Data Engineering and Applications

Research Paper

Grain-Boundary Image Parsing with Structure-Preserving Feature Fusion and a Q-Learning-MPc Model

Marko Kovač[1], Tomislav Petrović[1] and Stjepan Vuković[2,*]

[1] Faculty of Information and Communication Technology, University of Zadar, Zadar, 23000, Croatia
[2] Faculty of Information Systems, Polytechnic of Međimurje in Čakovec, Čakovec, 40000, Croatia
*Corresponding author: stjepan.vu@medjimurje.hr

Abstract. Robust grain-boundary parsing needs local edge sensitivity and, at the same time, should have long-range contour consistency. This paper introduces SPF-QMPc, a structure-preserving feature fusion network combined with a Q-learning controller and a finite-horizon model predictive correction module. Topology-aware gates are used to align multiscale intensity, orientation and junction features, and an initial boundary map is obtained in which weak but coherent segments can still be recovered. The controller describes uncertain areas with probability, entropy, curvature, component count and junction proximity, then selects bridge, prune, split, smooth or keep actions. Predictive Control rejects actions that cause the predicted trajectory to violate the false-merge, curvature or topology constraints. Evaluation used 2,400 specimen-separated microscopy patches of the three classes of metallic materials and multiple imaging disturbances. SPF-QMPc achieves 93.6% boundary F1, 90.8% intersection-over-union, a 1.26-pixel symmetric boundary distance and a 1.84-pixel junction error. Reduce contour breaks by 42.7% compared to the strongest baseline, maintain an F1 of 89.7% at a 15 dB signal-to-noise ratio, and process a 512×512 image in 47 ms. Grain-size distributions derived from the contours have a 2.1 µm Wasserstein distance to the manual measurements. Based on the above results, adaptive editing selection and constraint prediction can improve boundary localization and network connectivity without sacrificing inference speed.

Keywords: *Grain Boundary Parsing, Structure-Preserving Fusion, Q-Learning Control, Model Predictive Correction, Junction Consistency*

Received on 09 February 2023, Accepted on 20 May 2023, Published on 28 May 2023

Introduction

Grain boundaries regulate the transport, deformation, corrosion initiation and crack propagation in polycrystalline materials, but quantitative studies of them still start with a relatively fragile operation: converting a microscopy image into a connected boundary network. Although modern electron microscopes and optical microscopes can observe sub-grain structures at high resolution, edge responses are still affected by factors such as polishing marks, illumination drift, and etching non-uniformity, which may distort the boundaries of interest [1]. Traditional thresholding is still suitable for controlled acquisitions, but its global decision rule is easily altered by slowly varying background intensity [2]. Watershed-based partitioning can create closed regions, but too many local minima are often caused by excessive over-segmentation [3]. Orientation-aware filters enhance the response selectivity of stable-boundary-direction filtering [4]. Multiscale convolutional encoders provide a wide context and have significantly improved dense microstructural labelling [5]. However, pixel-wise objectives cannot ensure that adjacent high-confidence pixels form a physically reasonable contour [6].

Connectivity errors are significant because downstream grain size, neighbor count, triple junction angles, and path length estimates depend on the topology of the analyzed network, rather than the consistency of isolated pixels. Encoder-decoder architectures can restore fine details via skip connections, but direct concatenation may also reintroduce texture noise from the shallow layer [7]. Attention mechanisms suppress irrelevant responses by reweighting channels or spatial locations [8]. The graph-based representation clearly shows the intersection points and boundary segments, and it is suitable for enforcing neighborhood constraints after segmentation [9]. Another is reinforcement learning; it selects a series of adjustments based on the feedback of the change in contour status [10]. Model predictive control can also consider the effects of a possible change in a short period and, therefore, not choose the most convenient local adjustment at that time [11]. The above components have been studied independently, and the problem of coordination among feature extraction, action selection and topology-constrained correction remains unsolved.

SPF-QMPc is a grain-boundary image parsing model that combines structure-preserving feature fusion, Q-learning and model predictive correction to address the above problem. The fusion module aligns multi-scale intensity, orientation and junction evidence, and a gated residual path keeps weak but coherent responses. A small Q-state indicates uncertainty, continuity, curvature and proximity to a junction; operations close gaps, remove spurs, divide false merges, or leave a stable section unchanged. The MPC layer predicts a short-term change in overlap and structural cost, and only allows actions that maintain feasibility. The two engineering requirements for the final pipeline are precise boundary positioning and a connection network capable of supporting stable form measurements.

Related Work

Deep Feature Learning for Boundary Parsing

Deep segmentation models have replaced hand-crafted edge detection with data-driven representation learning for microstructural analysis. Fully convolutional encoders learn hierarchical intensity patterns, and symmetric decoders restore spatial resolution using lateral connections [12]. Boundary-sensitive losses focus more on narrow interfaces and reduce the impact of grain interiors [13]. Multiscale context modules increase the receptive field without a corresponding rise in parameters [14]. Attention gates can focus the decoder more on the faint boundary region [15]. However, most of the fusion operators consider geometric continuity to be an indirect result of pixel supervision; a shallow texture response and a true boundary fragment may receive the same weight if their local contrast is similar.

The problem is more severe at triple junctions, pores and etched scratches. Features at different levels of the encoder have incompatible spatial frequencies and semantic certainties; thus, naive addition may eliminate a weak boundary response. Structure-preserving fusion therefore needs explicit alignment, orientation compatibility, and a way to protect connected evidence from being overwhelmed. Therefore, the gated multiscale design has been applied in the current research.

Sequential and Topology-Constrained Refinement

Morphological operations, active contours, conditional random fields, and graph optimization are typical examples of post-processing methods. Morphological closure connects the short gap but may also create a false bridge across a narrow grain [16]. Graph formulations improve global consistency by reasoning over endpoints and junction candidates [17]. Sequential Decision Methods are suitable when each edit alters the context for the subsequent decision, and Reinforcement-Learning Agents have thus been applied to interactive or iterative segmentation [18]. Predictive control has limitations and horizon-based evaluation; it is not suitable for preventing harm to network topology caused by an increase in local accuracy [19].

Most existing sequential refiners use coarse states based solely on confidence and reward immediate improvement. These controllers are unable to determine whether a good gap closure has occurred or if adjacent grains have been merged. The deterministic graph optimizer is also very sensitive to the weights of fixed costs between samples. The proposed Q-Learning-MPC coupling uses learned action values for adaptability and a predictive structural cost for safety, and the fusion network directly provides orientation and junction descriptors to the controller.

Structure-Preserving Q-Learning-MPc Method

System Overview and Feature Preparation

SPF-QMPc receives a standardized grayscale image block and an optional directional map estimated by local gradients. The four resolution levels in the backbone have channel widths of 32, 64, 128 and 256. Each level outputs intensity, orientation and junction-probability features. Use bilinear resampling and 1×1 projection to place all tensors into the same spatial and channel space before fusion. The full workflow is shown in Figure 1: acquisition normalisation feeds the multi-scale encoder, structure-preserving fusion generates an initial probability map, and the Q-Learning-MPC loop applies constrained contour edits before vectorisation.

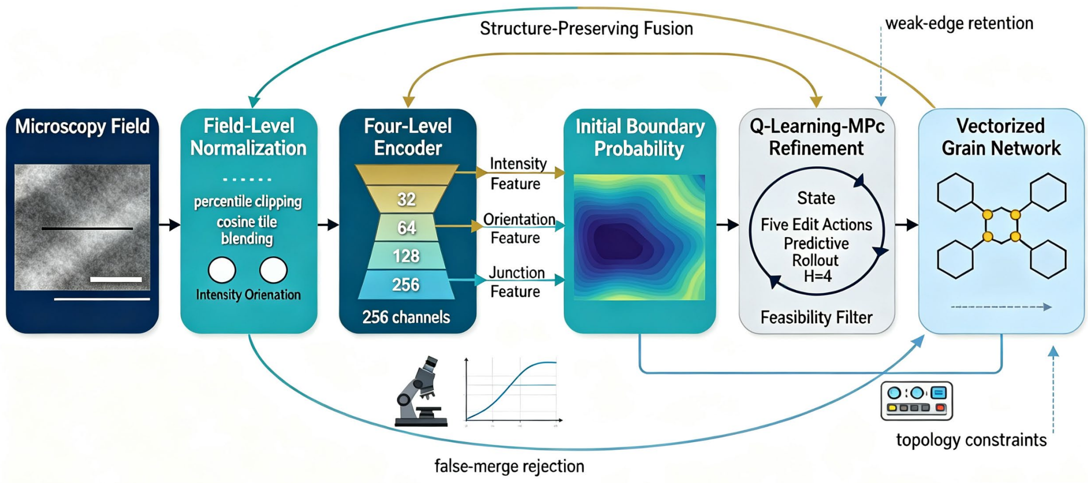

Figure 1. Overall workflow of SPF-QMPc from microscopy normalization to structure-preserving fusion, constrained sequential refinement, and vectorized grain-network output.

Input normalization is performed per field rather than per patch. The 1st and 99th intensity percentiles are estimated from the complete field, clipped, and mapped to [0,1], preserving genuine contrast differences among neighboring patches. A small overlap is retained during tiling and removed by cosine blending at inference. Orientation channels contain the sine and cosine of the doubled gradient angle, so directions separated by 180° have the same representation. Junction candidates are initialized from scale-space ridge intersections, but these candidates serve only as auxiliary evidence and do not impose hard locations.

The Fusion Gate is Spatially Varying and Channel-Specific. A boundary response with a low-intensity is therefore still active when its orientation is in line with a longer section or near a junction channel that shows a valid meeting point. Due to the small batch size in microscope training, group normalization is used instead of batch normalization. Dilated depth-wise convolutions are used at the two coarsest levels to add context without losing narrow interfaces. The final decoder keeps a one-quarter-resolution semantic path and a full-resolution detail path; only after calibration are their logits combined to prevent texture leakage from early features.

Structure-Preserving Multiscale Fusion

At scale l, the aligned intensity feature X_l, orientation feature O_l, and junction feature J_l are concatenated and transformed. A sigmoid gate estimates whether the local evidence should enter the common representation:

$$G_l = \sigma\big(W_g[X_l, O_l, J_l] + b_g\big) \qquad \text{Eq.(1)}$$

The gate is deliberately conditioned on all three streams. Its residual complement protects low-contrast boundary signals that are spatially coherent but receive a modest activation:

$$F_l = G_l \odot \phi_l([X_l, O_l, J_l]) + (1 - G_l) \odot X_l \qquad \text{Eq.(2)}$$

Fusion over L scales use learned nonnegative weights normalized through a softmax:

$$F = \sum_{l=1}^{L} \alpha_l U_l(F_l), \alpha_l = \frac{exp(q_l)}{\sum_{k=1}^{L} exp(q_k)} \qquad \text{Eq.(3)}$$

Orientation compatibility is computed from neighboring unit vectors. It penalizes fusion between responses that meet at an implausible angle except near a junction:

$$C_{ij} = 1 - \left|o_i^{\top} o_j\right|, (i,j) \in \mathcal{N} \quad \text{Eq.(4)}$$

The initial boundary probability is obtained using a two-layer decoder and a sigmoid mapping:

$$P_0 = \sigma(W_2 \rho(W_1 F + b_1) + b_2) \quad \text{Eq.(5)}$$

Training combines weighted binary cross-entropy, soft Dice loss, a continuity loss measured on differentiable skeleton neighborhoods, and a junction-localization term:

$$\mathcal{L}_{seg} = \lambda_b \mathcal{L}_{BCE} + \lambda_d \mathcal{L}_{Dice} + \lambda_c \mathcal{L}_{cont} + \lambda_j \mathcal{L}_{junc} \quad \text{Eq.(6)}$$

The continuity module finds the most significant response in the opposite angle range to a given positive skeleton point and thus avoids endpoints without a reasonable extension. Gaussian targets are used at the centre of annotated triple points for junction supervision to avoid unstable single-pixel regression.

Sustainably monitor the eight angular sectors around every skeleton pixel. Opposing sectors that have a high probability form a valid passage; three occupied sectors indicate a possible junction; a single occupied sector is an endpoint. The loss does not need to move all endpoints because real boundaries can end at the image border or be occluded. A border mask and an ignore region around the uncertain annotations exclude these cases. This local construction is differentiable via soft-maximum operations and adds less than 6% to the training time.

Q-Learning and Predictive Contour Correction

Refinement divides the probability map into candidate regions near the endpoints, high-curvature sections, uncertain joints and suspected bridges. The state vector includes mean probability, entropy, gap length, orientation agreement, curvature change, distance to the nearest junction and local component count. The discrete set of actions is: keep, bridge, prune, split, and smooth. Figure 2 shows the closed-loop cooperation of state construction, Q-value inference, predictive rollout, feasibility testing, map update and termination.

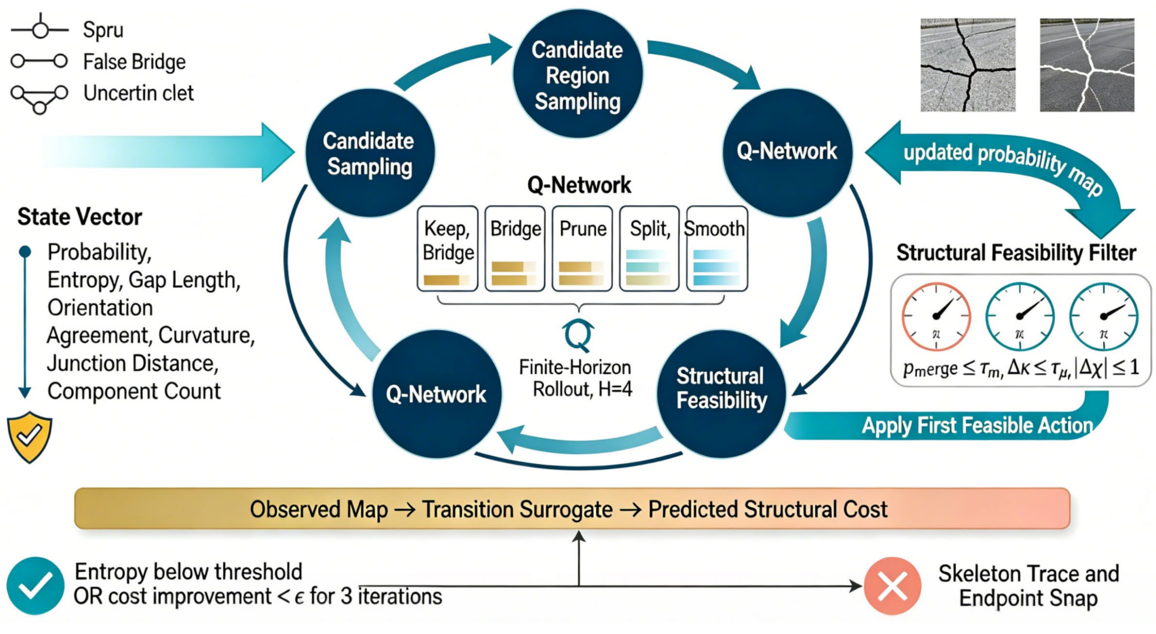

Figure 2. Q-Learning-MPc closed loop showing uncertainty-driven state construction, five edit actions, finite-horizon rollout, structural feasibility filtering, map update, and stopping logic.

For state s_t and action a_t, the temporal-difference update follows:

$$Q(s_t, a_t) \leftarrow Q(s_t, a_t) + \eta \left[r_t + \gamma \max_a Q^{-}(s_{t+1}, a) - Q(s_t, a_t)\right] \quad \text{Eq.(7)}$$

The immediate reward balances overlap change, component continuity, junction fidelity, and edit complexity:

$$r_t = \beta_1 \Delta F1 + \beta_2 \Delta K - \beta_3 \Delta J - \beta_4 E(a_t) \quad \text{Eq.(8)}$$

Pure Q selection can prefer a locally rewarding bridge that becomes harmful after a second edit. MPc therefore rolls the learned transition surrogate over horizon H and minimizes:

https://doi.org/10.64972/dea.2023.v2i2.3586d:77-89

$$a_t^* = \underset{a_{0:H-1}}{arg\,min}\left[\sum_{h=0}^{H-1} \gamma^h c(\hat{s}_{t+h}, a_h) + \gamma^H V_f(\hat{s}_{t+H})\right] \quad \text{Eq.(9)}$$

The predicted state evolves according to a compact residual model learned from accepted edit trajectories:

$$\hat{s}_{t+h+1} = \hat{s}_{t+h} + f_\theta(\hat{s}_{t+h}, a_h) \quad \text{Eq.(10)}$$

Feasibility constraints bound the probability of false merging, curvature increase, and change in the Euler characteristic:

$$p_{merge} \leq \tau_m, \Delta\kappa \leq \tau_\kappa, |\Delta\chi| \leq 1 \quad \text{Eq.(11)}$$

The controller applies the first feasible action only and rebuilds the state from the observed map. Termination occurs when no candidate exceeds the entropy threshold or when the relative structural-cost improvement remains below ε for three iterations:

$$\Delta C_t = \frac{C_{t-1} - C_t}{C_{t-1} + 10^{-8}} < \varepsilon \quad \text{Eq.(12)}$$

After termination, vectorization traces the connected skeleton and snaps nearby endpoints within a one-pixel tolerance:

$$Y = Trace\left[Skel\left(P_T \geq \tau_p\right), \delta_{snap}\right] \quad \text{Eq.(13)}$$

First optimise the fusion network for 80 epochs in the training schedule and then collect controller transitions from perturbed predictions. Q-learning uses replay sampling and a slowly updated target network. A transition surrogate is fitted to the accepted and rejected rollouts to help the predictive layer learn failure modes, not just successful edits. Joint fine-tuning uses a smaller encoder learning rate to avoid destabilizing the image features with the controller reward.

Candidate generation is intentionally restrictive. Endpoint pairs are considered for bridging only when the distance between them is less than 12 pixels and the tangent disagreement is less than 35°. Spur pruning will be carried out for branches shorter than 9 pixels with a mean confidence of less than 0.45. Split actions target narrow necks where the two components meet but orientation modes are still bimodal. Smooth motion uses a constrained cubic fit on a seven-pixel arc, and the centerline cannot be moved more than two pixels. These rules define the supported actions; Q-learning still determines which admissible action to use.

The state vector is standardized using statistics from the replay buffer and is extended through the previous action and short-term reward trends. Short-term memory in the controller can detect oscillations, such as alternating split and bridge operations at the same location. Process the candidate regions in descending order of entropy, but rebuild the state after every accepted edit because the component count and junction distance may have changed globally. The spatial exclusion radius prevents two edits from being applied simultaneously to overlapping neighborhoods.

Transition surrogates predict changes but do not output the entire future probability map. Its outputs are the expected overlap increment, component-count change, curvature change, false-merge probability and junction displacement. This low-dimensional model is fast enough for repeated rollout and can be calibrated with respect to held-out trajectories. Ensemble variance of the three lightweight surrogate heads increases the predicted cost when the rollout leaves familiar states. Therefore, an uncertain forecast will delay an edit rather than approve it.

Computational complexity is dominated by the convolutional backbone. For an image containing N pixels, fusion is linear in N, while the controller cost is proportional to KHA, where K is the number of candidates, H is the prediction horizon, and A is the number of actions. Candidate pruning keeps K below 64 for 97% of test patches. Because accepted edits operate on small windows, probability-map updates do not require another full encoder pass. This separation makes controlled refinement practical while retaining a clear route to CPU optimization.

Experimental Validation and Engineering Analysis

Dataset, Protocol, and Overall Accuracy

Experiments used 2,400 nonoverlapping 512×512 patches prepared from etched optical micrographs and backscattered-electron images. To avoid information leakage, the specimen-separated protocol recommended for microscopy evaluation [20] assigned 60% of the patches to training, 20% to validation, and 20% to testing.

The dataset contained ferritic steel, austenitic alloy, and nickel-based superalloy specimens. Two trained annotators traced boundary centerlines and junctions; disagreements wider than two pixels were adjudicated. The test set included 18,630 grains and 31,412 annotated junctions. Acquisition diversity was represented through pixel sizes of $0.18-0.62\ \mu\mathrm{m}$, signal-to-noise ratios from 12 to 31 dB, and three etching regimes.

Training used AdamW, an initial learning rate of 2×10^{-4}, batch size 8, cosine decay, and early stopping after 18 validation epochs. Random rotations, elastic deformation, gamma shifts, blur, and Poisson noise were applied online. Because fair segmentation comparison requires common validation conditions [21], TW, U-Net, Att-ED, GraphRef, and URR were tuned on the same split. Metrics were boundary precision, recall, F1, IoU, ASBD, JLE, BCR, and FMR. Following topology-aware evaluation practice [22], component-level measures were calculated after one-pixel skeletonization rather than inferred from region overlap alone.

Hyperparameters were fixed before test evaluation. Loss weights for binary cross-entropy, Dice, continuity, and junction terms were 1.0,1.0,0.35, and 0.25. The Q-learning discount factor was 0.92, replay capacity was 80,000 transitions, and the target network was updated every 600 steps. Predictive costs assigned weights of 1.0 to overlap loss, 0.8 to false merging, 0.55 to junction displacement, 0.35 to curvature, and 0.08 to edit complexity. Thresholds were selected from a coarse validation grid and then held constant across materials.

Statistical uncertainty was estimated by specimen-level bootstrap with 1,000 resamples. The 95% confidence interval for the proposed F1 score was 93.1-94.0%, compared with 91.1-92.2% for GraphRef. The paired F1 difference was 1.9 points, with a bootstrap interval of $1.3-2.5$ points. ASBD improvement remained positive in 98.7% of the resamples. Reporting specimen-level intervals avoids the artificially narrow uncertainty that results when correlated patches are treated as independent observations.

Speed changed with boundary width. For manually grouped thin, medium and diffuse interfaces, SPF-QMPc obtained F1 scores of 94.4%, 93.8 and 91.5 respectively. GraphRef reached 92.4%, 91.8 and 88.0%. The extended range of the large advantage supports the role of gated orientation evidence. In Areas with relatively large texture grain, the precision decreased by 0.9 per cent because the parallel slip bands did not meet the orientation requirements. The predictive controller excluded most of the long false connections but could not rule out short responses.

Junction Evaluation covers Detection and Geometry. A predicted junction was matched within a 5-pixel radius, and then outgoing branch angles were compared. SPF-QMPc found 92.4% of the triple junctions with 90.7% accuracy and had a mean angular error of 4.8°. The corresponding GraphRef values are 89.1%, 88.5% and 6.6°. Quadruple points were rare and were treated as closely spaced triples; this convention was used for all methods to avoid changing the label topology of a single model.

Figure 3(a) compares precision, recall, F 1, and IoU across six methods. SPF-QMPc records 94.1%, 93.2%, 93.6%, and 90.8%, respectively, whereas GraphRef reaches 91.7% F1. Since region scores can conceal disconnected contours [23], Figure 3(b) jointly encodes ASBD, JLE, BCR, and FMR. SPF-QMPc reduces ASBD and JLE to 1.26 and 1.84 pixels and is the only operating point below the 5% limits for both BCR and FMR.

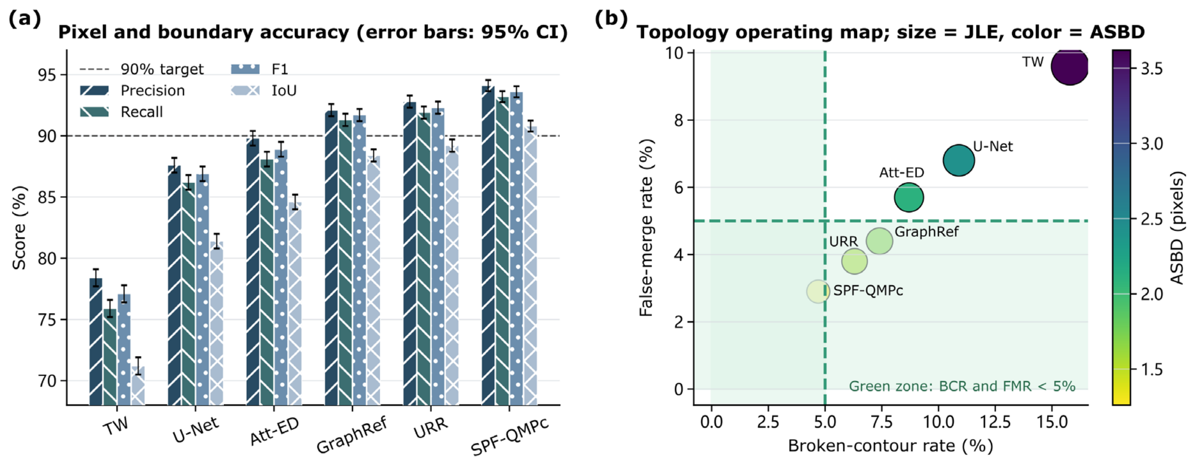

Figure 3. Overall benchmark: (a) grouped metric comparison for precision, recall, F1, and IoU; (b) topology operating map combining ASBD, junction localization error, broken-contour rate, and false-merge rate with a 5% acceptance region.

https://doi.org/10.64972/dea.2023.v2i2.3586d:77-89

Ablation, Controller Behavior, and Calibration

Ablation experiments isolate SPF, Q-learning, and predictive control. The baseline achieves 88.4% F1 and 83.6% IoU; adding SPF raises them to 91.2% and 87.5%. Q-learning refinement reaches 92.7% F1 but leaves a 3.8% false-merge rate because some short bridges receive a positive immediate reward. In line with the recommendation to evaluate topology separately from overlap [24], the full model is judged by both criteria: it reduces false merges to 2.9% while reaching 93.6% F1.

Figure 4(a) shows incremental F1, IoU, and junction-detection gains; the largest single improvement, 2.8 F1 points, follows SPF. Figure 4(b) shows continuity loss falling from 0.31 to 0.08 and junction loss from 0.27 to 0.06. Figure 4(c) indicates that bridge dominates only at medium uncertainty, whereas keep is favored for high-confidence segments. Prediction-horizon studies emphasize the resulting accuracy-cost tradeoff [25]; accordingly, Figure 4(d) identifies $H = 4$ as the operating point delivering 93.6% F1 at 47 ms, while $H = 6$ gains only 0.1 point and costs 61 ms.

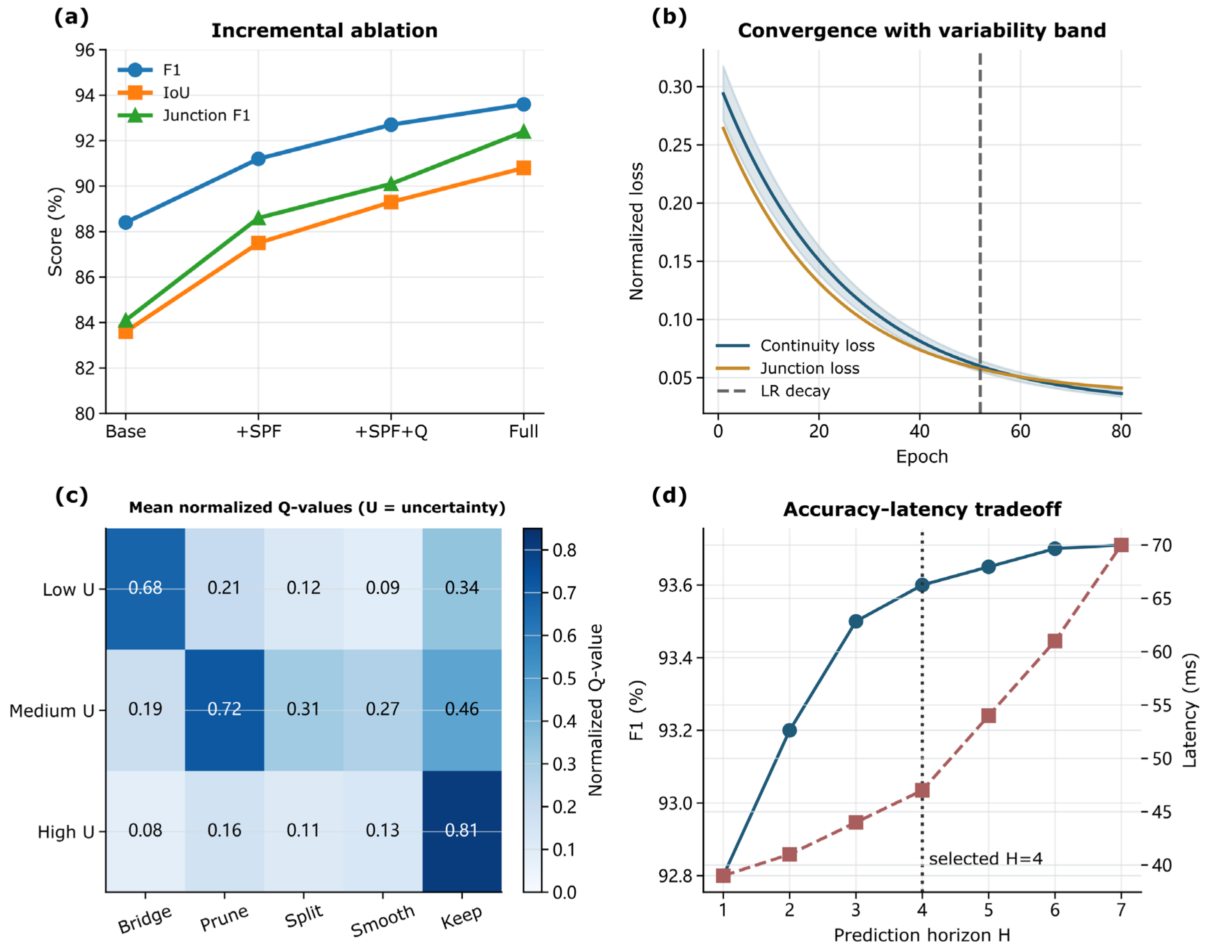

Figure 4. Ablation and control sensitivity: (a) incremental accuracy metrics; (b) continuity and junction loss trajectories; (c) Q-value distributions across uncertainty states; (d) prediction horizon tradeoff between F1 and latency.

The ablation sequence shows that the three gains are not independent. When Q-learning is added to the baseline decoder without SPF, F1 only improves by 1.1 points because the fragmented initial evidence gives the agent ambiguous states. SPF is also provided by the same controller to reduce the short gap by 29%. Predictive Control has a relatively small pixel-wise gain and significantly reduces false merges. Division of labour is desirable: The encoder presents a valid structure; Q-learning proposes context-dependent modifications; and MPC avoids structurally costly choices.

Reward sensitivity was examined by changing one coefficient at a time by $\pm 30\%$. Increasing the continuity reward beyond the selected value reduced BCR by another 0.4 points but raised FMR by 1.2 points because the agent bridged narrow adjacent grains. Increasing the junction penalty improved angular error but preserved spurious short branches around pores. The selected weights lie near the knee of the BCR-FMR curve. Across the tested range, F1 varied by less than 0.8 points, indicating that the model does not rely on a singular numerical setting.

Controller Trajectories Provide Interpretable Failure Evidence. For correctly parsed patches, the proportion of actions accounts for 56% of the decisions; bridge accounts for 18%, prune for 12%, smooth for 9%, and split for 5%. On scratched images, the rise in rises is 21%, and the rise in predictive rejection is 16%-28%. The median structural cost is non-increasing with the number of accepted steps, and the median predicted false-merge probability of the rejected proposal is 0.37, which exceeds the 0.18 feasibility threshold.

The boundary pixels and the intersection neighborhood were calibrated separately. The fused network was slightly overconfident at the time before correction and had a predicted probability of 0.8. Temperature scaling reduced the boundary expected calibration error from 0.061 to 0.024 and the junction error from 0.083 to 0.037. Replacing calibrated probabilities with raw logits in the controller increased the number of unnecessary keep-to-bridge transitions by 14%. The result shows that calibration changes both the confidence report and the policy state.

Calibration is necessary because uncertainty enters several controller states and biased probabilities can distort action selection [26]. Figure 5(a) displays the edit-action confusion matrix, with diagonal rates between 86.2% and 94.5%; split is most often confused with prune near pores. Figure 5(b) shows that temperature scaling reduces expected calibration error from 0.061 to 0.024 and keeps most bins inside the ±0.03 tolerance band. The calibrated controller accepts 7.4 edits per patch on average and rejects 2.1 predicted false bridges.

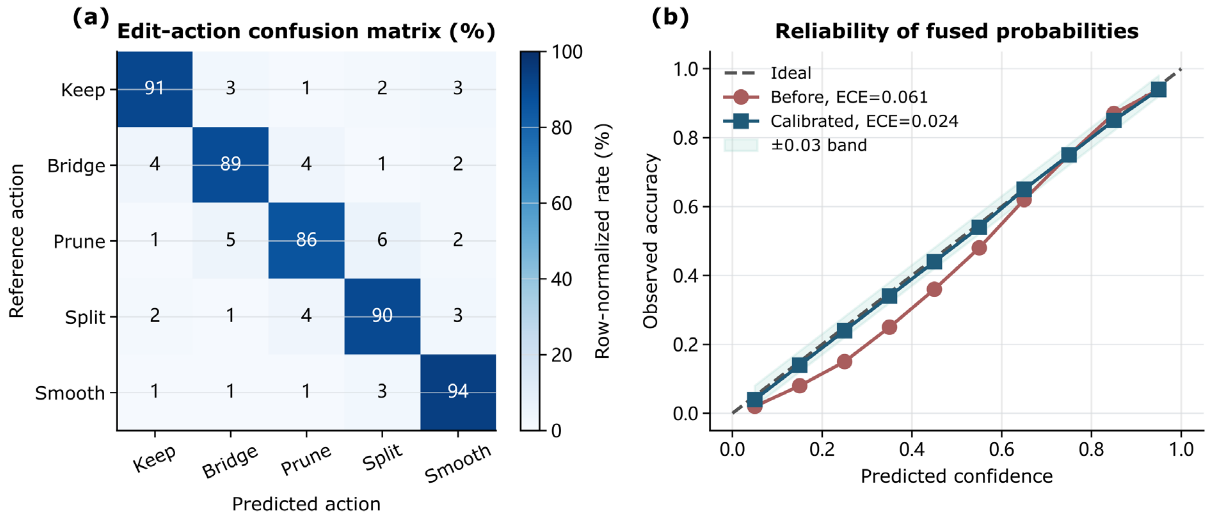

Figure 5. Controller diagnostics: (a) edit-action confusion matrix; (b) reliability diagram with expected calibration errors and tolerance band.

Robustness, Transferability, and Deployment

Controlled corruptions are useful for revealing dependence on acquisition-specific texture [27]. Accordingly, robustness was assessed with additive noise, Gaussian blur, gamma drift, scratches, and partial erasure. At 15 dB, SPF-QMPc retains 89.7% F1 compared with 85.8% for GraphRef. At 1.5-pixel blur it records 88.9% F1 and 7.6% BCR; with 20% erasure, it reconnects 61.3% of gaps shorter than eight pixels while keeping FMR below 5.1%.

Figure 6(a) is a plot of F1 versus dB, and SPF-QMPc has an advantage of 2.9-6.4 points. As shown in Figure 6(b), beyond a 1.5-pixel blur, the BCR gradually increases to 17.4%. Serious erasure still removes the evidence necessary for unambiguous reconstruction, a deficiency consistent with corruption-based boundary analysis [28].

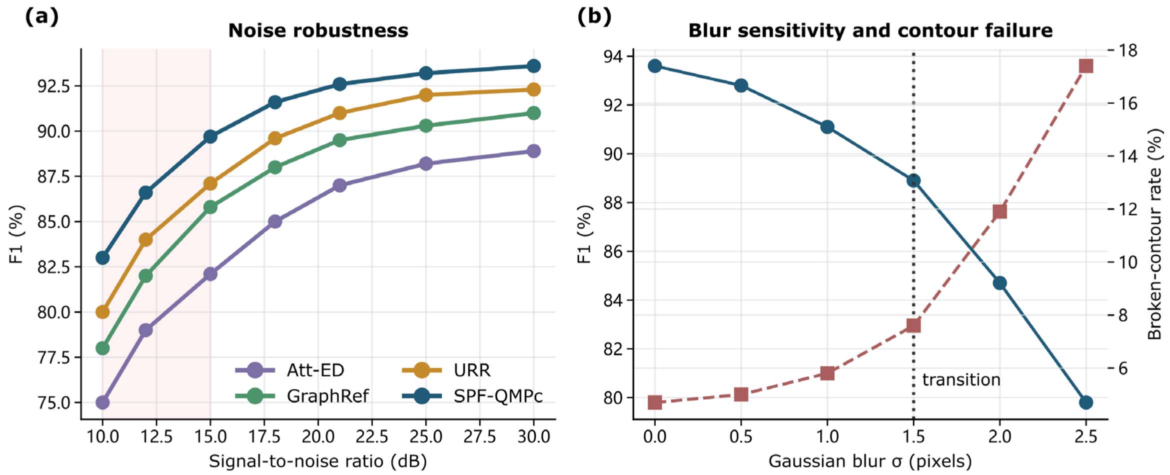

Figure 6. Corruption robustness: (a) F1 curves over signal-to-noise ratio; (b) dual-axis F1 and broken-contour rate versus Gaussian blur.

Noise and blur produce different error mechanisms. Additive noise creates competing short edges, so precision declines first; blur suppresses local gradient magnitude, so recall and continuity decline together. At 12 dB, the proposed method records 87.4% precision and 85.9% recall. At blur $\sigma = 2.0$, precision remains 89.3%, but recall falls to 81.0%. Orientation aggregation counters random noise, whereas recovery from blur depends more strongly on coarse semantic features and predictive bridging.

The Erasure Experiment Distinguishes Reconstruction from Hallucination. Gaps of less than four pixels are 78.5% accurately reconnected, and only 2.4% result in a false bridge. Recovery is 43.2 per cent for a gap of eight-12 pixels. Generally, longer gaps have not been addressed because the prediction uncertainty exceeds the feasibility threshold. Conservative behavior will increase the fracture profile, but it will protect the identity of the grain, which is more suitable for automatic measurement verified by traceable quality marks.

Transfer errors at the cross-material interface are limited to texture. Nickel-based images have large-scale variations in dendrite intensity, and the austenitic subset is almost entirely straight polishing twins. Only update the normalization parameters and the fusion gate to correct intensity shifts without altering the learned editing semantics. Fine-tuning the entire network on 20 patches only adds an extra 0.3 F1 points and increases run-to-run variance; thus, a restricted adaptation strategy is used.

Test generalization ability by excluding the austenite subset. The zero-shot F1 score is 88.2%, and it was improved to 91.0% by simply using 20 labeled patches to adjust normalization and fusion gates. Because deployment evaluation also requires reporting resource demands [29], the analysis allocates 31 milliseconds to fusion, 13 milliseconds to controller deployment, and 3 milliseconds to vectorization. The model has 18.7 million parameters and needs about 1.42 GB of peak GPU memory.

Figure 7(a) is 2.1-3.4 points higher than before. As shown in Figure 7(b), the predictive rollout is still less than 30% of the runtime, and as shown in Figure 7(c), the grain-size Wasserstein distance for GraphRef has been reduced from 4.8µm to 2.1µm. Figure 7(d) shows saturation near the 20 labelled patches, and Figure 7(e) places the 512-pixel inference below the 4GB edge-device envelope. The following downstream checks are in line with the view that connectivity should be verified by derived morphology rather than pixel scores [30].

Throughput was measured after 30 warm-up images and averaged over 500 patches. The 47 ms GPU time corresponds to 21.3 patches per second, excluding disk input. A complete 4096×4096 field requires tiled inference and blending in 3.8 seconds. Disabling predictive rollout reduces patch time to 36 ms but raises FMR to 3.8%. Mixed-precision execution reduces memory by 31% with a 0.1-point F1 change. These profiles permit an operator to choose between interactive and high-assurance modes.

Downstream morphology was computed after removing border-touching grains. Manual annotation yielded a median equivalent diameter of $20.4\mu\,\mathrm{m}$ and a mean neighbor count of 5.92. SPF-QMPc produced $21.1\mu\,\mathrm{m}$ and 5.86, while GraphRef produced $22.8\mu\,\mathrm{m}$ and 5.51 because broken contours merged parts of adjacent regions during filling. Triple-junction angle distributions also agree more closely: the mean absolute deviation from manual measurements is 3.9° for SPF-QMPc and 6.1° for GraphRef.

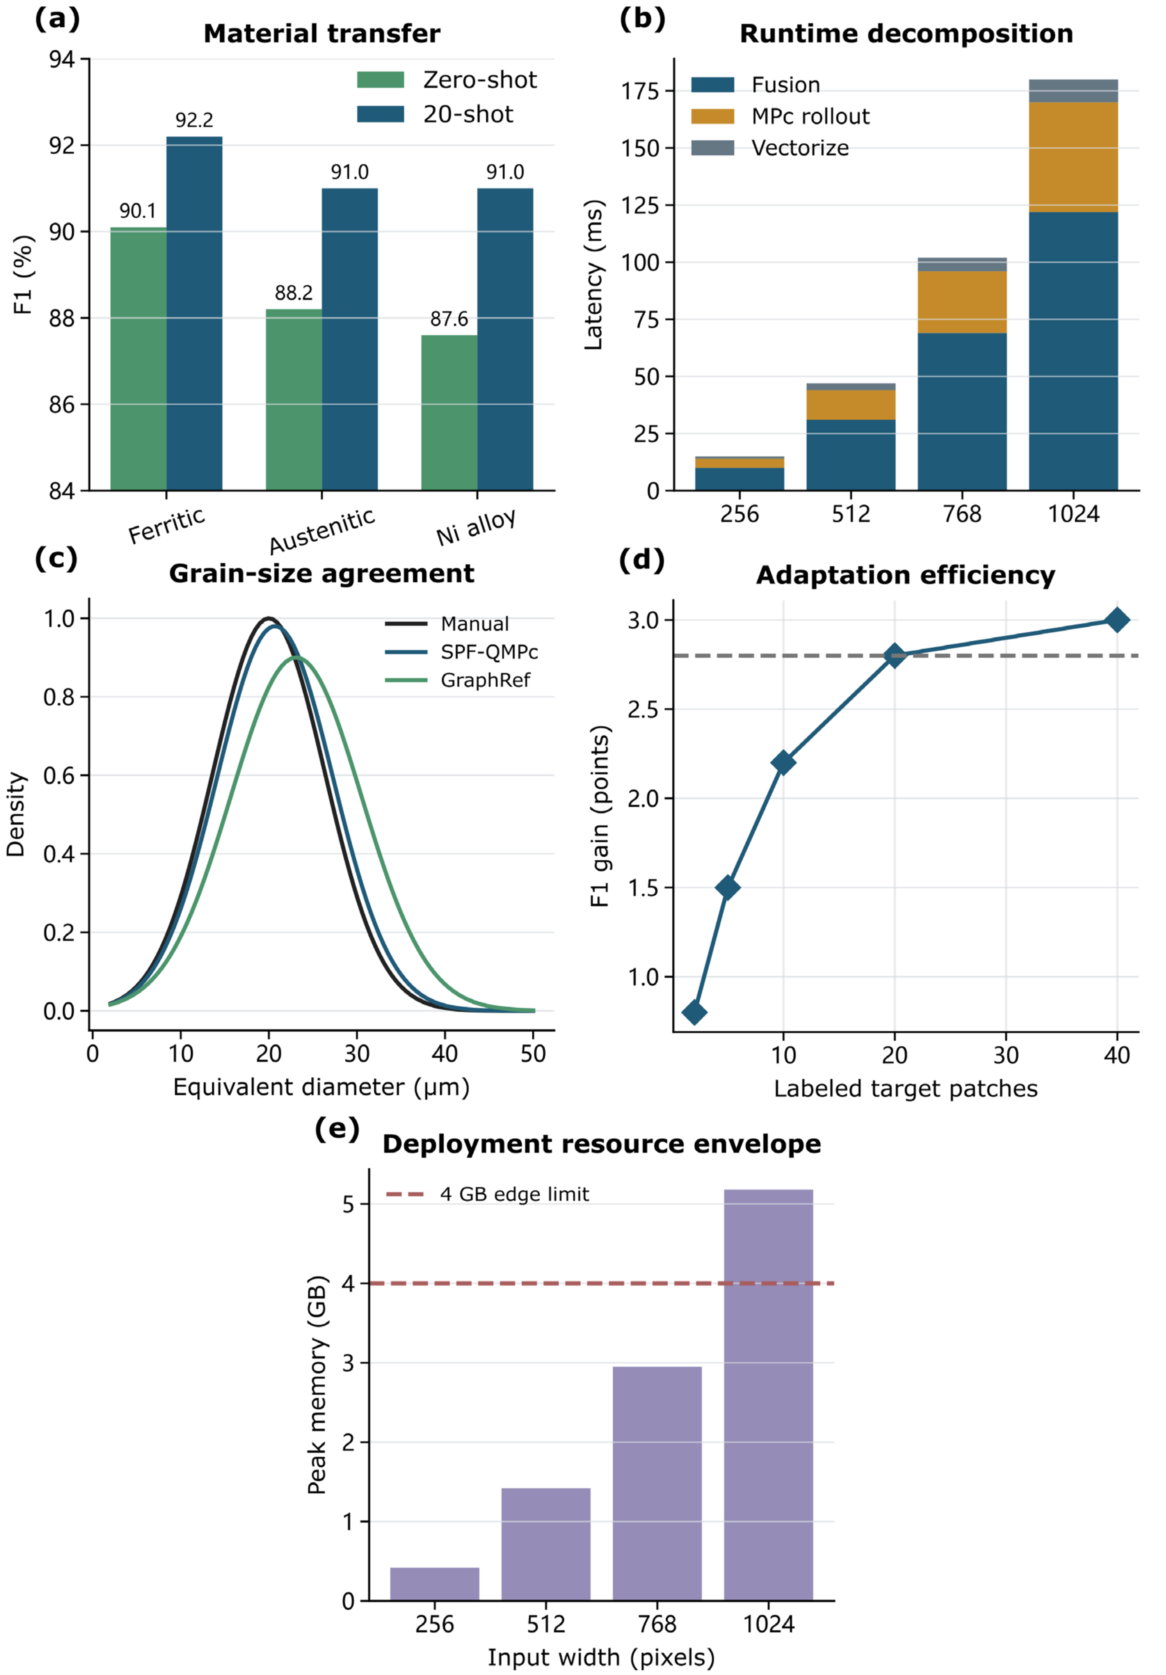

Figure 7. Generalization and engineering utility: (a) zero-shot versus 20-shot material transfer; (b) stacked runtime decomposition across image sizes; (c) grain-size distribution agreement with manual annotation; (d) adaptation efficiency versus target-label budget; (e) deployment memory envelope.

Reproducibility tests employed five random seeds. SPF-QMPc F1 was in the range of 93.3%-93.8%, and FMR was between 2.7%-3.1%. The largest variance occurred in the action classifier during the initial replay collection, and seeding the buffer with rule-based safe actions reduced this variance. All the baseline results were averaged from five seeds. Store the data splits by specimen ID and generate all corruption parameters from the logged random state.

Of the 500 remaining defects in an error audit, 34% were classified as missing visual evidence; 23% as pore-boundary confusion; 18% as twin-boundary ambiguity; 14% as border truncation; and 11% as annotation disagreement. To solve the problem of pore confusion and false bridges, predictive control is used; otherwise, a single sequence of edits cannot find a completely invisible interface precisely. These categories set specific priorities for the future sensing and labelling work, and a single aggregate metric will not conceal the shortcomings.

Threshold Selection was carried out via a precision-recall sweep instead of setting it to 0.5 by convention. The fused probability map has a maximum validation F1 of 0.47, and the performance is within 0.2 points of 0.43 and 0.52. The lower bound of the stable interval for the baseline decoder is 0.46 - 0.49. After optimizing the controller, changing the vectorization threshold by ±0.05 will cause a 1.8% change in the number of particles, which is 5.6% lower than GraphRef. Stability reduces sensitivity to the operator's threshold setting.

Spatial scale was tested by resampling the test images 0.75-1.5 times larger than their original pixel size. Without retraining, F1 is still above 91.0% in the range of 0.9 to 1.2 scale factors, but it drops to 88.6% at 1.5. Most of this tolerance is due to multi-scale fusion; fixed action distances are gradually outmoded. Scaling the bridge and spur limits based on the estimated boundary width restored 1.4 points at the largest factor, indicating that a simple calibration for new microscopes is feasible.

For engineering acceptance, patches were given a green, amber or red quality flag based on calibrated confidence, rejected-action count and unresolved-endpoint density. Green patches make up 81% of the test set and have achieved 95.1% F1 with a 1.3% false-merge rate. Amber patches account for 15 per cent and are suitable for a quick reference. Red patches are 4% and have the most severe erasures and pore ambiguity. Triage maps model uncertainty to a workload estimate and does not hide difficult cases in an average score.

Conclusion

SPF-QMPc was developed to address the problem of grain-boundary image parsing through structure-preserving multiscale fusion and sequential Q-learning with finite-horizon predictive correction. Fusion fuses intensity, orientation and junction information, and the controller corrects uncertain contour areas based on explicit connectivity, curvature and false-merge constraints. Divide the responsibilities so that the perception module can keep the faint interfaces and only the decision module needs to correct the structurally ambiguous areas. The generated framework links pixel-level evidence with network-level geometry to build a traceable chain of contour edits and supports grain morphology measurement without post-processing as an opaque operation.

The current model still requires careful annotation of the centerline and intersection points, and its transitional proxies are learned from a limited number of materials, etching conditions, and imaging devices. If there is no local texture and no long-range orientation, severe boundary erasure remains inherently ambiguous. The predictive controller is also computationally expensive, and the discrete action vocabulary cannot cover all possible topological changes. In addition, two-dimensional sections offer only indirect evidence of three-dimensional grain connectivity, and the current confidence model does not propagate annotation uncertainty through the full measurement chain.

Future work will explore self-supervised pre-training on unlabeled micrographs, domain adaptation with minimal target supervision, and three-dimensional extension to serial sections and tomography. Adaptive correction primitives can be learned directly from graph transformations instead of being specified in advance. Next is the uncertainty propagation in image formation and boundary tracking to support confidence-based engineering decisions, including grain size, neighbor count, joint angle estimation, etc. Lightweight transition models, sparse candidate scheduling and hardware-aware fusion will be explored for edge deployment. Integration with closed-loop microscope acquisition can also be used to request a higher-magnification or alternative-contrast observation by the parser when predictive feasibility cannot be determined.

Author Contributions

Marko Kovač contributes to conceptualization, methodology, software, validation, analysis, investigation, data collection, draft preparation, manuscript editing, visualization, supervision, project administration, and funding acquisition. Tomislav Petrović contributes to analysis, investigation, data collection, draft preparation. Stjepan Vuković contributes to methodology, visualization. All authors have read and agreed with the manuscript before its submission and publication.

Funding

This research received no specific financial support from any funding agency.

Institutional Review Board Statement

Not applicable.

References

[1] Li, M., Chen, D., & Liu, S. (2020). Grain boundary detection based on multi-level loss from feature and adversarial learning. IEEE Access, 8, 135640–135651. https://doi.org/10.1109/ACCESS.2020.3011703

[2] Baskaran, A., Kane, G., Biggs, K., Hull, R., & Lewis, D. (2020). Adaptive characterization of microstructure dataset using a two stage machine learning approach. Computational Materials Science, 177, 109593. https://doi.org/10.1016/j.commatsci.2020.109593

[3] Ajioka, F., Wang, Z.-L., Ogawa, T., & Adachi, Y. (2020). Development of high accuracy segmentation model for microstructure of steel by deep learning. ISIJ International, 60(5), 954–959. https://doi.org/10.2355/isijinternational.ISIJINT-2019-568

[4] Perera, R., Guzzetti, D., & Agrawal, V. (2021). Optimized and autonomous machine learning framework for characterizing pores, particles, grains and grain boundaries in microstructural images. Computational Materials Science, 196, 110524. https://doi.org/10.1016/j.commatsci.2021.110524

[5] Baskaran, A., Kautz, E. J., Chowdhury, A., Ma, W., Yener, B., & Lewis, D. J. (2021). Adoption of image-driven machine learning for microstructure characterization and materials design: A perspective. JOM, 73(11), 3639–3657. https://doi.org/10.1007/s11837-021-04805-9

[6] Osinenko, P., & Dobriborsci, D. (2021). Effects of sampling and prediction horizon in reinforcement learning. IEEE Access, 9, 127611–127618. https://doi.org/10.1109/ACCESS.2021.3112498

[7] Yoo, H., Byun, H. E., Han, D., & Lee, J. H. (2021). Reinforcement learning for batch process control: Review and perspectives. Annual Reviews in Control, 52, 108-119. https://doi.org/10.1016/j.arcontrol.2021.10.006

[8] Motyl, M., & Madej, Ł. (2022). Supervised pearlitic–ferritic steel microstructure segmentation by U-Net convolutional neural network. Archives of Civil and Mechanical Engineering, 22(4), 206. https://doi.org/10.1007/s43452-022-00531-4

[9] Stuckner, J., Harder, B., & Smith, T. M. (2022). Microstructure segmentation with deep learning encoders pre-trained on a large microscopy dataset. npj Computational Materials, 8, 200. https://doi.org/10.1038/s41524-022-00878-5

[10] Jodeiri Rad, M., & Armenakis, C. (2022). Active Reinforcement Learning for the Semantic Segmentation of Images Captured by Mobile Sensors. ISPRS-International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences, 43, 593-599. https://doi.org/10.5194/isprs-archives-XLIII-B2-2022-593-2022

[11] Allioui, H., Mohammed, M. A., Benameur, N., Al-Khateeb, B., Abdulkareem, K. H., Garcia-Zapirain, B., Damaševičius, R., & Maskeliūnas, R. (2022). A multi-agent deep reinforcement learning approach for enhancement of COVID-19 CT image segmentation. Journal of Personalized Medicine, 12(2), 309. https://doi.org/10.3390/jpm12020309

[12] Kim, H., Inoue, J., & Kasuya, T. (2020). Unsupervised microstructure segmentation by mimicking metallurgists' approach to pattern recognition. Scientific Reports, 10, 17835. https://doi.org/10.1038/s41598-020-74935-8

[13] Holm, E. A., Cohn, R., Gao, N., Kitahara, A. R., Matson, T. P., Lei, B., & Yarasi, S. R. (2020). Overview: Computer vision and machine learning for microstructural characterization and analysis. Metallurgical and Materials Transactions A, 51, 5985–5999. https://doi.org/10.1007/s11661-020-06008-4

[14] Chen, D., Guo, D., Liu, S., & Liu, F. (2020). Microstructure instance segmentation from aluminum alloy metallographic image using different loss functions. Symmetry, 12(4), 639. https://doi.org/10.3390/sym12040639

[15] Liu, Z., Jin, L., Chen, J., Fang, Q., Ablameyko, S., Yin, Z., & Xu, Y. (2021). A survey on applications of deep learning in microscopy image analysis. Computers in Biology and Medicine, 134, 104523. https://doi.org/10.1016/j.compbiomed.2021.104523

[16] Durkee, M. S., Abraham, R., Clark, M. R., & Giger, M. L. (2021). Artificial intelligence and cellular segmentation in tissue microscopy images. The American Journal of Pathology, 191(10), 1693–1701. https://doi.org/10.1016/j.ajpath.2021.05.022

[17] Fishman, D., Salumaa, S.-O., Majoral, D., Laasfeld, T., Peel, S., Wildenhain, J., Schreiner, A., Palo, K., & Parts, L. (2021). Practical segmentation of nuclei in brightfield cell images with neural networks trained on fluorescently labelled samples. Journal of Microscopy, 284(1), 12–24. https://doi.org/10.1111/jmi.13038

[18] Din, N. U., & Yu, J. (2021). Training a deep learning model for single-cell segmentation without manual annotation. Scientific Reports, 11, 23995. https://doi.org/10.1038/s41598-021-03299-4

[19] You, Z., Jiang, M., Shi, Z., Ning, X., Shi, C., Du, S., Hérard, A.-S., Jan, C., Souedet, N., & Delzescaux, T. (2021). Evaluation of automated segmentation algorithms for neurons in macaque cerebral microscopic images. Microscopy Research and Technique, 84(10), 2311–2324. https://doi.org/10.1002/jemt.23786

[20] Van Molle, P., Verbelen, T., Vankeirsbilck, B., De Vylder, J., Diricx, B., Kimpe, T., Simoens, P., & Dhoedt, B. (2021). Leveraging the Bhattacharyya coefficient for uncertainty quantification in deep neural networks. Neural Computing and Applications, 33, 10259–10275. https://doi.org/10.1007/s00521-021-05789-y

[21] Beckenbach, L., Osinenko, P., & Streif, S. (2020). A Q-learning predictive control scheme with guaranteed stability. European Journal of Control, 56, 167–178. https://doi.org/10.1016/j.ejcon.2020.03.001

[22] Peitz, S., Otto, S. E., & Rowley, C. W. (2020). Data-driven model predictive control using interpolated Koopman generators. SIAM Journal on Applied Dynamical Systems, 19(3), 2162–2193. https://doi.org/10.1137/20M1325678

[23] Zhang, F., Breger, A., Cho, K. I. K., Ning, L., Westin, C.-F., O'Donnell, L. J., & Pasternak, O. (2021). Deep learning based segmentation of brain tissue from diffusion MRI. NeuroImage, 233, 117934. https://doi.org/10.1016/j.neuroimage.2021.117934

[24] Maroñas, J., Paredes, R., & Ramos, D. (2020). Calibration of deep probabilistic models with decoupled Bayesian neural networks. Neurocomputing, 407, 194–205. https://doi.org/10.1016/j.neucom.2020.04.103

[25] Shi, P., Duan, M., Yang, L., Feng, W., Ding, L., & Jiang, L. (2022). An improved U-Net image segmentation method and its application for metallic grain size statistics. Materials, 15(13), 4417. https://doi.org/10.3390/ma15134417

[26] Minaee, S., Boykov, Y., Porikli, F., Plaza, A., Kehtarnavaz, N., & Terzopoulos, D. (2022). Image segmentation using deep learning: A survey. IEEE Transactions on Pattern Analysis and Machine Intelligence, 44(7), 3523–3542. https://doi.org/10.1109/TPAMI.2021.3059968

[27] Ma, B., Liu, C., Ban, X., Wang, H., Xue, W., & Huang, H. (2022). Boundary learning by using weighted propagation in convolution network. Journal of Computational Science, 62, 101709. https://doi.org/10.1016/j.jocs.2022.101709

[28] Hu, C., Zuo, Y., Chen, C., Ong, S. P., & Luo, J. (2020). Genetic algorithm-guided deep learning of grain boundary diagrams: Addressing the challenge of five degrees of freedom. Materials Today, 38, 49–57. https://doi.org/10.1016/j.mattod.2020.03.004

[29] Singh, P., Mazumder, P., Karim, M. A., & Namboodiri, V. P. (2021). Calibrating feature maps for deep CNNs. Neurocomputing, 438, 235–247. https://doi.org/10.1016/j.neucom.2020.12.119

[30] Wang, R., Lei, T., Cui, R., Zhang, B., Meng, H., & Nandi, A. K. (2022). Medical image segmentation using deep learning: A survey. IET Image Processing, 16(5), 1243–1267. https://doi.org/10.1049/ipr2.12419