

Label Efficient Segmentation of SiC–C–Si Phases in SiC-based Ceramics from Optical Micrographs Using a Compact U-Net

A.V. Kaledin

kaledin@issp.ac.ru

Osipyan Institute of Solid State Physics RAS

S.M. Peshcherova

Irkutsk State University

D.O. Stryukov

Osipyan Institute of Solid State Physics RAS

S.L. Shikunov

Osipyan Institute of Solid State Physics RAS

Yu.N. Zubareva

Osipyan Institute of Solid State Physics RAS

V.N. Kurlov

Osipyan Institute of Solid State Physics RAS

I.S. Petrushin

Irkutsk State University

Research Article

Keywords: SiC-based ceramics, Optical microscopy, Weak supervision, U, nobreakdash-Net, Microstructure segmentation, Phase quantification

Posted Date: October 19th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-10241263/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: No competing interests reported.

Label Efficient Segmentation of SiC–C–Si Phases in SiC-based Ceramics from Optical Micrographs Using a Compact U-Net

A.V. Kaledin^{1*}, S.M. Peshcherova², D.O. Stryukov¹, S.L. Shikunov¹,
Yu.N. Zubareva¹, V.N. Kurlov¹, I.S. Petrushin²

¹Osipyan Institute of Solid State Physics RAS, Academician Osipyan 2, Chernogolovka,
142432, Russia.

²Faculty of Business Communication and Informatics, Irkutsk State University, Karl
Marx 1, Irkutsk, 664003, Russia.

*Corresponding author(s). E-mail(s): kaledin@issp.ac.ru;

Abstract

Purpose: This work introduces a label-efficient workflow for semantic segmentation of silicon carbide, free silicon, and carbon in SiC-based ceramic from reflected-light micrographs, aimed at reproducible phase-fraction quantification under varying photometric conditions. **Methods:** Weak supervision is derived via unsupervised clustering in the CIE L*a*b* (CIELAB) color space, followed by light morphological refinement to obtain proxy masks for training. A compact four-stage U-Net (encoder–decoder with skip connections) is then trained on RGB micrographs using a common preprocessing pipeline. Final performance is assessed against expert-annotated ground-truth masks, reporting per-pixel agreement, mean Intersection over Union (IoU), and silicon phase-fraction error. **Results:** The U-Net architecture demonstrates a substantial improvement over multilevel Otsu thresholding in low-contrast fields and provides fully automatic, reproducible segmentation. Although marker-controlled watershed segmentation can achieve comparable or, in some cases, lower phase-fraction errors, the U-Net approach provides a significantly faster end-to-end workflow suitable for routine analysis. The resulting error in silicon phase fraction is typically below 2 vol.% for most fields, with larger deviations observed only in the coarsest microstructures. End-to-end processing of a full-resolution micrograph (9–16 Mpx) requires approximately 0.3–0.6 s on a mid-range GPU and 3–6 s on an 8-core CPU. **Conclusion:** The proposed workflow offers a scalable and interpretable approach for rapid microstructural assessment and process screening in SiC-based ceramics. Future work will focus on extended expert validation and detailed, figure-rich case studies.

Keywords: SiC-based ceramics; Optical microscopy; Weak supervision; U-Net; Microstructure segmentation; Phase quantification

1 Introduction

Silicon carbide (SiC)-based ceramics are extensively employed in high-temperature and wear-resistant applications due to their unique combination of low density ($2.5\text{--}3.2\text{ g cm}^{-3}$), exceptional hardness ($25\text{--}30\text{ GPa}$) and outstanding oxidation resistance up to 1600°C . This oxidation resistance arises from the in situ formation of a protective silica scale [1–4]. To optimize these properties, various processing methods have been developed, each balancing densification temperature, additive chemistry and dimensional precision. Recrystallized SiC (RSiC), typically sintered above 2200°C , achieves high purity ($> 96\text{ vol.}\%$) but remains porous ($20\text{--}50\text{ vol.}\%$), resulting in moderate mechanical strength [5, 6]. Hot-pressed SiC (HPSiC) attains superior flexural strength due to pressure-assisted densification but is limited in fabricating complex geometries [7–9]. Pressureless-sintered SiC (SSiC) can achieve comparable properties; however, it undergoes $5\text{--}15\%$ linear shrinkage during processing, complicating near-net-shape manufacturing [10–13].

Liquid-silicon infiltration (LSI) of porous preforms presents an alternative approach, minimizing shrinkage and enabling premachining prior to infiltration. In reaction-bonded SiC (RB-SiC), the preform comprises both SiC and carbon, whereas in siliconized graphite, it consists solely of carbon [14–21]. LSI facilitates the production of SiC-based ceramics across a broad density range and at lower raw-material cost, yielding materials with $75\text{--}92\text{ vol.}\%$ SiC and flexural strengths between $317\text{--}389\text{ MPa}$ [18, 19, 22]. The resulting microstructure typically contains three main phases: silicon carbide (predominantly β -SiC at processing temperatures), free silicon ($5\text{--}30\text{ vol.}\%$), and residual carbon ($1\text{--}50\text{ vol.}\%$). The thermal and electrical conduction, oxidation behavior, and mechanical performance are strongly influenced by the volume fractions and connectivity of these phases. For structural applications, a high SiC content with minimal residual silicon and carbon is preferred, whereas tribological components may benefit from residual graphite. Applications requiring specific electrical or thermal properties demand careful control of the SiC-carbon balance to tailor resistivity and thermal conductivity.

Quantitative phase segmentation from reflected-light micrographs is critical for process control and materials optimization. Although manual thresholding and interactive editing remain widespread, these methods are labour-intensive and vulnerable to operator bias and variations in photometric variability. Traditional image-processing, including global or multilevel thresholding and gradient-driven watershed segmentation, often underperform when gray-level distributions overlap or when the silicon network exhibits thin necks connecting to bright pools. Deep learning has revolutionized automated microstructure analysis; however, many existing approaches focus on binary or high-contrast multiphase systems, require dense pixel-level annotation, or emphasize geometric descriptors over chemically resolved phases. Previous studies on cast Al-Si alloys, 3D FIB-SEM segmentation, and oxide ceramics highlight these limitations, as do workflows relying heavily on post-processing or physics-based supervision [23–27]. Recent reviews identify chemically resolved, label-efficient segmentation of low-contrast multiphase ceramics as a persistent challenge [28–31].

This work addresses this gap by introducing a label-efficient workflow that generates reproducible SiC/C/Si phase maps directly from reflected-light images without requiring dense annotation. Weak pixel labels (“proxy masks”) are obtained through unsupervised k -means clustering in the CIE $L^*a^*b^*$ (CIELAB) color space, followed by light morphological refinement. These proxy masks are used exclusively for training a compact four-stage U-Net under a uniform pre-processing protocol. For final evaluation, a held-out set of expert-annotated masks serves as the ground truth, providing an unbiased comparison of all methods. Our contributions are: (i) a chemically resolved and label-efficient segmentation pipeline tailored to low-contrast SiC/C/Si microstructures imaged via LSI; (ii) a physically interpretable diagnostic framework linking photometric separability and interphase geometry to known failure modes of classical methods; and (iii) a unified evaluation protocol demonstrating stable phase-fraction estimates with inference times compatible with routine industrial process monitoring.

2 Materials and Methods

2.1 Sample Preparation and Microscopy

Graphite blocks (Technocarb, Chelyabinsk, Russia) were crushed in a hammer mill and sieved into three particle size fractions ($< 50 \mu\text{m}$, $50\text{--}100 \mu\text{m}$ and $100\text{--}200 \mu\text{m}$). These fractions were dry-blended with phenolic resin powder in a drum mixer (Fig. 1a). The resulting mixtures were cold-pressed in a steel die, thermally cured, and pyrolysed to yield porous carbon preforms (Fig. 1b). All preforms were machinable to near-net-shape. Vacuum liquid-silicon infiltration (LSI) was carried out at 1450°C under vacuum of 10^{-2} Pa (Fig. 1d). Following infiltration and minor machining (Fig. 1e), volumetric shrinkage remained below 1 %. The final microstructure comprised silicon carbide (SiC) formed around the original graphite particles, residual unreacted carbon (C), and free silicon (Si).

Polished cross-sections were prepared by sequential grinding using boron carbide abrasives (P120, P150, P320, P600, P1500), followed by polishing with diamond suspensions of $10 \mu\text{m}$, $5 \mu\text{m}$ and $1 \mu\text{m}$ particle sizes. Reflected-light bright-field micrographs were acquired at nominal magnification of $50\times$ and $100\times$, captured as 8-bit RGB images with a resolution of 3488×2616 pixels. Absolute dimensions were calibrated with scale bars, and all subsequent image analyses were performed using pixel units. Variations in focus and illumination between fields were minimized through per-image white balancing and contrast normalization.

A total of 900 non-overlapping fields (300 per material variant) were recorded. Representative microstructure variations, illustrating a range from fine to coarse morphologies at two magnifications, are shown in Figure 2. These variations motivated the adoption of a label-efficient workflow: initial weak labels were generated via unsupervised clustering in CIELAB color space and subsequently refined using a compact U-Net neural network for supervised segmentation.

2.2 Segmentation Difficulty Descriptors

To quantitatively evaluate segmentation difficulty with reproducible and physically meaningful metrics, we assessed (i) inter-class separability in CIELAB color space, (ii) illumination non-uniformity at the field scale, (iii) boundary complexity of the interphase network, and (iv) morphological statistics reflecting feature widths and heterogeneity. All descriptors were calculated in pixel units and, where applicable, converted to physical units using the microscope calibration scale s (in $\mu\text{m px}^{-1}$).

Let $\mathbf{x}_i = [L_i, a_i, b_i]^T$ denote the CIELAB vector of pixel i . For class c , the centroid and within-class variance were

$$\boldsymbol{\mu}_c = \frac{1}{n_c} \sum_{i \in c} \mathbf{x}_i, \quad \sigma_c^2 = \frac{1}{n_c} \sum_{i \in c} \|\mathbf{x}_i - \boldsymbol{\mu}_c\|_2^2.$$

Inter-class color separation between classes c_1 and c_2 was measured by the Euclidean distance

$$\Delta E_{ab}^*(c_1, c_2) = \|\boldsymbol{\mu}_{c_1} - \boldsymbol{\mu}_{c_2}\|_2, \quad (1)$$

and the contrast-to-noise ratio

$$\text{CNR}(c_1, c_2) = \frac{\|\boldsymbol{\mu}_{c_1} - \boldsymbol{\mu}_{c_2}\|_2}{\sqrt{\sigma_{c_1}^2 + \sigma_{c_2}^2}}. \quad (2)$$

To complement the moment-based analysis with a distributional perspective, the projections of $\mathbf{x}_i$ onto $\mathbf{u} = (\boldsymbol{\mu}_{c_1} - \boldsymbol{\mu}_{c_2}) / \|\boldsymbol{\mu}_{c_1} - \boldsymbol{\mu}_{c_2}\|_2$, were compared across the estimated one-dimensional class densities $p_c(t)$, obtained using Gaussian kernel density estimation with bandwidth selected according to Silverman's rule, and computed the histogram overlap

$$\text{OVL}(c_1, c_2) = \int_{-\infty}^{\infty} \min\{p_{c_1}(t), p_{c_2}(t)\} dt, \quad (3)$$

where $\text{OVL}(c_1, c_2) \in [0, 1]$.

Smaller ΔE_{ab}^* and CNR in (1)–(2) together with larger OVL in (3) indicate poorer separability.

Illumination non-uniformity was characterized by the coefficient of variation of a low-frequency illumination field $F(x, y)$, obtained by Gaussian




Fig. 1: Manufacturing route of SiC ceramics via liquid silicon infiltration (LSI) of porous carbon preforms. (a) blending graphite powder with a phenolic resin binder; (b) cold pressing, thermal curing, and pyrolysis to produce a porous carbon preform; (c) near-net-shape machining; (d) vacuum LSI at 1450 °C and 10^{-2} Pa, during which carbon is converted to SiC and residual porosity is filled with silicon; (e) final machining and polishing. Color code: grey — SiC, black — C, blue — Si.

smoothing of the luminance channel with a standard deviation larger than the characteristic grain size ($\sigma \approx 20\text{--}30$ pixels) after masking non-specimen regions and specular highlights:

$$CV = \frac{\text{std}(F(x, y))}{\text{mean}(F(x, y))} \times 100\%. \quad (4)$$

The boundary complexity was quantified by the density of the interface length L/A . The inter-phase boundaries were extracted from the union of pairwise class interfaces in manually corrected masks. The sub-pixel perimeter L was calculated using a marching square estimator and normalized by the field area A .

Morphology statistics were evaluated for the silicon phase. The local thickness $t(\mathbf{p})$ along the pruned medial axis was defined as $2DT(\mathbf{p})$, where DT is the Euclidean distance transform to the nearest phase boundary. Distributions of t were summarized by the median t_{50} and percentile range $t_{10}\text{--}t_{90}$ in pixels. These descriptors informed method selection: lower phase separability and stronger illumination variability justified the use of CIELAB clustering and photometric normalization; higher L/A values and small-thickness tails supported conservative morphological filtering of

weak labels and motivated the inclusion of skip connections in the segmentation network.

2.3 Image Processing, Weak Labels and Segmentation Model

Per-image white balance and contrast normalization were performed in RGB before conversion to the CIELAB color space. Since dense annotation of 900 micrographs was impractical, proxy (weak) labels were generated using k -means clustering with $k = 3$ and k -means++ initialization, retaining the run with the minimum within-cluster sum of squares:

$$\begin{aligned} \mathcal{J} &= \sum_{c=1}^3 \sum_{\mathbf{x}_i \in \mathcal{C}_c} \|\mathbf{x}_i - \boldsymbol{\mu}_c\|_2^2, \\ \boldsymbol{\mu}_c &= \frac{1}{|\mathcal{C}_c|} \sum_{\mathbf{x}_i \in \mathcal{C}_c} \mathbf{x}_i. \end{aligned} \quad (5)$$

Binary masks were refined by an opening-closing morphological sequence with a circular structuring element of radius $r = 5$ pixels. Regions smaller than 0.05 % of the image area were excluded. This choice of r effectively suppressed speckle noise while preserving narrow silicon bridges. To assess the suitability of




Fig. 2: Representative reflected-light micrographs used in this study. (a-c) Finer morphology (50 $\times$); (d-f) coarser morphology (100 $\times$). SiC appears grey, carbon dark and silicon bright.

the resulting proxy masks for supervised learning, a held-out set of thirty fields with manually corrected masks was used as a reference.

A compact four-stage U-Net architecture was trained to segment SiC, C and Si directly from the color micrographs. The encoder comprised two Conv-BN-ReLU blocks per level, followed by 2×2 max pooling, with feature channels 32, 64, 128 and 256. The decoder was symmetric, employing transpose-convolution up-sampling and skip connections to maintain precise interphase boundaries. A dropout layer ($p = 0.2$) was inserted at the bottleneck to improve generalization.

The dataset was stratified by material variant into training, validation, and test subsets (720/90/90 images). During training, 256×256 crops were sampled on the fly with random horizontal and vertical flips, rotations within $\pm 10^\circ$,

and mild photometric jittering (contrast $\pm 5\%$, brightness $\pm 3\%$). To mitigate the influence of noisy proxy labels near interphase boundaries, smoothed class-indicator targets $\varepsilon = 0.1$ were applied. Denoting by $\tilde{y}_{ijc}$ the target for pixel (i, j) and class c , the loss function was the categorical cross-entropy

$$\mathcal{L} = -\frac{1}{HW} \sum_{i=1}^H \sum_{j=1}^W \sum_{c=1}^3 \tilde{y}_{ijc} \log p_{ijc}, \quad (6)$$

$$\tilde{y}_{ijc} = \begin{cases} 1 - \varepsilon, & c = c^*(i, j), \\ \varepsilon/2, & c \neq c^*(i, j), \end{cases}$$

where p_{ijc} are the softmax probabilities and $c^*(i, j)$ is the proxy class. Optimization used the

Adam algorithm with a learning rate of 10^{-4} and $(\beta_1, \beta_2) = (0.9, 0.999)$, a batch size of 4, and early stopping after 10 epochs without validation improvement (up to a maximum of 100 epochs).

During inference, full-resolution images were segmented using a sliding window of 256×256 pixels with a 128-pixel overlap, blending probability maps across overlapping regions. On an NVIDIA RTX 3060 (12 GB), a full micrograph of 9–16 Mpx was processed in approximately 0.3–0.6 s; on an 8-core CPU, in 3–6 s.

For the final quantitative evaluation, 50 fields from the test set were manually annotated by an expert. These expert masks served as the reference for computing pixel-wise accuracy, per-class intersection-over-union, and absolute phase-fraction error for phase c ,

$$\Delta\phi_c = |\phi_c^{\text{pred}} - \phi_c^{\text{ref}}|, \quad (7)$$

$$\phi_c = \frac{1}{HW} \sum_{i=1}^H \sum_{j=1}^W \mathbf{1}\{\hat{y}_{ij} = c\},$$

providing complementary insights into pixel-level accuracy, interphase delineation, and quantitative phase fraction estimation. All reported metrics in Section 3 are based on these expert masks.

The study design is summarized in Fig. 3, illustrating all stages — from material preparation and imaging to weak-label generation, model training, and evaluation — performed under identical photometric conditions.

The source code used in this study is publicly available on GitHub: <https://github.com/ipetrushin/Ceramics-analysis>.

3 Results and Discussion

3.1 Segmentation Performance

Quantitative descriptors from Section 2.1 were calculated relative to expert masks and are summarized alongside the performance of all methods in Table 1. The silicon phase fraction, ϕ_{Si} , varies from 6.56% (Fig. 2a) to 23.83% (Fig. 2f), confirming that a single global threshold is inadequate. The inter-class color separation, $\Delta E_{ab}^*(\text{Si}, \text{SiC})$, ranges from 9.87 (Fig. 2d) to 29.54 (Fig. 2e); larger values indicate more reliable global color cues for discriminating silicon from silicon carbide. The interface length density, L/A , is the




Fig. 3: Study design employed in this work. The diagram provides an overview of the complete workflow, encompassing material preparation, imaging, preprocessing, and weak-label generation, followed by data splitting with augmentation, model training, full-image inference, and quantitative evaluation.

sub-pixel perimeter of the union of pairwise class boundaries per 10^3 pixels squared; it is highest for the fine network in Fig. 2b (58.17) and lowest for the coarsest structure in Fig. 2f (26.22). Silicon–ligament thickness is computed along the phase skeleton as $t(\mathbf{p}) = 2d(\mathbf{p})$, where $d(\mathbf{p})$ is the Euclidean distance to the nearest phase boundary. The median t_{50} ranges from 6.32 px (Fig. 2b,d) to 20.88 px (Fig. 2f), while the wide intervals t_{10} – t_{90} (e.g., 2.00–23.32 px in Fig. 2b and 2.83–90.44 px in Fig. 2f) reveal persistent coexistence of thin necks with thick pools – a known failure mode for threshold- or edge-driven segmentation routines.

Table 1 reveals clear, physically interpretable trends. As the morphology coarsens from Fig. 2a to Fig. 2f, the interface density L/A generally decreases (from 41.53 to 26.22), while the median silicon thickness t_{50} increases (from 8.94 to 20.88 px). At the same time, the wide intervals t_{10} – t_{90} across all fields indicate the coexistence of thin silicon necks and thick silicon regions, even in coarse

microstructures. This heterogeneity poses a challenge for segmentation approaches based solely on thresholding or edge detection.

The performance of classical methods in Table 1 correlates with these descriptors. When ΔE_{ab}^* is small and t_{10} is low (Fig. 2a and Fig. 2d), multilevel Otsu thresholding fails dramatically: its accuracy drops to 32.50% and 28.85%, respectively, with silicon fraction errors of 64.53 and 68.33 vol.%. This failure cannot be attributed only to the removal of thin silicon necks. In these low-contrast fields, multilevel Otsu thresholding fails to separate silicon as an independent phase because the color distributions of Si and SiC strongly overlap. As a result, the segmentation effectively collapsed into a two-phase partition, with silicon being merged with silicon carbide or incorrectly mapped during class assignment. This explains the large silicon phase-fraction errors observed for Fig. 2a and Fig. 2d. In contrast, marker-controlled watershed segmentation on the luminance gradient is much more robust in these low-contrast fields, achieving accuracies above 93% and errors below 0.5 vol.% for Fig. 2a and Fig. 2d. For coarser structures (Fig. 2e,f), both Otsu and watershed perform well (accuracy > 91%, error < 6.1 vol.%), though the wide t_{10} – t_{90} intervals indicate that local failures due to thin connections remain possible.

These observations carry direct methodological implications. The median t_{50} captures the geometric scale and correlates with the overall segmentation ease, yet it is the left-tail thickness t_{10} that predicts the onset of connectivity loss. Interface density L/A acts as an independent stressor, elevating the number of boundary decisions per unit area and thus the likelihood of local topological errors. A small ΔE_{ab}^* diminishes the discriminative power of a global threshold, fostering systematic SiC to Si substitutions in low-contrast regions. Collectively, the descriptors and the outcomes in Table 1 elucidate performance variability across morphologies and magnifications in classical methods (Fig. 2a–f), underscoring the inherent fragility of rule-based pipelines on these micrographs.

These observations motivate a learning-based approach. The following subsection evaluates a compact U-Net on the same fields and compares

its performance with the classical methods, focusing on phase-fraction accuracy and topological preservation.

3.2 Learning-Based Segmentation Using a Compact U-Net

To enable a direct and fair comparison with classical methods, we evaluate the U-Net on the same 50 expert-annotated fields used in Section 3.1. All methods are applied with identical preprocessing and the same sliding-window inference protocol. The expert masks provide the ground truth for computing accuracy, IoU, and phase-fraction error. In this subsection, we compare our compact U-Net model with multilevel Otsu thresholding in CIELAB space and marker-controlled watershed segmentation on ∇L^* (both applied with the same preprocessing steps described in Section 2.1). This approach ensures that the reported metrics reflect algorithmic performance rather than annotator variability.

The proposed network is a four-stage encoder–decoder comprising two Conv–BN–ReLU blocks per resolution level with symmetric skip connections (Fig. 4). Inference is performed at full resolution using a 256×256 sliding window with a 128-px overlap. Probability maps are blended across overlaps, and the final class label for each pixel is assigned as the arg-max of the softmax probabilities. Using identical preprocessing for all methods minimizes method-specific photometric bias tuning and ensures a fair comparison.

We report pixel-wise accuracy, mean IoU, and the absolute silicon phase-fraction error $|\Delta\phi_{\text{Si}}|$ (vol.%) relative to the expert masks. Additionally, we evaluate the repeatability of the U-Net output by measuring the spread of $|\Delta\phi|$ and mean IoU under test-time augmentations (rotations $\pm 10^\circ$, flips, contrast $\pm 5\%$, brightness $\pm 3\%$). For process monitoring applications, repeatability is particularly critical, as the segmentation method must provide consistent quantitative results under identical imaging conditions.

Table 1 summarizes the comparison across the six fields shown in Fig. 2(a–f). The U-Net achieves competitive accuracy and mIoU across all fields. In the most challenging cases for Otsu (Fig. 2a and Fig. 2d), where Otsu fails completely, the U-Net maintains accuracy above 76% and 81%, respectively, and keeps the silicon fraction error

Table 1: Microstructural descriptors (computed from expert masks) and segmentation performance of classical methods and the proposed U-Net on the same six representative fields (test set, expert masks as ground truth). Metrics: accuracy (%); IoU (for U-Net); absolute silicon phase-fraction error $|\Delta\phi|$ (vol.%).

	Fig. 2a	Fig. 2b	Fig. 2c	Fig. 2d	Fig. 2e	Fig. 2f
ϕ_{Si} (vol.%)	6.56	9.88	20.81	7.36	10.59	23.83
ΔE_{ab}^* (Si-SiC)	10.32	18.95	18.86	9.87	29.54	17.16
L/A (px/ 10^3px^2)	41.53	58.17	53.05	35.05	37.53	26.22
t_{50} (px)	8.94	6.32	10.00	6.32	10.00	20.88
$t_{10}-t_{90}$ (px)	2.83–32.56	2.00–23.32	2.00–46.39	2.00–26.00	2.00–48.41	2.83–90.44
Otsu Acc. (%)	32.50	97.76	95.66	28.85	91.48	96.96
Otsu $ \Delta\phi $ (vol.%)	64.53	0.44	1.05	68.33	6.10	0.99
Watershed Acc. (%)	93.93	95.19	95.04	94.34	93.54	97.09
Watershed $ \Delta\phi $ (vol.%)	0.41	1.26	1.37	0.48	4.03	1.15
U-Net Acc. (%)	76.15	95.69	74.70	81.15	92.09	93.62
U-Net mIoU	0.481	0.898	0.589	0.543	0.814	0.855
U-Net $ \Delta\phi $ (vol.%)	1.36	0.58	0.81	0.94	5.51	4.63




Fig. 4: Architecture of the compact U-Net employed in this study. The network consists of four stages featuring double Conv-BN-ReLU blocks, a symmetric decoder with skip connections, and a final 1×1 convolution followed by a softmax layer producing three output classes.

below 1.4 vol.%. Compared with watershed, which already performs well in these fields, the U-Net does not universally improve the pixel-wise metrics or the silicon phase-fraction error. However, watershed is not an end-to-end automatic procedure in routine use: it requires parameter tuning and processing time can reach an hour per full-resolution micrograph. In contrast, after training, the U-Net segments a full-resolution micrograph automatically within only a few seconds. Thus,

the main advantage of the proposed network is not a universal improvement over watershed in all accuracy metrics, but a substantial reduction in processing time while maintaining acceptable phase-fraction accuracy and high repeatability. For coarser structures (Fig. 2e,f), the U-Net’s silicon fraction error (5.51 and 4.63 vol.%) is larger than that of both Otsu (6.10 and 0.99 vol.%) and watershed (4.03 and 1.15 vol.%). This indicates that the current U-Net training, based on proxy labels, may not fully capture the variability of coarse morphologies. Nevertheless, its mIoU remains high (0.814 and 0.855), suggesting good boundary delineation. The U-Net also exhibits excellent repeatability: under test-time augmentations, the spread of $|\Delta\phi|$ is within ± 0.3 – 0.6 vol.% (not tabulated), which is critical for process monitoring.

From an application standpoint, the U-Net provides a fully automatic segmentation with inference times of 0.3–0.6 s per full micrograph on a GPU, substantially faster than manual threshold tuning, and offers a robust solution for routine phase quantification when combined with quality control measures.

4 Conclusions

This study introduced a label-efficient pipeline for semantic segmentation of SiC–C–Si phases in siliconized graphite from reflected-light micrographs. Weak supervision was achieved through

unsupervised clustering in the CIELAB color space, followed by light morphological refinement, while a compact four-stage U-Net was trained to segment the three phases directly from color images. All methods were evaluated under identical preprocessing conditions and against expert-annotated ground truth masks (50 fields), ensuring an unbiased and reproducible comparison. The results show that the U-Net is particularly advantageous relative to global thresholding in low-contrast cases where Otsu fails to separate silicon as an independent phase. Compared with marker-controlled watershed, the U-Net does not provide a universal improvement in all pixel-level metrics, but it offers a major practical advantage: fully automatic processing within seconds, whereas watershed-based analysis requires parameter adjustment and can take tens of minutes per field when manual correction is needed. For coarser structures, the U-Net’s phase-fraction error was slightly larger than that of classical methods, but its high mIoU and low test-time variability indicate reliable boundary delineation and suitability for process monitoring. These results align with the diagnostic indicators: low ΔE_{ab}^* and small t_{10} increase the likelihood of threshold-based failures, whereas the encoder–decoder architecture with skip connections effectively integrates local texture and mesoscale context. The pipeline is computationally efficient: full micrograph processing requires 0.3–0.6 s on a GPU, representing an order-of-magnitude speedup over manual editing. Two limitations should be acknowledged. First, the evaluation was performed on a limited set of 50 expert masks; extending the annotated dataset would further strengthen statistical significance. Second, performance on the coarsest field revealed sensitivity to acquisition conditions, highlighting the importance of routine color-constancy checks and lightweight domain adaptation when imaging parameters vary. Future work will focus on extending the weakly supervised framework to additional constituents (e.g., oxide films or pores), developing volumetric inference from serial sections or tomography, and integrating segmentation outputs into digital process twins to accelerate the design and optimization of SiC-based components.

References

- [1] Roy J, Chandra S, Das S, Maitra S (2014) Oxidation behaviour of silicon carbide—a review. *Rev Adv Mater Sci* 38:29–39.
- [2] Ruys AJ (2023) *Silicon Carbide Ceramics: Structure, Properties and Manufacturing*. Elsevier, Amsterdam.
- [3] Carter CB, Norton MG (2013) *Ceramic Materials: Science and Engineering*, 2nd edn. Springer, New York. <https://doi.org/10.1007/978-1-4614-3523-5>
- [4] Grinchuk P, Tretyak MS, Chuprasov VV (2020) Thermal protection material on the base of silicon-carbide ceramics. *IOP Conf Ser Mater Sci Eng* 709:044112. <https://doi.org/10.1088/1757-899X/709/4/044112>
- [5] Zhang JF, Zhou XN, Zhi Q, Zhao S, Huang X, Zhang NL (2020) Microstructure and mechanical properties of porous SiC ceramics by carbothermal reduction and subsequent recrystallization sintering. *J Asian Ceram Soc* 8:255–264. <https://doi.org/10.1080/21870764.2020.1728045>
- [6] Sung HW, Lee GW, Yun ES, Kim DJ (2020) Novel process for recrystallized silicon carbide through β - α phase transformation. *Ceram Int* 46:21920–21924. <https://doi.org/10.1016/j.ceramint.2020.04.106>
- [7] Kim YH, Kim YW, Lim KY, Lee SJ (2019) Mechanical and thermal properties of silicon carbide ceramics with yttria–scandia–magnesia. *J Eur Ceram Soc* 39:144–149. <https://doi.org/10.1016/j.jeurceramsoc.2018.09.021>
- [8] Kim GD, Kim YW, Yong SM, Jung WK (2022) Improving specific stiffness of silicon carbide ceramics by adding boron carbide. *J Eur Ceram Soc* 42:6827–6835. <https://doi.org/10.1016/j.jeurceramsoc.2022.08.003>
- [9] Shen Y, Wang X, Xing J (2025) A method for preparing electrically conductive SiC ceramics through hot-press

- p>
sintering followed by nitrogen pressure heat treatment.
- Coatings*
- 15:618.
- <https://doi.org/10.3390/coatings15060618>
- [10] Malik R, Kim YH, Kim YW (2020) Effect of additive content on the mechanical and thermal properties of pressureless liquid-phase sintered SiC. *J Asian Ceram Soc* 8:448–459. <https://doi.org/10.1080/21870764.2020.1749376>
- [11] Oh Y, Heo M, Kim YW, Kim HS (2025) Effect of carbon content on electrical, thermal, and mechanical properties of pressureless sintered SiC ceramics. *J Am Ceram Soc* 108:e20562. <https://doi.org/10.1111/jace.20562>
- [12] Ribeiro S, Genova LA, Ribeiro GC, Oliveira MR, Bressiani AHA (2016) Effect of heating rate on the shrinkage and microstructure of liquid phase sintered SiC ceramics. *Ceram Int* 42:17398–17404. <https://doi.org/10.1016/j.ceramint.2016.08.039>
- [13] Yeom JA, Kim YW, Jung WK, Cheong DI, Kang ES (2023) Pressureless sintering of SiC ceramics with improved specific stiffness. *J Eur Ceram Soc* 43:3941–3949. <https://doi.org/10.1016/j.jeureeram Soc.2023.03.011>
- [14] Wang Y, Tan S, Jiang D (2004) The effect of porous carbon preform and the infiltration process on the properties of reaction-formed SiC. *Carbon* 42:1833–1839. <https://doi.org/10.1016/j.carbon.2004.03.018>
- [15] Zhou Y, Sha W, Liu Y, Lyu Y, Huang Y (2022) Influence of carbon source on microstructural and mechanical properties of high-performance reaction-bonded silicon carbide. *Materials* 15:5250. <https://doi.org/10.3390/ma15155250>
- [16] Li W, Zhang G, Cui C, Bao J, Guo C, Xu C, Zhang W, Zhu W (2022) Structure evolution and properties modification for reaction-bonded silicon carbide. *Materials* 15:8721. <https://doi.org/10.3390/ma15248721>
- [17] Liu R, Chen Y, Chen Z, Hai W, Liu M (2025) Effect of gradation on thermal and mechanical properties of reaction-bonded silicon carbide ceramics. *Int J Appl Ceram Technol* 22:e14964. <https://doi.org/10.1111/ijac.14964>
- [18] Filonov KN, Kurlov VN, Klassen NV, Samoilov VM, Vodovozov AN (2009) New shaped ceramics based on silicon carbide. *Bull Russ Acad Sci Phys* 73:1377–1379. <https://doi.org/10.3103/S1062873809100165>
- [19] Zhao YY, Wang WM, Xia HY, Wang JP (2018) Preparation of siliconized graphite by liquid silicon infiltration of porous carbon materials. *Mater Sci Forum* 922:55–61. <https://doi.org/10.4028/www.scientific.net/MSF.922.55>
- [20] Xue R, Wang J, Chen Y, Zhang Z, Xiao Z, Xia H (2023) A hybrid graphite powder PFC@G for reserving higher carbon content of self-lubrication graphite/SiC composites by RMI. *J Am Ceram Soc* 106:6317–6329. <https://doi.org/10.1111/jace.19241>
- [21] Liu S, Zhang B, Cai L, Wang W, Liang T, Lei M (2024) High-temperature and high-pressure tribological properties of siliconized graphite for water-lubricated thrust bearing application in main coolant pump. *Lubricants* 12:159. <https://doi.org/10.3390/lubricants12050159>
- [22] Gerhardt R (ed) (2011) Properties and Applications of Silicon Carbide. InTech, Rijeka. <https://doi.org/10.5772/615>
- [23] Ershov AE, Shikunov SL, Kurlov VN (2017) Method of calculating the phase composition of SiC–Si–C materials obtained by silicon infiltration of carbon matrices. *Tech Phys* 62:903–910. <https://doi.org/10.1134/S1063784217060093>
- [24] Shikunov SL, Kurlov VN (2017) SiC-based composite materials obtained by siliconizing carbon matrices. *Tech Phys* 62:1869–1876. <https://doi.org/10.1134/S1063784217120222>
- [25] Strohmamm T, Bugelnig K, Breitbarth E, Wilde F, Steffens T, Germann H, Requena G (2019) Semantic segmentation of synchrotron tomography of multiphase Al-Si alloys using

a convolutional neural network with a pixel-wise weighted loss function. *Sci Rep* 9:19611. <https://doi.org/10.1038/s41598-019-56008-7>

- [26] Hirabayashi Y, Iga H, Ogawa H, Tokuta S, Shimada Y, Yamamoto A (2024) Deep learning for three-dimensional segmentation of electron microscopy images of complex ceramic materials. *npj Comput Mater* 10:46. <https://doi.org/10.1038/s41524-024-01226-5>
- [27] Aksoy D, Xin HL, Rupert TJ, Bowman WJ (2025) Human perception-inspired grain segmentation refinement using conditional random fields. *Mater Charact* 230:115694. <https://doi.org/10.1016/j.matchar.2025.115694>
- [28] Zhang K, Chen W, Liu WK, Brinson LC, Apley DW (2025) A framework for supervised and unsupervised segmentation and classification of materials microstructure images. *Acta Mater* 301:121588. <https://doi.org/10.1016/j.actamat.2025.121588>
- [29] Della Mea GB, Ovalle C, Laiarinandrasana L, Decenciere E, Dokládal P (2024) Physics informed self-supervised segmentation of elastic composite materials. *Comput Methods Appl Mech Eng* 432:117355. <https://doi.org/10.1016/j.cma.2024.117355>
- [30] Ragone M, Shahabazian-Yassar R, Mashayek F, Yurkiv V (2023) Deep learning modeling in microscopy imaging: A review of materials science applications. *Prog Mater Sci* 138:101165. <https://doi.org/10.1016/j.pmatsci.2023.101165>
- [31] Baskaran A, Kautz EJ, Chowdhary A, Ma W, Yener B, Lewis DJ (2021) Adoption of image-driven machine learning for microstructure characterization and materials design: a perspective. *JOM* 73:3639–3657. <https://doi.org/10.1007/s11837-021-04805-9>

Statements and Declarations

Funding

The work was supported by the Russian Science Foundation, Project No. 25-29-00394.

Competing Interests

The authors have no relevant financial or non-financial interests to disclose.

Author Contributions

A.V. Kaledin and S.M. Peshcherova: Conceived and designed the project, analysed and interpreted the results and prepared the original draft. **D.O. Stryukov:** Contributed to concept, analysis, and editing. **S.L. Shikunov and Yu.N. Zubareva:** Responsible for the material production and discussions. **V.N. Kurlov and I.S. Petrushin:** Guided and supervised the complete work, analysis, final editing, discussions. All authors read and approved the final manuscript.

Data Availability Statement

The datasets generated during and/or analyzed during the current study are available from the corresponding author on reasonable request.

Ethical Approval

This article does not contain any studies with human participants or animals performed by any of the authors.