiplied by their corresponding class-specific F1-scores and optimized weights, after which the class with the highest aggregated weighted score was selected as the final segmentation

result. Given the complexity involved in simultaneously optimizing multiple network weights, Bayesian optimization was employed to achieve globally optimal parameters. Since the selected networks already exhibited consistently high and stable accuracy for fibers and matrix regions, further performance gains in these categories were unnecessary and potentially detrimental if overemphasized. Thus, to prevent fibers and matrix from dominating the voting process, the optimization range for their corresponding weights was limited to between 0.1 and 1. Moreover, to simplify the optimization task and enhance computational efficiency, a single shared weight was assigned to all five networks for each of the fiber and matrix classes. In contrast, individual network weights for the patch class were independently optimized over a wider range (0.1 to 3), enabling the ensemble to amplify the influence of a network particularly effective in segmenting patches, without excessively diminishing contributions from other networks. Consequently, the optimization involved seven parameters in total, constituting a high-dimensional search space ideally suited to Bayesian optimization methods.

During optimization, the F1-score for the patch regions in the final voting outcome was chosen as the objective function. The optimization process was conducted over 60 iterations, with an initial baseline established by setting all seven weights uniformly to 1. Figure 11a illustrates the progression of F1-scores for each of the three classes over the course of optimization. Clearly, after 60 iterations, the F1-score for patch prediction improved from the initial uniform-voting baseline of 0.595 to 0.638, representing a substantial increase of approximately 33% compared to the average single-network patch F1-score of 0.479 (cf. Figure 11b). Notably, modest performance gains for fibers and matrix regions were also observed. Table 3 provides the optimized values of the seven voting weights obtained from this process.

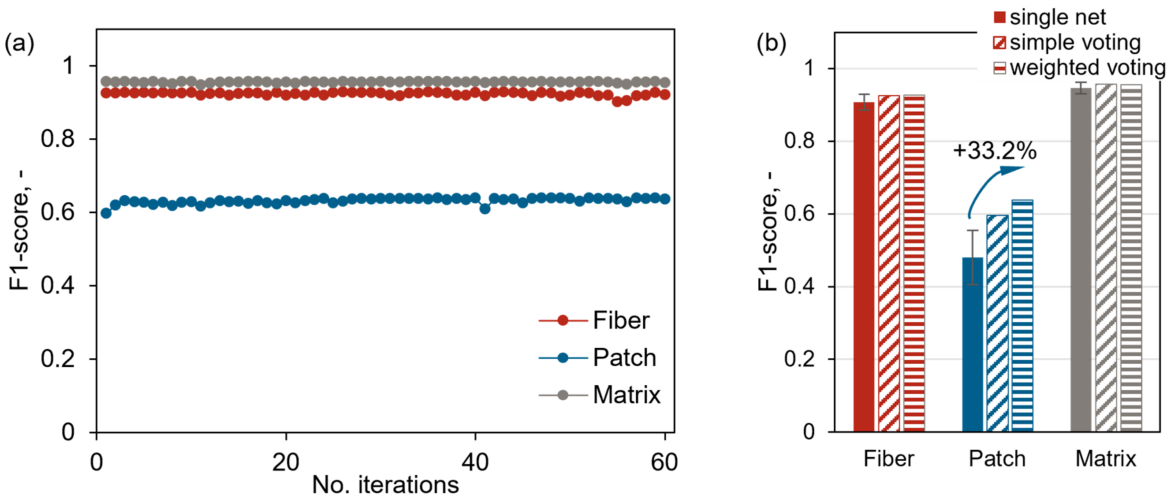


Figure 11. Progression of F1-score for three labels over 60 iterations of Bayesian optimization (**a**) and comparison of F1-scores for the three labels using a single network, simple voting, and weighted voting (**b**).

Table 3. Bayesian-optimized voting weights for five selected networks.

Weight	Fiber	Matrix	$Patch_1$	$Patch_2$	$Patch_3$	$Patch_4$	$Patch_5$
Value	0.212	0.216	0.280	1.763	0.336	0.365	2.750

To validate and assess the optimized weighted voting approach, and to ensure that the optimization did not lead to overfitting on the training data, an additional set of five images was prepared as independent test samples. These test images were not included within the original set of 63 training images, thereby allowing an unbiased evaluation of the

generalization capability of both the trained networks and the weighted voting ensemble. Figure 12 visually compares segmentation results from each of the five individual networks, the optimized weighted voting ensemble, the corresponding manually annotated ground truth, and the original optical micrographs. As shown in Figure 12, the predictions from each individual network were consistently accurate in segmenting fiber regions, exhibiting only minor differences across localized regions. Notably, all networks successfully differentiated between fibers and visually similar bright regions—primarily graphite flakes, as exemplified by large bright areas in the upper left and lower left regions of the optical image—and correctly avoided mislabeling these as fibers. Conversely, significant variations were observed in the identification and delineation of patch regions among the individual networks, with varying degrees of conservativeness or aggressiveness. However, the weighted voting ensemble effectively integrated these divergent predictions, resulting in a balanced segmentation that closely resembled the manual annotation. Figure 13 presents the confusion matrices for the optimized ensemble on the five test images, alongside the per-class precision, recall, and F_1-scores. The results demonstrate that, while maintaining fiber and matrix precision and recall above 0.93, the weighted voting scheme markedly boosts patch segmentation performance: patch precision rises to 0.917, recall to 0.737, and F_1 to 0.817. Examination of the confusion matrices further reveals minimal off-diagonal errors, indicating that most misclassifications occur between patch and matrix regions rather than with fibers. These improvements confirm that the Bayesian-optimized ensemble achieves a superior balance between precision and sensitivity for the most challenging class without compromising overall segmentation fidelity.

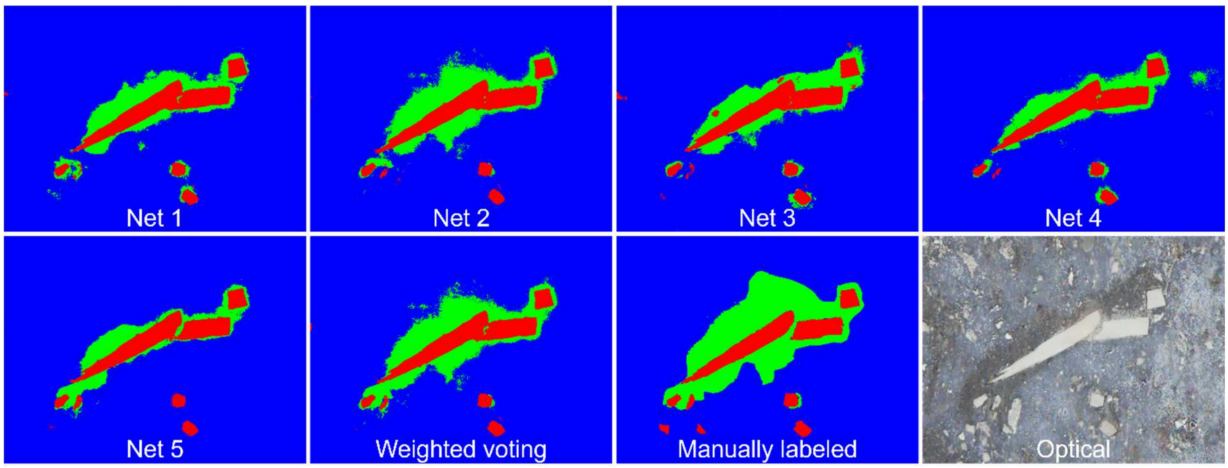


Figure 12. Example of visualized segmentation results of five single U-Net networks, weighted voting results, manually labeled ground truth, and original optical image.

Furthermore, segmentation accuracy for each of the three classes across all five test images was quantitatively evaluated, both individually for each network and collectively for the weighted voting ensemble, using the accuracy calculation defined as follows:

$$Accuracy_c = \frac{\sum_x II(y_x = c \ \wedge \ \hat{y}_x = c)}{\sum_x II(y_x = c)} \tag{13}$$

where y_x denotes the true label of pixel x, $\hat{y}_x$ denotes the predicted label, $\wedge$ is logic "and", and $II(\cdot)$ is the indicator function that equals 1 if its argument is true and 0 otherwise. The average accuracies calculated from this formula are presented in Figure 14. Compared to the average performance of the individual networks, the weighted voting ensemble demonstrated a slight improvement in accuracy for fibers, increasing from 0.943 to 0.953, while accuracy for the matrix class remained nearly unchanged at approximately 0.986. Crucially,

accuracy for the patch class showed a substantial enhancement, rising from 0.543 to 0.730, representing an increase of approximately 34.5%. Additionally, the weighted voting results exhibited enhanced stability and reduced variability, as evidenced by noticeably shorter error bars compared to individual network predictions.

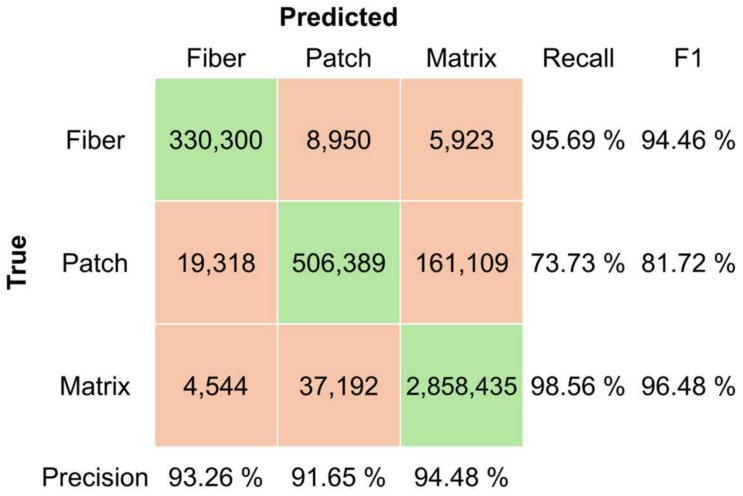


Figure 13. Confusion matrix of the optimized ensemble on the five test images.

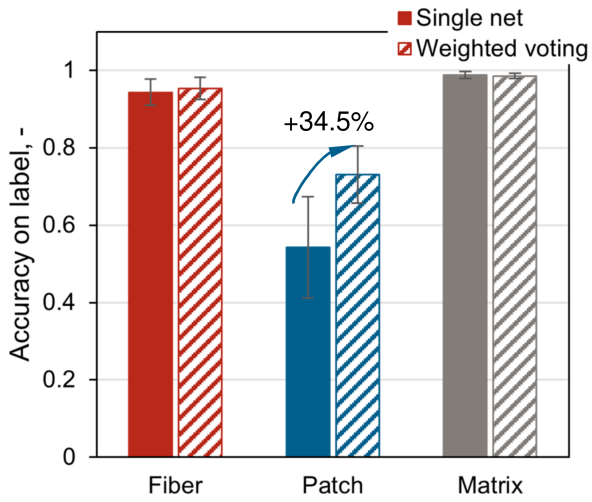


Figure 14. Calculated average accuracy on three labels of segmentations from single net and from weighted voting ensemble.

Given that the ultimate aim of this research is to apply neural networks to quantitatively investigate the influence of third-body patches on tribological performance, the accurate measurement of fiber and patch areas remains the primary focus. Therefore, the final weighted voting results for all test images and their corresponding manual annotations were visualized and prepared for subsequent morphometric analysis. Figure 15 compares the segmentation results for the remaining four test images, with the first row showing the ensemble predictions obtained from the optimized weighted voting scheme and the second row presenting the manually annotated ground truths. As illustrated, the ensemble results generated by neural network segmentation closely match the ground truths overall, demonstrating particularly high accuracy in fiber identification. Although minor discrepancies exist within certain patch regions, these do not significantly impact overall segmentation quality.

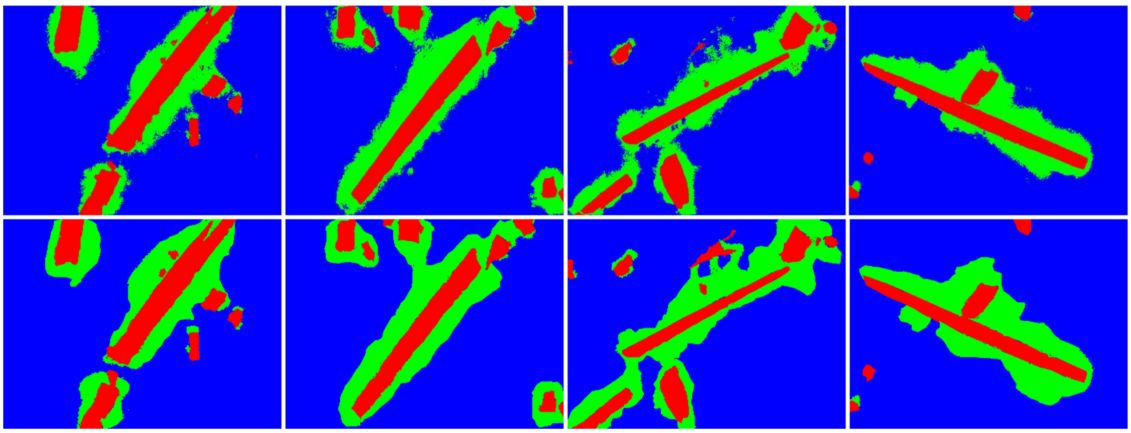


Figure 15. Comparison of weighted voting segmentation results (**first row**) with corresponding manual annotations (**second row**) for the four remaining test images.

To further validate the practical applicability of our proposed method in tribological studies, the areas of fibers and patches in the test dataset were quantitatively calculated from manual annotations, individual network predictions, and weighted voting results, respectively. Figure 16a shows the actual calculated fiber and patch areas identified by these three methods, while Figure 16b compares the area measurement mean absolute percentage errors (MAPE) for individual networks and the weighted ensemble relative to the manual ground truths. Clearly, the weighted voting results were significantly closer to the manual annotations than individual network predictions. Specifically, the error percentage in fiber area measurements, already low at an average of 2.3% from individual networks, improved to only 0.8% with weighted voting, reflecting practically negligible deviation. More notably, the patch area error dramatically improved from an unacceptable 63.3% with individual network predictions to just 7.2% after weighted voting. Although this error remains somewhat larger than that for fiber regions, it is considered acceptable. The relatively higher error in patch segmentation partly arises from the inherent limitations of pixel-level predictions, as some pixels near the ambiguous patch boundaries may not be confidently assigned to the patch region, increasing the complexity for subsequent image processing and analysis steps. Moreover, unlike fibers, patch boundaries are inherently ambiguous and cannot be defined with absolute clarity; even experienced researchers cannot guarantee perfect accuracy. Therefore, an error margin of approximately 7.2% relative to manual annotation is considered entirely acceptable within the scope of this research.

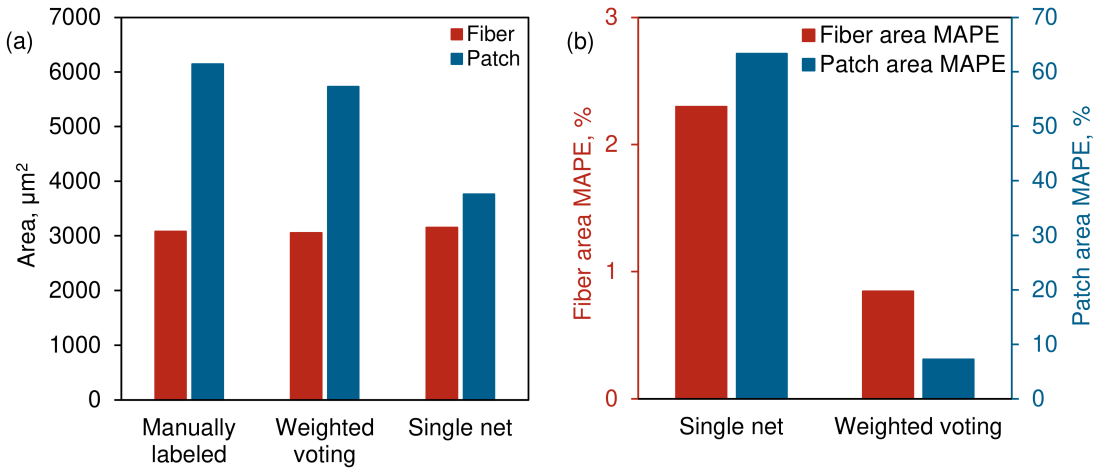


Figure 16. Comparison between different segment methods of (**a**) fiber and patch area measurements, and (**b**) percentage errors of fiber and patch areas relative to manual annotations.

4. Conclusions

In this study, a U-Net-based convolutional neural network was developed and optimized to perform precise, automated segmentation of worn-surface micrographs from Pin-on-Disc experiments on a PEEK-based composite. By leveraging high-resolution optical images combined with quantitative height data, several independent U-Net models were trained through five-fold cross-validation. Subsequently, a multi-network weighted voting system was implemented based on the five best-performing models. Bayesian optimization was employed to specifically fine-tune the weights assigned to different labels across networks, aiming primarily to enhance segmentation accuracy of patch regions while preserving performance for fiber and matrix areas. The main conclusions drawn from this study are as follows:

1. By carefully selecting optimal hyperparameters, adopting an appropriate U-Net architecture, and employing data augmentation along with cross-validation to mitigate overfitting, the resulting networks demonstrated high segmentation performance for fiber and matrix regions, achieving F1-scores of 0.912 and 0.946, respectively. However, segmentation accuracy for patch regions remained relatively low, with an F1-score of only 0.404.
2. Implementing Bayesian optimization to optimize and allocate seven individual weights for three segmentation classes across five selected U-Net models significantly improved patch region segmentation. The F1-score increased from an initial single network average of 0.404 to 0.638, representing an improvement of approximately 33%. Meanwhile, slight improvements in fiber and matrix segmentation accuracy were also observed, indicating that the weighted voting ensemble effectively harmonized the strengths and mitigated the weaknesses of individual networks.
3. Compared with single network segmentations, the optimized weighted voting scheme considerably reduced the MAPE of fiber area measurements from 2.3% to 0.8%, and dramatically improved patch area measurements from an initially unusable 63.3% down to a substantially more acceptable 7.2%. These improvements clearly indicate that our automated segmentation method provides sufficiently accurate morphological data for quantitative tribological analysis.

Ultimately, our findings demonstrate that a properly configured U-Net ensemble combined with Bayesian-optimized voting weights offers a high-precision tool for automated and accurate segmentation of complex worn surfaces. The proposed framework allows rapid and precise extraction of fiber and patch characteristics, laying a robust foundation for more efficient research workflows, real-time image analysis, and data-driven materials design in tribology. However, third-body patch boundaries lack clear definition as they transition into the surrounding matrix with low contrast, so even carefully produced ground-truth labels reflect a degree of subjectivity. Employing a multi-annotator consensus labeling strategy could reduce this subjectivity and further improve patch-region accuracy by providing the model with more robust boundary definitions. Future work will also validate the proposed segmentation framework on additional polymer-based composites to quantitatively assess cross-material robustness.

Author Contributions: Conceptualization, L.L.; Methodology, Y.Z. and L.L.; Investigation, Y.Z.; Resources, L.L.; Data curation, Y.Z.; Writing—original draft, Y.Z.; Writing—review & editing, L.L.; Visualization, Y.Z.; Supervision, L.L.; Project administration, L.L.; Funding acquisition, L.L. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by the German Research Foundation (DFG) through grant 499376717.

Data Availability Statement: The data presented in this study are available on request from the corresponding author.

Acknowledgments: We acknowledge the German Research Foundation (DFG) for the financial support of this research work (project number: 499376717). We also gratefully acknowledge Evonik Operations GmbH, Germany, Venator Germany GmbH, Germany, Kronos International Inc., Germany, SGL Carbon Fibers Ltd., Germany, and Superior Graphite Europe, Sweden for providing the experimental materials.

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

References

1. Voss, H.; Friedrich, K. On the wear behaviour of short-fiber-reinforced PEEK composites. *Wear* **1987**, *116*, 1–18. [CrossRef]
2. Friedrich, K.; Zhang, Z.; Schlarb, A.K. Effects of various fillers on the sliding wear of polymer composites. *Compos. Sci. Technol.* **2005**, *65*, 2329–2343. [CrossRef]
3. Myshkin, N.K.; Pesetskii, S.S.; Grigoriev, A.Y. Polymer tribology: Current state and applications. *Tribol. Ind.* **2015**, *37*, 284–290.
4. De Baets, P.; Glavatskih, S.; Ost, W.; Sukumaran, J. Polymers in tribology: Challenges and opportunities. In Proceedings of the 1st International Conference on Polymer Tribology, Bled, Slovenia, 11–12 September 2014; Volume 8.
5. Lu, Z.; Friedrich, K.; Küntzler, B. On sliding friction and wear of polyetheretherketone (PEEK) composites at elevated temperatures. *J. Synth. Lubr.* **1995**, *12*, 103–114. [CrossRef]
6. Zhang, M.Q.; Rong, M.Z.; Yu, S.L.; Wetzel, B.; Friedrich, K. Effect of particle surface treatment on the tribological performance of epoxy based nanocomposites. *Wear* **2002**, *253*, 1086–1093. [CrossRef]
7. Briscoe, B.J.; Yao, L.H.; Stolarski, T.A. Friction and Wear of Poly(Tetrafluoroethylene)-Poly(Etheretherketone) Composities; An Initial Appraisal of the Optimum Composition. *Wear Mater. Int. Conf. Wear Mater.* **1985**, *108*, 725–741.
8. Bahadur, S.; Polineni, V.K. Tribological studies of glass fabric-reinforced polyamide composites filled with CuO and PTFE. *Wear* **1996**, *200*, 95–104. [CrossRef]
9. Rodriguez, V.; Sukumaran, J.; Schlarb, A.K.; De Baets, P. Influence of solid lubricants on tribological properties of polyetheretherketone (PEEK). *Tribol. Int.* **2016**, *103*, 45–57. [CrossRef]
10. Jiang, Z.; Gyurova, L.A.; Schlarb, A.K.; Friedrich, K.; Zhang, Z. Study on friction and wear behavior of polyphenylene sulfide composites reinforced by short carbon fibers and sub-micro TiO_2 particles. *Compos. Sci. Technol.* **2008**, *68*, 734–742. [CrossRef]
11. Zhang, X.; Pei, X.; Wang, Q. Friction and Wear Behavior of Basalt-Fabric-Reinforced/Solid-Lubricant-Filled Phenolic Composites. *J. Appl. Polym. Sci.* **2010**, *117*, 3428–3433. [CrossRef]
12. Li, F.; Hu, Y.; Hou, X.; Hu, X.; Jiang, D. Thermal, mechanical, and tribological properties of short carbon fibers/PEEK composites. *High Perform. Polym.* **2018**, *30*, 657–666. [CrossRef]
13. Friedrich, K.; Lu, Z.; Häger, A.M. Overview on polymer composites for friction and wear application. *Theor. Appl. Fract. Mech.* **1993**, *19*, 1–11. [CrossRef]
14. Li, E.Z.; Guo, W.L.; Wang, H.D.; Xu, B.S.; Liu, X.T. Research on tribological behavior of PEEK and glass fiber reinforced PEEK composite. *Phys. Procedia* **2013**, *50*, 453–460. [CrossRef]
15. Li, J.; Zhang, L.Q. The research on the mechanical and tribological properties of carbon fiber and carbon nanotube-filled PEEK composite. *Polym. Compos.* **2010**, *31*, 1315–1320. [CrossRef]
16. Wang, Q.; Xue, Q.; Liu, H.; Shen, W.; Xu, J. The effect of particle size of nanometer ZrO_2 on the tribological behaviour of PEEK. *Wear* **1996**, *198*, 216–219. [CrossRef]
17. Chang, L.; Zhang, Z.; Zhang, H.; Schlarb, A.K. On the sliding wear of nanoparticle filled polyamide 66 composites. *Compos. Sci. Technol.* **2006**, *66*, 3188–3198. [CrossRef]
18. Padhan, M.; Marathe, U.; Bijwe, J. Tribology of Poly(etherketone) composites based on nano-particles of solid lubricants. *Compos. Part B Eng.* **2020**, *201*, 108323. [CrossRef]
19. Friedrich, K. Polymer composites for tribological applications. *Adv. Ind. Eng. Polym. Res.* **2018**, *1*, 3–39. [CrossRef]
20. Chang, L.; Friedrich, K. Enhancement effect of nanoparticles on the sliding wear of short fiber-reinforced polymer composites: A critical discussion of wear mechanisms. *Tribol. Int.* **2010**, *43*, 2355–2364. [CrossRef]
21. Briscoe, B.J.; Sinha, S.K. Chapter 1. Tribological applications of polymers and their composites: Past, present and future prospects. *Tribol. Interface Eng. Ser.* **2008**, *55*, 1–14. [CrossRef]
22. Harsha, A.P.; Wäsche, R.; Hartelt, M. Friction and wear studies of polyetherimide composites under oscillating sliding condition against steel cylinder. *Polym. Compos.* **2015**, *38*, 1–212. [CrossRef]
23. Guo, L.; Zhang, G.; Wang, D.; Zhao, F.; Wang, T.; Wang, Q. Significance of combined functional nanoparticles for enhancing tribological performance of PEEK reinforced with carbon fibers. *Compos. Part A Appl. Sci. Manuf.* **2017**, *102*, 400–413. [CrossRef]

24. Pei, X.Q.; Bennewitz, R.; Schlarb, A.K. Mechanisms of friction and wear reduction by carbon fiber reinforcement of PEEK. *Tribol. Lett.* **2015**, *58*, 42. [CrossRef]
25. Zhang, X.; Pei, X.; Wang, Q. Study on the friction and wear behavior of surface-modified carbon nanotube filled carbon fabric composites. *Polym. Adv. Technol.* **2011**, *22*, 2157–2165. [CrossRef]
26. Lin, L.; Zhao, Y.; Xu, Y.; Sun, C.; Schlarb, A.K. Advanced recycled carbon fiber (rCF) reinforced PEEK composites—Excellent alternatives for high-performance tribomaterials. *Mater. Today Sustain.* **2022**, *20*, 100227. [CrossRef]
27. Zhang, X. Study on the Tribological Properties of Carbon Fabric Reinforced Phenolic Composites Filled with Nano-Al_2O_3. *J. Macromol. Sci. Part B Phys.* **2017**, *56*, 568–577. [CrossRef]
28. Godet, M. The third-body approach: A mechanical view of wear. *Wear* **1984**, *100*, 437–452. [CrossRef]
29. Denape, J. Third body concept and wear particle behavior in dry friction sliding conditions. *Key Eng. Mater.* **2015**, *640*, 1–12. [CrossRef]
30. Krejčí, P.; Petrov, A. A mathematical model for the third-body concept. *Math. Mech. Solids* **2018**, *23*, 420–432. [CrossRef]
31. Martinez Ostormujof, T.; Purushottam Raj Purohit, R.R.P.; Breumier, S.; Gey, N.; Salib, M.; Germain, L. Deep Learning for automated phase segmentation in EBSD maps. A case study in Dual Phase steel microstructures. *Mater. Charact.* **2022**, *184*, 111638. [CrossRef]
32. Petkov, V.I.; Pakkam Gabriel, V.R.; Fernberg, P. Semantic segmentation of progressive micro-cracking in polymer composites using Attention U-Net architecture. *Tomogr. Mater. Struct.* **2024**, *5*, 100028. [CrossRef]
33. Gola, J.; Britz, D.; Staudt, T.; Winter, M.; Schneider, A.S.; Ludovici, M.; Mücklich, F. Advanced microstructure classification by data mining methods. *Comput. Mater. Sci.* **2018**, *148*, 324–335. [CrossRef]
34. Gola, J.; Webel, J.; Britz, D.; Guitar, A.; Staudt, T.; Winter, M.; Mücklich, F. Objective microstructure classification by support vector machine (SVM) using a combination of morphological parameters and textural features for low carbon steels. *Comput. Mater. Sci.* **2019**, *160*, 186–196. [CrossRef]
35. Baak, N.; Hajavifard, R.; Lücker, L.; Vasquez, J.R.; Strodick, S.; Teschke, M.; Walther, F. Micromagnetic approaches for microstructure analysis and capability assessment. *Mater. Charact.* **2021**, *178*, 111189. [CrossRef]
36. Pitz, E.; Pochiraju, K. A neural network transformer model for composite microstructure homogenization. *Eng. Appl. Artif. Intell.* **2024**, *134*, 108622. [CrossRef]
37. Li, Y.; Guan, J.; Guo, L. Peridynamic-driven feature-enhanced Vision Transformer for predicting defects and heterogeneous materials locations: Applications of deep learning in inverse problems. *Eng. Appl. Artif. Intell.* **2025**, *151*, 110677. [CrossRef]
38. Falk, T.; Mai, D.; Bensch, R.; Çiçek, Ö.; Abdulkadir, A.; Marrakchi, Y.; Böhm, A.; Deubner, J.; Jäckel, Z.; Seiwald, K.; et al. U-Net: Deep learning for cell counting, detection, and morphometry. *Nat. Methods* **2019**, *16*, 67–70. [CrossRef] [PubMed]
39. Zhou, Z.; Siddiquee, M.M.R.; Tajbakhsh, N.; Liang, J. UNet++: A Nested U-Net Architecture. In *Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support*; Springer International Publishing: Berlin/Heidelberg, Germany, 2018. [CrossRef]
40. Ronneberger, O.; Fischer, P.; Brox, T. U-Net: Convolutional Networks for Biomedical Image Segmentation. In *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015*; Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics); Springer International Publishing: Berlin/Heidelberg, Germany, 2015; pp. 234–241. [CrossRef]
41. Oktay, O.; Schlemper, J.; Folgoc, L.L.; Lee, M.; Heinrich, M.; Misawa, K.; Mori, K.; McDonagh, S.; Hammerla, N.Y.; Kainz, B.; et al. Attention U-Net: Learning Where to Look for the Pancreas. *arXiv* **2018**. [CrossRef]
42. Warren, P.; Raju, N.; Prasad, A.; Hossain, M.S.; Subramanian, R.; Kapat, J.; Manjooran, N.; Ghosh, R. Grain and grain boundary segmentation using machine learning with real and generated datasets. *Comput. Mater. Sci.* **2024**, *233*, 112739. [CrossRef]
43. Shi, P.; Duan, M.; Yang, L.; Feng, W.; Ding, L.; Jiang, L. An Improved U-Net Image Segmentation Method and Its Application for Metallic Grain Size Statistics. *Materials* **2022**, *15*, 4417. [CrossRef]
44. Wang, N.; Guan, H.; Wang, J.; Zhou, J.; Gao, W.; Jiang, W.; Zhang, Y.; Zhang, Z. A deep learning-based approach for segmentation and identification of δ phase for Inconel 718 alloy with different compression deformation. *Mater. Today Commun.* **2022**, *33*, 104954. [CrossRef]
45. Bertoldo, J.P.C.; Decencière, E.; Ryckelynck, D. A modular U-Net for automated segmentation of X-ray tomography images in composite materials. *Front. Mater.* **2021**, *8*, 761229. [CrossRef]
46. Li, H.; Wei, C.; Cao, Z.; Zhang, Y.; Li, X. Deep learning-based microstructure analysis of multi-component heterogeneous composites during preparation. *Compos. Part A Appl. Sci. Manuf.* **2024**, *186*, 108437. [CrossRef]
47. Dong, J.; Kandemir, A.; Hamerton, I. Microstructural characterisation of fibre-hybrid polymer composites using U-Net on optical images. *Compos. Part A Appl. Sci. Manuf.* **2025**, *190*, 108569. [CrossRef]
48. Firdaus-Nawi, M.; Noraini, O.; Sabri, M.Y.; Siti-Zahrah, A.; Zamri-Saad, M.; Latifah, H. DeepLabv3+_Encoder-Decoder with Atrous Separable Convolution for Semantic Image Segmentation. *Pertanika J. Trop. Agric. Sci.* **2011**, *34*, 137–143.

49. Lin, L.; Schlarb, A.K. Development and optimization of high-performance PEEK/CF/Nanosilica hybrid composites. *Polym. Adv. Technol.* **2021**, *32*, 3150–3159. [CrossRef]
50. *ASTM G99-17*; Standard Test Method for Wear Testing with a Pin-on-Disk Apparatus. ASTM: West Conshohocken, PA, USA, 2023.
51. Gholamalinezhad, H.; Khosravi, H. Pooling Methods in Deep Neural Networks, a Review. *arXiv* **2020**. [CrossRef]
52. Tasci, E.; Uluturk, C.; Ugur, A. A voting-based ensemble deep learning method focusing on image augmentation and preprocessing variations for tuberculosis detection. *Neural Comput. Appl.* **2021**, *33*, 15541–15555. [CrossRef]
53. Coupé, P.; Mansencal, B.; Clément, M.; Giraud, R.; de Senneville, B.D.; Ta, V.T.; Lepetit, V.; Manjon, J.V. AssemblyNet: A large ensemble of CNNs for 3D whole brain MRI segmentation. *NeuroImage* **2020**, *219*, 117026. [CrossRef]
54. Dang, T.; Nguyen, T.T.; Moreno-García, C.F.; Elyan, E.; McCall, J. Weighted Ensemble of Deep Learning Models based on Comprehensive Learning Particle Swarm Optimization for Medical Image Segmentation. In Proceedings of the 2021 IEEE Congress on Evolutionary Computation, CEC 2021—Proceedings, Kraków, Poland, 28 June–1 July 2021; IEEE: Piscataway, NJ, USA, 2021; pp. 744–751. [CrossRef]
55. Bergstra, J.; Bardenet, R.; Bengio, Y.; Kégl, B. Algorithms for hyper-parameter optimization. *Adv. Neural Inf. Process. Syst.* **2011**, *24*, 1–9.
56. Mockus, J. The Bayesian Approach to Global Optimization. In *System Modeling and Optimization*; Springer Science and Business Media: Berlin/Heidelberg, Germany, 1981.
57. Pelikan, M.; Goldberg, D.E.; Cantú-Paz, E. Hierarchical problem solving by the Bayesian optimization algorithm. In Proceedings of the 2nd Annual Conference on Genetic and Evolutionary Computation, Las Vegas, NV, USA, 10–12 July 2000.

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.