

EM3M: An Electron Micrograph Dataset for Microstructural Segmentation and Generation

Nan Wang*, Zhiyi Xia*, Yiming Li*, Siyuan Zhang, Shi Tang,
Zoe Fan, Xi Fang, Haoyi Tao, Guolin Ke, and Yanhui Hong**

DP Technology, Beijing, China
{wangnan01, hongyanhui}@dp.tech

Abstract. Quantitative microstructural characterization is fundamental to materials science, and electron micrographs (EMs) provide indispensable high-resolution insights. However, progress in deep learning-based analysis of EMs has been hampered by the scarcity of large-scale, expert-annotated public datasets. To address this issue, we introduce EM3M, a large-scale and multimodal dataset for instance-level understanding of EMs. EM3M comprises 5,091 high-quality EMs, approximately **3 million** instance segmentation annotations, and image-level textual descriptions with disentangled attributes. The dataset is constructed through a rigorous multi-stage curation and validation pipeline, with comprehensive statistical analyses to ensure reliability and reproducibility. Building upon these curated image-text pairs, we further provide a text-to-image diffusion model that serves as a controllable data augmentation engine, demonstrating that synthetic augmentation consistently improves downstream segmentation performance. To establish a systematic benchmark, we evaluate representative instance segmentation methods on EM3M. Our results reveal that conventional detection-based and query-based methods struggle with the extreme instance densities and textural complexities inherent in EMs. We additionally provide an optimized flow-based baseline to facilitate fair comparison and future research. EM3M¹, the generative engine², and an online demo³ are publicly available to support future research in automated materials analysis.

Keywords: Dataset and Benchmark · Electron Microscopy · Instance Segmentation · Microstructural Characterization

1 Introduction

Quantitative microstructural characterization underpins advances in materials design, diagnostics, and fabrication. Electron microscopy techniques such as

* Equal contribution.

** Corresponding author.

¹ Dataset: <https://huggingface.co/datasets/UniParser/EM3M>

² Generation: <https://huggingface.co/UniParser/EM3M-Gen>

³ Segmentation demo: <https://www.bohrium.com/apps/uni-aims>




Fig. 1: Visualization of the EM3M. **Left:** Diverse raw EMs from different materials and morphologies. **Right:** A WO_3 sample exhibiting dense particulate features, along with its instance segmentation masks containing **2,445** annotated particles, and structured textual annotations describing material type, morphology, diameter distribution, and other attributes.

scanning electron microscopy (SEM) and transmission electron microscopy (TEM) are widely employed to capture microstructural details, enabling precise measurement of grain size, density, and distribution [7, 10, 19]. However, large-scale EM analysis remains challenging due to the need for domain expertise, costly licensing, and limited transferability of existing tools, such as ImageJ [17] and MountainsSEM.

Despite the promise of deep learning for automated EM analysis [16], most models are typically developed for specific tasks and struggle to generalize across the vast heterogeneity of real-world EM data. Previous studies have applied learning-based methods to nanomaterial classification [3], microstructural segmentation [5, 42, 49, 52, 58], and micrograph generation [4, 46, 54]. However, these efforts are often confined to particular imaging modalities, operational modes, or material types, and therefore lack the universality required to address the broader spectrum of EMs. Additionally, costly EM acquisition, expert annotation, and IP/privacy restrictions [39, 43, 50] limit the availability of large, diverse public datasets, which in turn hinders the development of Segment Anything-style foundation models (SAM) [31] for materials science.

The multimodal domain presents additional challenges. A comprehensive understanding of material microstructures requires correlating visual patterns with precise scientific interpretations—yet existing datasets fall short. For instance, MMSci [33] has harvested vast volumes of micrographs and associated text from scientific literature, but its figures are post-processed for publication and its captions are often noisy and incomplete. Datasets such as MicroVQA [8] and MicroBench [37] include electron microscopy modalities but focus on biological and medical domains, leaving a gap for materials-oriented multimodal resources.

To address these issues, we present EM3M, a large-scale and multimodal EM dataset designed for instance-level materials understanding. Through ex-

Table 1: Summary of EM datasets. *Estimated upper bound of EMs in MMSci. [†]Images were extracted from literature figures rather than raw experimental data. Abbreviations: Mixed–synthetic and experimental data; Ins–instance segmentation; Cls–classification; # Ins.–number of annotated instances.

Dataset	Source	# Images	Label	# Ins.	Scene	Text	Available
EMPS [58]	Literature [†]	465	Ins	11596	Diverse	×	✓
Aversa <i>et al.</i> [3]	Experiment	22000	Cls	-	Diverse	×	✓
Stuckner <i>et al.</i> [52]	Mixed	110,861	Cls	-	Diverse	×	×
Okunev <i>et al.</i> [42]	Experiment	26	Ins	5852	Single	×	✓
Bin <i>et al.</i> [49]	Experiment	237	Ins	605	3 scenes	×	×
Bals <i>et al.</i> [5]	Experiment	93	Ins	-	3 scenes	×	×
MMSci [33]	Literature	$\ll$ 100,000*	-	-	Diverse	✓	✓
EM3M (Ours)	Experiment	5,091	Ins	2,985,660	Diverse	✓	✓

tensive academic and industrial collaborations, we curated a collection of 5,091 high-quality electron micrographs, covering a wide range of material types, fabrication processes, and imaging modalities across SEM and TEM instruments, which capture representative microstructural morphologies. Crucially, via a dedicated human-in-the-loop pipeline, we have annotated approximately **3 million** instance masks in total, achieving unprecedented annotation density with certain images exceeding 2,000 instances. Fig. 1 illustrates the diversity of our dataset and the richness of its annotations. Unlike MMSci, EM3M is built on raw experimental micrographs and uses an expert-defined protocol to produce structured, disentangled attribute descriptions for each image. To further enhance the dataset’s utility, we leverage the curated image-text pairs to train a text-to-image diffusion model that functions as a controllable augmentation engine, enabling scalable synthesis of realistic micrographs for the annotation pipeline. We also note that the decoupled attributes provide the secondary benefit of creating novel samples through recombination. While our primary focus is dataset construction, the structured descriptions provide a natural foundation for future materials-oriented vision–language modeling.

Finally, we benchmark various instance segmentation methods on EM3M, revealing that models originally designed for natural images [6, 9, 12, 15, 21, 22, 29, 32, 55] suffer accuracy degradation and become computationally and memory-intensive under extreme instance density. In contrast, micrograph-specific approaches that rely on flow- or embedding-based intermediate representations remain both accurate and resource efficient [14, 20, 25, 48, 51]; we group these into a “bottom-up” family. Building on these insights, we revisit and improve the flow-based framework and train it on a mixture of synthetic and real micrographs, establishing a strong baseline called UniAIMS (**U**nified **AI**-assisted **M**icroscopy **S**egmentation). We summarize the contribution of this work as follows:

- We propose the EM3M, a large-scale and multimodal EM dataset for microstructural segmentation and generation. It comprises 5,091 high-quality

EMs, about 3 million instance segmentation labels, and image-level attribute-disentangled textual descriptions.

- We release a text-to-image diffusion model trained on the EM3M to serve as a powerful data augmentation tool. Quantitative evaluation metrics and downstream segmentation tasks validate the effectiveness of the generated data.
- We benchmark various instance segmentation methods on EM3M, and improve the flow-based framework on both synthetic and real micrographs, establishing a strong EM-tailored baseline that serves as a foundation for future microstructural research.

2 Related Work

2.1 Material EM Dataset

Tab.1 summarizes representative EM datasets used in prior work. The EMPS dataset [58] contains 465 EMs with pixel-level segmentation masks, featuring only particle-type microstructures with low instance density and simple scenes. Although Aversa *et al.* [3] introduced the first large-scale annotated SEM dataset comprising 22,000 images, and Stuckner *et al.* [52] used over 100,000 EMs for pre-training, their datasets provide only image-level labels without pixel-level annotations. Additionally, several studies [5, 42, 49] still rely on datasets too small for comprehensive instance-level learning. Other recent efforts draw on EM images extracted from scientific literature, such as MMSci [33] and OmniSci [53]. This dataset contains post-processed images tailored for publication figures, exhibiting reduced realism and lacking raw structural complexity. Moreover, given the scarcity of manually annotated EM datasets, a number of studies [34, 36, 40, 54] have explored synthetic data generation to alleviate the annotation burden and improve model performance. However, they typically focus on simplistic, single-material scenarios with limited structural diversity and fail to capture the high-density, heterogeneous microstructural complexity found in real-world EMs.

2.2 Instance Segmentation

Instance segmentation aims to delineate and distinguish individual object instances within an image. Detection-based two-stage methods, such as Mask R-CNN [22] and its variants [9, 12], are intuitive but often degrade in crowded or highly overlapping scenes. In contrast, one-stage methods [6, 29, 38, 55] prioritize efficiency by directly predicting instance masks in a single pass, often at the cost of mask quality. More recently, query-based transformers [15, 21, 32] have achieved state-of-the-art performance by reformulating the task as a direct set prediction problem. However, their reliance on complex bipartite matching for label assignment and a fixed number of learned object queries make them both computationally intensive and difficult to optimize, especially in scenarios with extremely high instance counts. To address this scalability challenge, a fundamentally different bottom-up, proxy-based paradigm has proven highly effective,

particularly for dense imagery in biological microscopy [14, 20, 28, 48, 51]. These methods predict intermediate pixel-level proxies, such as the geometric polygons in StarDist [48] and CPP-Net [14], or vector fields like the distance maps in HoverNet [20, 25] and the gradient flows in Cellpose [44, 51], which are then used to group pixels into individual instances. This bottom-up design bypasses the memory-intensive proposal matching overhead during training, and hence scales better in extreme instance-density scenarios. Lastly, interactive methods like SAM [31] excel at user-guided segmentation, but their reliance on prompts makes them less effective for fully automatic segmentation in dense scenes.

3 EM3M Dataset

3.1 Data Engine

Data collection and processing. We constructed EM3M by aggregating real-world EMs sourced from three primary channels: (1) high-quality, original EMs from academic collaborations, (2) automated web crawling, and (3) 300 representative images selected from Aversa *et al.* [3] and newly annotated for instance segmentation. The consolidated data underwent a rigorous curation pipeline, including perceptual hash deduplication and a quality filtering process. This process discarded low-resolution images (<200px) and employed a GPT-based [1] assessment to remove severely blurred images where any embedded text was illegible. After curation, the final dataset comprises approximately 84% images from academic collaborations, 10% from web crawling, and 6% from the public dataset.

Specifically, for training our text-to-image diffusion model, we performed an additional data sanitation step to remove textual artifacts, which could otherwise cause hallucinations and degrade image quality. To eliminate embedded metadata that introduces such artifacts, we trained a YOLO11 [29] model to detect and remove these regions, ensuring the training data contains only pure microstructural content.

Structured description generation. To enable fine-grained and structured descriptions of EMs, we designed a descriptive framework based on nine largely disentangled dimensions that characterize the key aspects of a micrograph. These dimensions include metadata (material and microscopy type), visual appearance attributes (color configuration), and microstructural characteristics (morphology, density, distribution, layering in Z-axis, surface texture, and diameter distribution, as shown in Fig. 1). Subsequently, we developed a semi-automated annotation pipeline leveraging Gemini [18] and GPT. Each model independently extracts nine image attributes into a unified JSON schema, and then cross-validates the peer’s output to identify inconsistent predictions. Detailed prompts are provided in the supplementary material Sec. 10. Subsequently, all outputs undergo a mandatory human review. Discrepancies identified between the models are resolved by materials science experts to ensure high-fidelity, attribute-decoupled annotations.

Table 2: Inter-annotator agreement and evaluation metrics across stages.

Metric	Inter-annotator agreement			Evaluation across stages (vs. GT)		
	A vs B	B vs C	A vs C	Pseudo Label	Peer Review	Final Validation
KA	0.823	0.810	0.816	0.695	0.801	0.818
PQ	0.795	0.779	0.783	0.553	0.738	0.784

Segmentation annotation and quality assessment. We adopted an iterative, human-in-the-loop workflow: domain experts first annotated a high-quality seed set to train an initial segmentation model. To maximize the model’s performance and the quality of its predictions in each iteration, we leveraged a diffusion model to augment the training data with diverse, realistic samples. This enhanced model then generated pseudo-labels for the remaining data, which were subsequently refined by trained annotators. To ensure label quality, all annotations underwent a rigorous, multi-stage quality control pipeline, including peer review and final validation by a materials science expert. Additional details are available in the supplementary material Sec. 11. To rigorously validate annotation reliability and the efficacy of this pipeline, we established a stratified subset ($N = 300$). This subset was independently re-annotated by three annotators to serve as a benchmark. We measured inter-annotator agreement and assessed the reduction of model bias using pixel-level Krippendorff’s Alpha (KA) [2] and Panoptic Quality (PQ). As shown in Tab 2, the results indicate consistency among experts, and the steady metric improvements observed across pipeline stages confirm that our protocol effectively mitigates model bias and guarantee expert-level annotation quality.

3.2 Dataset Analysis

Through the multi-stage pipeline described above, we constructed the EM3M dataset, comprising 5,091 EMs with 2,985,660 instance-level segmentation annotations. The dataset is randomly divided into 4,128 training images and 963 test images.

Comparison with other EM datasets. As summarized in Tab. 1, previous EM datasets either focus on classification tasks with a large number of images but no instance-level labels, or provide instance segmentation annotations for a very limited number of images and scenes. While MMSci offers a large-scale collection with textual data, it is sourced from literature, and it lacks precise instance-level annotations. Our EM3M dataset uniquely addresses these limitations by providing the first large-scale collection of nearly 3 million instance segmentation labels derived directly from raw experimental images. Furthermore, it is the only dataset that combines high-density instance labels with structured textual descriptions, providing a crucial resource for developing and evaluating future models for both complex segmentation and multimodal understanding in materials science.




Fig. 2: Statistical distribution of three representative attributes from EM3M, illustrating its diversity: (a) Material categories, (b) Morphologies, and (c) Particle densities.




Fig. 3: Statistical distribution of instances in EM3M (log scale). (a) Distribution of instance counts per image. (b) Distribution of instance areas.

Universality of EM3M. We visualized three key attributes—material categories, microstructure morphology, and density—by applying a coarse categorical scheme to accommodate the dataset’s broad coverage and aggregating the corresponding frequencies, as shown in Fig. 2. Despite a naturally long-tail distribution across individual categories, the combined coverage encompasses inorganic and organic materials, a wide spectrum of particle shapes (from highly irregular aggregates to well-defined geometries), and varying degrees of feature density (from isolated particles to heavily overlapped clusters). This breadth of diversity underpins the dataset’s value for developing robust computer-vision models across challenging EM scenarios. Additional attribute visualizations are available in the supplementary material Sec. 12.

Instance-level complexity and scale variation. As shown in Fig. 3, instance-level statistics reveal highly heterogeneous and challenging scene complexity. The number of annotated microstructures per image spans over four orders of magnitude: while many micrographs contain a few hundred microstructures, a non-negligible portion exhibits extreme crowding (over 1,000 instances), posing significant challenges for segmentation algorithms. Similarly, the distribution of instance areas—from ten-pixel segments to over 10^5 pixels on a logarithmic

scale—indicates substantial scale variation. Additional statistical visualizations are provided in the supplementary material Sec. 13.

4 Methodology

4.1 Benchmarking Baseline

To identify a suitable framework for microstructural analysis, we benchmark representative instance segmentation paradigms on EM3M. Our empirical results (Sec. 5.1) demonstrate that flow-based approaches, exemplified by Cellpose [44, 51], provide a more stable and competitive reference than detection- and query-based methods under extreme instance density. This advantage stems from their pixel-wise flow representation, which directly aggregates local gradients into object instances and is less constrained by proposal or query mechanisms in crowded scenarios. A summary of the principle is in the supplementary material Sec. 8.

Guided by this empirical observation, we construct a strong baseline, termed UniAIMS, based on the Cellpose flow formulation. To better match the characteristics of EM3M, we introduce several practical adaptations: we remove the built-in size model to accommodate substantial scale variation, adopt scale-robust data augmentation, enable distributed training with large-resolution inputs, and modernize the backbone for stronger feature representation. These adjustments are implementation-level refinements tailored to large-scale EM data, with further details provided in the supplementary material Sec. 9.

4.2 SDXL Finetune

Diffusion models [24, 27, 45, 59, 60], especially Stable Diffusion and its advanced variant SDXL, have recently achieved remarkable progress in text-guided image generation. Methods such as Textual Inversion [41] and DreamBooth [47] further adapt these pretrained models to specific subjects. Moreover, LoRA-based methods [26] significantly reduce the computational cost of fine-tuning, enabling efficient content and style customization.

In this work, we adopt the LoRA training scheme and introduce several learnable tokens to represent different categories within EM3M. Each token is initialized with a rare word and optimized during LoRA fine-tuning, allowing it to better capture category-specific representations. This training process is guided by a carefully prepared dataset of paired images and structured descriptions. On the image side, micrographs from EM3M are processed to remove any regions containing text, ensuring a clean, high-quality dataset free from textual artifacts (see Sec. 3.1 for details). On the text side, each image is paired with a structured prompt constructed from a set of quasi-orthogonal attributes (see Fig. 1). The prompt follows the consistent template shown below:

```
<microscopy_type> of <material>: <morphology>. <texture>. <density>. <distribution>.
<layering>. <diameters>. <color>.
```

Table 3: Baseline performance comparison of various instance segmentation methods on the sparse subset and the full dataset.

Class	Method	Backbone	Params	Sparse		All	
				$mAP@0.5$	$PQ@0.5$	$mAP@0.5$	$PQ@0.5$
Two-stage	Mask R-CNN [22]	ResNeXt101	101M	0.620	0.610	-	-
	Cascade R-CNN [9]	ResNeXt101	135M	0.624	0.622	-	-
	HTC [12]	ResNeXt101	137M	0.497	0.538	-	-
One-stage	YOLACT [6]	ResNeXt101	93M	0.538	0.537	-	-
	YOLO11-seg [29]	ResNeXt101	116M	0.619	0.605	-	-
	SOLOv2 [55]	ResNeXt101	113M	0.339	0.454	-	-
Query-based	FastInst [21]	ResNeXt101	93M	0.332	0.466	-	-
	Mask2Former [15]	Swin-Base	107M	0.296	0.301	-	-
	MaskDINO [32]	Swin-Base	115M	0.395	0.515	-	-
Bottom-up	Stardist [48]	ViT-Base	93M	0.670	0.559	0.668	0.582
	CPP-Net [14]	ViT-Base	96M	0.663	0.602	0.647	0.528
	HoverNet [20]	ViT-Base	104M	0.620	0.549	0.594	0.554
	CellViT [25]	ViT-Base	131M	0.608	0.625	0.672	0.600
	Cellpose [44]	ViT-Base	93M	0.650	0.692	0.711	0.691
UniAIMS (Ours)		ViT-Base	93M	0.818	0.769	0.800	0.704

The structured format is crucial for achieving fine-grained control over the synthesis process. By leveraging LoRA-based fine-tuning of both adapter weights and word embeddings, our model can generate plausible variations when sampling attributes based on their observed correlations, resulting in microstructures that adhere to realistic material properties. Furthermore, this attribute-level control allows for the exploratory synthesis of novel images by randomly recombining attributes, which can generate diverse, sometimes physically implausible, samples potentially useful for enriching the data distribution. Visualization examples of controllable synthesis, as well as the impact of synthetic data on instance segmentation performance, are provided in Sec. 5.

5 Experiment

5.1 Instance Segmentation

Comparison methods. For our benchmark, we select representative instance segmentation methods, including four major architectural paradigms. We evaluate classic detection-based approaches, including two-stage models that follow a “detect-then-segment” strategy (Mask R-CNN [22], Cascade R-CNN [9], HTC [12]) and faster one-stage models (YOLACT [6], YOLO11-seg [29], SOLOv2 [55]). We also include query-based architectures (Mask2Former [15], MaskDINO [32], FastInst [21]), which treat segmentation as a direct set prediction task. Finally, to specifically address the challenges of dense imagery, we assess a group of bottom-up, proxy-based methods, including StarDist [48] (polygon-based) and its variant CPPNet [14], HoverNet [20] and CellViT [25] (embedding-based), and Cellpose (flow-based) [51], which predict intermediate pixel-level representations before reconstructing instances.




Fig. 4: Ground truth and predictions of different methods on the sparse subset.

Implementation details. To ensure rigor and reproducibility, all benchmarked methods are implemented using their official public codebases, primarily through the MMDetection [13] and Detectron2 [56] framework, with specialized repositories integrated as needed. Following CellViT [25], we equip all bottom-up methods with a SAM-pretrained ViT-B backbone to ensure fair comparison. Correspondingly, other methods are benchmarked using backbones of comparable scale, such as ResNeXt101 [57] and Swin-B [35]. Given that detection-based and query-based architectures struggle with the extreme instance densities found in our dataset, we benchmarked them exclusively on a sparse subset containing images with fewer than 100 instances. Conversely, all bottom-up methods, including UniAIMS, were evaluated on the full dataset to demonstrate their scalability. A detailed justification for this evaluation strategy, along with other implementation specifics, is provided in the supplementary material Sec. 14.

Evaluation metrics. We evaluate instance segmentation performance using mean Average Precision (mAP) for mask detection accuracy and Panoptic Quality (PQ) [30] at 0.5 IoU for combined segmentation and recognition quality. Following prior work in bioimage analysis [20, 48, 51], we report AP as $AP = \frac{|TP|}{|TP| + |FP| + |FN|}$, where predictions with IoU above a threshold T are counted as true positives. Mean AP is obtained by averaging AP across images. This definition differs from the COCO-style AP, but has become a standard practice in microscopy instance segmentation.

Main results. Tab. 3 presents a comprehensive quantitative comparison of representative instance segmentation paradigms on EM3M, and the qualitative examples in Fig. 4 corroborate these trends while revealing characteristic failure modes. Additional visualizations and analysis are provided in the supplementary material Sec. 15.

Query-based models generally underperform on our benchmark: for instance, Mask DINO attains a mAP of only 0.395 on the sparse split, and frequently pro-




Fig. 5: Visualization of our prediction result for images with dense instances.

duces fragmented or non-cohesive masks that fail to separate adjacent instances in cluttered regions (Fig. 4, column 3). For microstructures with complex surface topographies and high-frequency texture, the direct set-prediction strategy in these architectures may be less effective at recovering continuous, fine-grained boundaries that are important for morphological analysis.

Detection-based methods (e.g., Mask R-CNN, Cascade R-CNN) remain competitive, which indicates their utility for simpler scenes with a lower density of objects. However, bottom-up approaches show a consistent advantage overall. Notably, Cellpose achieves a strong balance between instance recall and mask fidelity, as evidenced by the relatively small gap between its mAP and PQ scores. This motivated our decision to adopt its flow-based mechanism as the foundation for our model. A key vulnerability of bottom-up pipelines, however, is their iterative pixel-grouping post-processing, which can accumulate errors and occasionally lead to the complete omission of very large instances (Fig. 4, row 2).

Building upon the flow-based paradigm, our proposed UniAIMS mitigates these limitations by improving multi-scale processing and feature representation. The resulting model achieves a state-of-the-art mAP of 0.818 on the sparse subset and maintains robust performance (0.800 mAP) on the full, more challenging dataset. The quantitative and qualitative evidence strongly indicates




Fig. 6: Evaluation of the generated dataset quality. (a) Comparison of models pre-trained on SA1B versus our 30K generated images (Gen30K). (b) t-SNE visualization showing the substantial feature space overlap between real (orange) and generated (blue) data, indicating high domain fidelity.

that proxy-based designs, particularly when augmented for scale robustness, are exceptionally well-suited for dense, texture-rich electron micrographs. The superiority of our method is most pronounced in highly dense scenarios. As visualized in Fig. 5, UniAIMS delivers highly accurate segmentation for densely packed and irregularly shaped instances, with predictions that closely align with the ground truth, confirming the effectiveness of our approach in addressing these unique challenges.

Varying amounts of data. Tab. 4 demonstrates a clear positive correlation between the amount of real training data and model performance, with mAP steadily increasing from 0.767 (1K images) to 0.800 (4K images). A similar trend is observed for generated data, where performance also increases with data volume but shows diminishing gains at larger scales. Training solely on 30K generated images (0.791 mAP) yields performance approaching that of 3K real images (0.795 mAP), underscoring the high quality of our synthesized data. The best results, 0.809 mAP and 0.707 PQ, are obtained when real and generated data are jointly used, confirming that our synthetic data serves as a powerful and effective tool for data augmentation.

Pretraining with generated data. To further assess the utility of the synthesized dataset, we pretrain our model on 30K generated images and finetune it with varying amounts of real data, as illustrated in Fig. 6a. Compared with the SA1B-pretrained [31] counterpart, the model pretrained on our generated data consistently achieves superior performance across all real-data scales, demonstrating the strong transferability and pretraining potential of our synthetic dataset.

Cross-Dataset Evaluation and Benchmark Saturation. We evaluated UniAIMS on external public datasets (EMPS [58], Okunev *et al.* [42]). As shown in Tab. 5, UniAIMS achieves strong zero-shot performance, significantly outperforming existing baselines even without domain-specific training. Moreover, with

Table 4: Performance metrics for models trained on varying amounts of real data, generated data, and their combination.

Modality	Amount	$mAP@0.5$	$PQ@0.5$
Real	1 K	0.767	0.682
	2 K	0.786	0.696
	3 K	0.795	0.700
	4 K	0.800	0.704
Gen	5 K	0.760	0.669
	10 K	0.771	0.676
	20 K	0.785	0.686
	30 K	0.791	0.690
Gen+Real	30+4 K	0.809	0.707

Table 5: UniAIMS Performance on external datasets.

Dataset	Reported	Zero-shot	Finetune
EMPS [58]	0.92	0.89	0.92
Okunev [42]	0.75	0.81	0.84

Table 6: Generative metrics between the real datasets and the generative dataset.

Metric	Real	Base	Ours
FID(↓)	28.06	213.16	34.03
Style(↑)	0.985	0.669	0.958

fine-tuning, performance on these datasets saturates at near-perfect levels. This performance ceiling suggests that prior datasets may be insufficiently challenging for state-of-the-art models due to limited structural complexity. In contrast, EM3M poses substantially greater difficulty, underscoring its necessity as a robust benchmark to drive future advances in EM segmentation.

5.2 Synthesized Data

This section presents qualitative and quantitative analysis of the generative model (finetuned SDXL), focusing on aspects such as fidelity and distribution consistency.

Quantitative analysis. We quantitatively assess our generative model in Tab. 6, where *Real Data* denotes the processed real micrographs without text regions, and *SDXL (base)* refers to the synthetic dataset produced by the vanilla SDXL model without fine-tuning. To measure distributional consistency, we adopt the Fréchet Inception Distance (FID) [23]. Since the standard FID relies on an ImageNet-trained feature extractor that is suboptimal for EMs, we first establish a domain-specific reference by computing the FID between real images and their spatially augmented counterparts (rotation, scaling, and flipping). This reference value of 28.06 reflects the inherent sensitivity of FID within our domain and serves as a baseline for meaningful comparison. Our fine-tuned model achieves an FID of 34.03, substantially improving over the SDXL baseline (213.16) and approaching the domain reference, indicating strong generative fidelity. We further evaluate style similarity by extracting the [CLS] tokens of both real and generated images using a pre-trained DINO-ViT [11] and computing their cosine similarity. As shown in Tab. 6, the high Style score (0.958) confirms that our synthesized images closely align with the stylistic distribution of real micrographs. This, combined with the analysis from segmentation




Fig. 7: Visual samples of generated data. (a) demonstrates distributional consistency between real images and text-conditioned generated samples, while (b) illustrates novel out-of-distribution examples produced through attribute recombination.

experiments, further validates the effectiveness of the synthesized data and its generalization capability for downstream tasks.

Qualitative analysis. Fig. 6b illustrates the t-SNE visualization results comparing the synthesized and real datasets, where the highly overlapping distributions confirm their consistency. Furthermore, visual samples in Fig. 7(a) demonstrate that our model can produce high-quality, realistic EM data. Notably, generative EM samples from different categories show obvious distinctions and exhibit visual attributes highly similar to real samples. Meanwhile, Fig. 7(b) shows the capability of generative model to generate novel, out-of-distribution images via attribute recombination. This further demonstrates the model’s ability to learn representations across various categories and attributes. Overall,

both qualitative and quantitative analyses confirm the potential of our synthetic dataset to replace and augment real datasets.

6 Open Problems

Despite the strong performance of UniAIMS and the scale of EM3M, several important challenges remain unsolved. First, extremely dense micrographs containing heavily clustered and stacked instances continue to be difficult even for flow-based methods. Moreover, while AP/PQ@50 provides a useful benchmark, these metrics may be overly lenient in crowded scenarios and do not fully reflect fine-grained boundary quality. For example, although UniAIMS achieves 0.704 PQ@50, its performance drops substantially to 0.340 under the stricter PQ@90 criterion, indicating considerable room for improvement in precise instance delineation. Second, microstructures with fibrous textures, weak contrast, and ambiguous boundaries remain challenging due to the lack of reliable structural cues for instance separation (Fig. 13). Third, substantial intra-image scale variation, where particle sizes may span several orders of magnitude within a single micrograph, continues to challenge current architectures despite scale-aware training strategies. Furthermore, a noticeable gap remains between human annotation consistency and model performance: under the same PQ metric, average human agreement reaches 0.786, whereas the best model attains 0.704. Together, these findings suggest that future research should explore more effective dense-scene representations, boundary-aware learning objectives, scale-adaptive architectures, and evaluation protocols that better capture segmentation quality in highly crowded micrographs.

7 Conclusion

In this study, we introduce EM3M, a large-scale and multimodal electron microscopy dataset for microstructural segmentation and generation in materials science. Building on this foundation, we release a text-to-image diffusion model that serves as an effective data augmentation engine, with quantitative and downstream results validating the utility of the generated samples. Through comprehensive benchmarking, we demonstrate that an improved flow-based framework performs robustly on synthetic and real micrographs, establishing a strong EM-tailored baseline. Together, EM3M, the UniAIMS baseline, and the generative diffusion model form a unified toolkit that we expect will accelerate progress in automated materials analysis and material-oriented vision–language research.

References

1. Achiam, J., Adler, S., Agarwal, S., Ahmad, L., Akkaya, I., Aleman, F.L., Almeida, D., Altenschmidt, J., Altman, S., Anadkat, S., et al.: Gpt-4 technical report. arXiv preprint arXiv:2303.08774 (2023)

2. Amgad, M., Atteya, L.A., Hussein, H., Mohammed, K.H., Hafiz, E., Elsebaie, M.A., Alhusseiny, A.M., AlMoslemany, M.A., Elmatboly, A.M., Pappalardo, P.A., et al.: Nucls: A scalable crowdsourcing approach and dataset for nucleus classification and segmentation in breast cancer. *GigaScience* **11**, giac037 (2022)
3. Aversa, R., Modarres, M.H., Cozzini, S., Ciancio, R., Chiusole, A.: The first annotated set of scanning electron microscopy images for nanoscience. *Scientific data* **5**(1), 1–10 (2018)
4. Bals, J., Eppe, M.: Artificial scanning electron microscopy images created by generative adversarial networks from simulated particle assemblies. *Advanced Intelligent Systems* **5**(7), 2300004 (2023)
5. Bals, J., Eppe, M.: Deep learning for automated size and shape analysis of nanoparticles in scanning electron microscopy. *RSC advances* **13**(5), 2795–2802 (2023)
6. Bolya, D., Zhou, C., Xiao, F., Lee, Y.J.: Yolact: Real-time instance segmentation. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 9157–9166 (2019)
7. Brandon, D., Kaplan, W.D.: *Microstructural characterization of materials*. John Wiley & Sons (2013)
8. Burgess, J., Nirschl, J.J., Bravo-Sánchez, L., Lozano, A., Gupta, S.R., Galaz-Montoya, J.G., Zhang, Y., Su, Y., Bhowmik, D., Coman, Z., et al.: Microvqa: A multimodal reasoning benchmark for microscopy-based scientific research. In: *Proceedings of the Computer Vision and Pattern Recognition Conference*. pp. 19552–19564 (2025)
9. Cai, Z., Vasconcelos, N.: Cascade r-cnn: High quality object detection and instance segmentation. *IEEE transactions on pattern analysis and machine intelligence* **43**(5), 1483–1498 (2019)
10. Callister, W.D., Rethwisch, D.G.: *Fundamentals of materials science and engineering*. John Wiley & Sons (2022)
11. Caron, M., Touvron, H., Misra, I., Jégou, H., Mairal, J., Bojanowski, P., Joulin, A.: Emerging properties in self-supervised vision transformers. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 9650–9660 (2021)
12. Chen, K., Pang, J., Wang, J., Xiong, Y., Li, X., Sun, S., Feng, W., Liu, Z., Shi, J., Ouyang, W., et al.: Hybrid task cascade for instance segmentation. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 4974–4983 (2019)
13. Chen, K., Wang, J., Pang, J., Cao, Y., Xiong, Y., Li, X., Sun, S., Feng, W., Liu, Z., Xu, J., et al.: Mmdetection: Open mmlab detection toolbox and benchmark. *arXiv preprint arXiv:1906.07155* (2019)
14. Chen, S., Ding, C., Liu, M., Cheng, J., Tao, D.: Cpp-net: Context-aware polygon proposal network for nucleus segmentation. *IEEE Transactions on Image Processing* **32**, 980–994 (2023)
15. Cheng, B., Misra, I., Schwing, A.G., Kirillov, A., Girdhar, R.: Masked-attention mask transformer for universal image segmentation. *CVPR* (2022)
16. Choudhary, K., DeCost, B., Chen, C., Jain, A., Tavazza, F., Cohn, R., Park, C.W., Choudhary, A., Agrawal, A., Billinge, S.J., et al.: Recent advances and applications of deep learning methods in materials science. *npj Computational Materials* **8**(1), 59 (2022)
17. Collins, T.J.: Imagej for microscopy. *Biotechniques* **43**(sup1), S25–S30 (2007)
18. Comanici, G., Bieber, E., Schaeckermann, M., Pasupat, I., Sachdeva, N., Dhillon, I., Blistein, M., Ram, O., Zhang, D., Rosen, E., et al.: Gemini 2.5: Pushing the

- frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities. arXiv preprint arXiv:2507.06261 (2025)
19. Goldstein, J.I., Newbury, D.E., Michael, J.R., Ritchie, N.W., Scott, J.H.J., Joy, D.C.: Scanning electron microscopy and X-ray microanalysis. springer (2017)
 20. Graham, S., Vu, Q.D., Raza, S.E.A., Azam, A., Tsang, Y.W., Kwak, J.T., Rajpoot, N.: Hover-net: Simultaneous segmentation and classification of nuclei in multi-tissue histology images. *Medical image analysis* **58**, 101563 (2019)
 21. He, J., Li, P., Geng, Y., Xie, X.: Fastinst: A simple query-based model for real-time instance segmentation. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 23663–23672 (2023)
 22. He, K., Gkioxari, G., Dollár, P., Girshick, R.: Mask r-cnn. In: *Proceedings of the IEEE international conference on computer vision*. pp. 2961–2969 (2017)
 23. Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems* **30** (2017)
 24. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. *Advances in neural information processing systems* **33**, 6840–6851 (2020)
 25. Hörst, F., Rempe, M., Heine, L., Seibold, C., Keyl, J., Baldini, G., Ugurel, S., Siveke, J., Grünwald, B., Egger, J., et al.: Cellvit: Vision transformers for precise cell segmentation and classification. *Medical Image Analysis* **94**, 103143 (2024)
 26. Hu, E.J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., Chen, W., et al.: Lora: Low-rank adaptation of large language models. *ICLR* **1**(2), 3 (2022)
 27. Huang, G., Mao, J., Huang, F., Liu, F., Luo, X., Liang, Y., Lu, J., Wang, X., Liu, P., Fu, R., Huang, R., Huang, S.L.: Exposure bias can alleviate itself via directional and frequency rectification in flow matching (2026), <https://arxiv.org/abs/2606.28226>
 28. Jiang, C., Zhu, C., Guo, H., Tan, G., Liu, C., Li, K.: Mcbl-unet: A hybrid mamba-cnn boundary enhanced light-weight unet for placenta ultrasound image segmentation. *IEEE Journal of Biomedical and Health Informatics* (2025)
 29. Jocher, G., Qiu, J., Chaurasia, A.: Ultralytics YOLO (Jan 2023), <https://github.com/ultralytics/ultralytics>
 30. Kirillov, A., He, K., Girshick, R., Rother, C., Dollár, P.: Panoptic segmentation. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 9404–9413 (2019)
 31. Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., et al.: Segment anything. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 4015–4026 (2023)
 32. Li, F., Zhang, H., Xu, H., Liu, S., Zhang, L., Ni, L.M., Shum, H.Y.: Mask dino: Towards a unified transformer-based framework for object detection and segmentation. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 3041–3050 (2023)
 33. Li, Z., Yang, X., Choi, K., Zhu, W., Hsieh, R., Kim, H., Lim, J.H., Ji, S., Lee, B., Yan, X., et al.: Mmsci: A dataset for graduate-level multi-discipline multimodal scientific understanding. arXiv preprint arXiv:2407.04903 (2024)
 34. Lin, B., Emami, N., Santos, D.A., Luo, Y., Banerjee, S., Xu, B.X.: A deep learned nanowire segmentation model using synthetic data augmentation. *npj Computational Materials* **8**(1), 88 (2022)
 35. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 10012–10022 (2021)

36. López Gutiérrez, J.D., Abundez Barrera, I.M., Torres Gómez, N.: Nanoparticle detection on sem images using a neural network and semi-synthetic training data. *Nanomaterials* **12**(11), 1818 (2022)
37. Lozano, A., Nirschl, J., Burgess, J., Gupte, S.R., Zhang, Y., Unell, A., Yeung, S.: Micro-bench: A microscopy benchmark for vision-language understanding. *Advances in Neural Information Processing Systems* **37**, 30670–30685 (2024)
38. Lyu, C., Zhang, W., Huang, H., Zhou, Y., Wang, Y., Liu, Y., Zhang, S., Chen, K.: RtmDET: An empirical study of designing real-time object detectors. *arXiv preprint arXiv:2212.07784* (2022)
39. Ma, B., Yin, X., Ban, X., Huang, H., Zhang, N., Wang, H., Xue, W.: Data privacy protection in microscopic image analysis for material data mining. *arXiv preprint arXiv:2111.07892* (2021)
40. Mill, L., Wolff, D., Gerrits, N., Philipp, P., Kling, L., Vollnhals, F., Ignatenko, A., Jaremenko, C., Huang, Y., De Castro, O., et al.: Synthetic image rendering solves annotation problem in deep learning nanoparticle segmentation. *Small Methods* **5**(7), 2100223 (2021)
41. Mokady, R., Hertz, A., Aberman, K., Pritch, Y., Cohen-Or, D.: Null-text inversion for editing real images using guided diffusion models. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 6038–6047 (2023)
42. Okunev, A.G., Mashukov, M.Y., Nartova, A.V., Matveev, A.V.: Nanoparticle recognition on scanning probe microscopy images using computer vision and deep learning. *Nanomaterials* **10**(7), 1285 (2020)
43. Opatz, T., Feldhoff, K., Wiemer, H., Ihlenfeldt, S.: Sharing research data in collaborative material science and engineering projects. *Data* **10**(4), 53 (2025)
44. Pachitariu, M., Rariden, M., Stringer, C.: Cellpose-sam: superhuman generalization for cellular segmentation. *bioRxiv* pp. 2025–04 (2025)
45. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 10684–10695 (2022)
46. Rühle, B., Krumrey, J.F., Hodoroaba, V.D.: Workflow towards automated segmentation of agglomerated, non-spherical particles from electron microscopy images using artificial neural networks. *Scientific reports* **11**(1), 4942 (2021)
47. Ruiz, N., Li, Y., Jampani, V., Pritch, Y., Rubinstein, M., Aberman, K.: Dream-booth: Fine tuning text-to-image diffusion models for subject-driven generation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 22500–22510 (2023)
48. Schmidt, U., Weigert, M., Broaddus, C., Myers, G.: Cell detection with star-convex polygons. In: *Medical Image Computing and Computer Assisted Intervention - MICCAI 2018 - 21st International Conference, Granada, Spain, September 16-20, 2018, Proceedings, Part II*. pp. 265–273 (2018). https://doi.org/10.1007/978-3-030-00934-2_30
49. Shi, B., Patel, M., Yu, D., Yan, J., Li, Z., Petriw, D., Pruyn, T., Smyth, K., Passeport, E., Miller, R.D., et al.: Automatic quantification and classification of microplastics in scanning electron micrographs via deep learning. *Science of The Total Environment* **825**, 153903 (2022)
50. Shi, R., Wang, W., Li, Z., He, L., Sheng, K., Ma, L., Du, K., Jiang, T., Huang, T.: U-risc: an annotated ultra-high-resolution electron microscopy dataset challenging the existing deep learning algorithms. *Frontiers in Computational Neuroscience* **16**, 842760 (2022)

51. Stringer, C., Wang, T., Michaelos, M., Pachitariu, M.: Cellpose: a generalist algorithm for cellular segmentation. *Nature methods* **18**(1), 100–106 (2021)
52. Stuckner, J., Harder, B., Smith, T.M.: Microstructure segmentation with deep learning encoders pre-trained on a large microscopy dataset. *npj Computational Materials* **8**(1), 200 (2022)
53. Tao, H., Huang, C., Wang, N., Lyu, H., Zhang, L., Ke, G., Fang, X.: Omniscience: A large-scale multi-modal dataset for scientific image understanding. *arXiv preprint arXiv:2602.13758* (2026)
54. Vagenknecht, M., Soukup, J., Chen, A., Irizarry, R.: A deep learning solution for particle size analysis in low resolution inline microscopy images based on generative adversarial network. *Powder Technology* **426**, 118641 (2023)
55. Wang, X., Zhang, R., Kong, T., Li, L., Shen, C.: Solov2: Dynamic and fast instance segmentation. *Advances in Neural information processing systems* **33**, 17721–17732 (2020)
56. Wu, Y., Kirillov, A., Massa, F., Lo, W.Y., Girshick, R.: Detectron2. <https://github.com/facebookresearch/detectron2> (2019)
57. Xie, S., Girshick, R., Dollár, P., Tu, Z., He, K.: Aggregated residual transformations for deep neural networks. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 1492–1500 (2017)
58. Yildirim, B., Cole, J.M.: Bayesian particle instance segmentation for electron microscopy image quantification. *Journal of Chemical Information and Modeling* **61**(3), 1136–1149 (2021)
59. Zhao, Q., Li, Y., Sun, Q., Yan, Z.: Resilphase: Plug-and-play phase mapping and noise-resilient macro-trajectory extrapolation for diffusion acceleration (2026), <https://arxiv.org/abs/2606.26769>
60. Zhu, C., Shao, J., Lin, J., Wang, Y., Wang, J., Tang, J., Li, K.: fmri2ges: Co-speech gesture reconstruction from fmri signal with dual brain decoding alignment. *IEEE Transactions on Circuits and Systems for Video Technology* (2025)