

Addressing Data Scarcity in Materials Science Research with Deep Generative Models

Anonymous WACV Applications Track submission

Paper ID 6

Abstract

001 *Developments in deep learning have facilitated the auto-*
002 *matic visual analysis of scientific data, driving forward ex-*
003 *ploratory research. However, these approaches depend on*
004 *large amounts of expert-annotated data for effective train-*
005 *ing, which is difficult to come by in narrow application do-*
006 *main. In this work, we address the challenges that come*
007 *with performing visual analysis of high-speed x-ray phase*
008 *contrast images of the combustion of molten metal particles.*
009 *In this case, manual annotations of thousands of complex*
010 *frames is highly impractical. To address this, we propose*
011 *a synthetic data generation framework that eliminates the*
012 *need for large-scale manual labelling by generating image-*
013 *annotation pairs for the task of image segmentation. We*
014 *first train a denoising diffusion model with a small number*
015 *of annotated samples to generate image-binary mask pairs.*
016 *We use the predictions of a fine-tuned segmentation foun-*
017 *dation model to create a multi-class semantic annotations*
018 *for the synthetic dataset. We apply our framework on x-*
019 *ray phase contrast videos of particle combustion. From 200*
020 *manually annotated frames, we generate 10,000 synthetic*
021 *image-annotation pairs. We demonstrate that training se-*
022 *mantic segmentation models with our generated synthetic*
023 *data yields significant boost in performance.*

024 1. Introduction

025 Recent advancements in deep learning have greatly bene-
026 fitted scientific discovery in many fields of research [2, 20].
027 In particular, this has accelerated the processing of visual
028 information, enabling large scale statistical analysis of data,
029 having applications in fields such as cell microscopy [7, 24],
030 bioimage analysis [9], and high energy physics [5]. How-
031 ever, these data-driven approaches depend on the availabil-
032 ity of large amounts of data annotated for the required vi-
033 sual perception task. This becomes particularly challeng-
034 ing when applying general deep learning tools in narrow
035 domain scenarios such as those listed above. This is due

to the fact that domain-specific, expert-level knowledge
is required for annotating images, and obtaining a large
amount of annotated data becomes expensive and imprac-
tical. There are a few recent studies proposed synthetic
data generation pipelines. In [13], a framework for syn-
thetic data generation using large language models and text-
to-image diffusion models was proposed, employing cross-
attention maps to generate pseudo-semantic segmentation
labels. Similarly, [22] transforms diffusion models into
tools for generating labeled datasets, introducing a percep-
tion decoder to produce high-quality labels like segmenta-
tion masks and depth maps alongside synthetic samples.
In [4], authors proposed a framework to generate pseudo
binary segmentation labels utilizing off-the-shelf medical
foundation models.

In this work, we address this challenge as it applies
to materials science research. Specifically, we look at a
large dataset of high-speed x-ray phase contrast imaging
(XPCI) videos [21], that captures in-situ behavior of molten
metal particles as they combust. Elemental and composite
metal powders are useful for combustion applications due
to their high combustion enthalpies and burn temperatures
[15]. This dataset characterizes the combustion properties
of reactive metal powders that could be used in munitions
for rapid neutralization of chemical warfare agents (CWAs).
Predicting particle attributes such as behavior, shape, tra-
jectory, and frequency of occurrence thus becomes an im-
portant task for research of these processes. However, due
to the complexity of the combustion scenes tracking the
evolution of particles' combustion behavior cannot be done
without using sophisticated and high-throughput data driven
methods.

In order to leverage deep learning tools for visual reason-
ing tasks in the absence of a large-scale annotated dataset,
we propose a framework for the generation of synthetic
XPCI images of particle combustion events along with their
semantic segmentation annotation masks. Specifically, we
train a denoising diffusion probabilistic model (DDPM)
[11] on a small number of manually annotated images and
generate a large dataset of image and binary mask pairs.

036
037
038
039
040
041
042
043
044
045
046
047
048
049
050
051
052
053
054
055
056
057
058
059
060
061
062
063
064
065
066
067
068
069
070
071
072
073
074
075

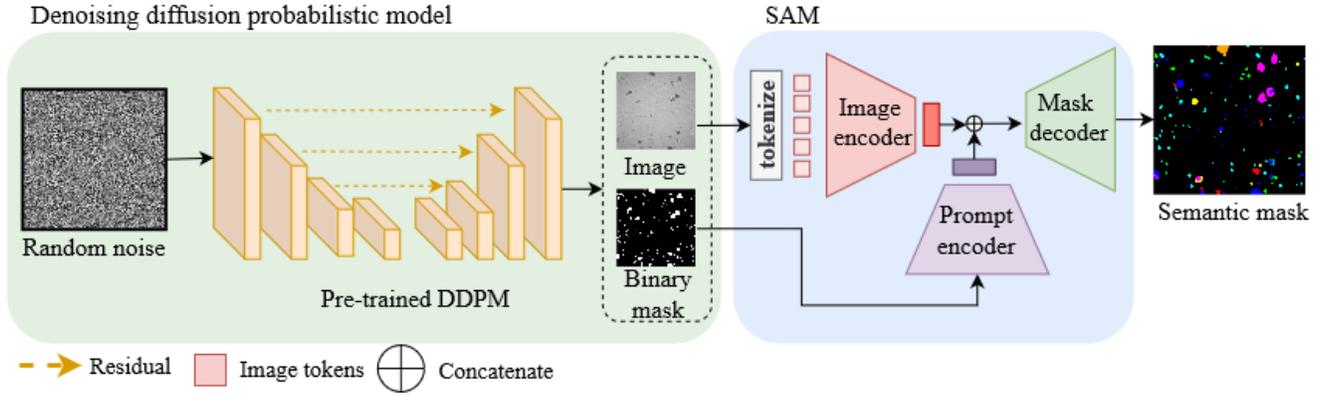


Figure 1. An illustration of the synthetic dataset creation pipeline using a pre-trained diffusion model and the fine-tuned SAM. In this framework, we employ a DDPM to synthesize realistic images together with their corresponding binary masks, which are then fed into the fine-tuned SAM model; SAM uses these binary masks as prompts to generate pseudo-annotated synthetic images with detailed semantic segmentation masks.

076 Here, the binary mask indicates the location of all types of
 077 particles present in the image. We then leverage the founda-
 078 tion Segment Anything Model (SAM) [12] fine-tuned on
 079 the same small set of manually annotated images to predict
 080 semantic categories of each particle, using the binary mask
 081 as a prompt to SAM. Through this process, we create a large
 082 dataset of synthetic images and their corresponding seman-
 083 tic pseudo-annotations. We demonstrate the effectiveness
 084 of the additional synthetic training data by training state-
 085 of-the-art segmentation baselines [3, 10, 16]. We show that
 086 training these models on a combination of real and synthetic
 087 images results in a significant boost in performance.

- 088 • We propose a framework for the generation of realistic
 089 images-mask pairs of x-ray phase-contrast images of par-
 090 ticle combustion events to address the scarcity of anno-
 091 tated data for the task of semantic segmentation.
- 092 • With the addition of our generated synthetic data in the
 093 training set, we demonstrate boosts in segmentation per-
 094 formance of TransUNet [3], SwinUNETR [10] and EM-
 095 CAD [16] across multiple metrics.

096 2. Methodology

097 The objective of this work is to generate a large number of
 098 high-quality synthetic images of particle combustion events
 099 and their corresponding semantic segmentation masks. An
 100 illustration of the framework can be seen in Figure 1.

101 First, a diffusion-based generation model is trained on
 102 a small number of manually annotated frames to gener-
 103 ate realistic synthetic images along with their binary seg-
 104 mentation masks. In order to achieve realistic binary mask
 105 generation via learning joint distribution, we concatenate
 106 target images with the corresponding ground truth binary
 107 masks and define the diffusion objective with this concate-
 108 nated representation. This process is detailed in Section

2.1. Through this step, we obtain synthetic frames of parti-
 cle combustion events and their corresponding binary masks
 showing particle locations without class labels.

Next, we assign semantic categories to each particle
 mask generated in the previous step. Towards this, we adapt
 SAM with the same small set of annotated real images. This
 step enhances the model’s capability to accurately annotate
 the synthetic images, creating reliable pseudo-annotations
 that closely mimic expert annotations. Figure 1 illustrates
 the overall framework, from synthetic image and binary
 mask generation to leveraging adapted SAM to predict se-
 mantic categories. Figure 2 illustrates the SAM adaptation
 framework.

Finally, we train a segmentation network with the syn-
 thetic dataset generated in the previous steps. This strategy
 compensates for the limited availability of real annotated
 images in materials science, ultimately improving the seg-
 mentation performance. Each step of the framework is de-
 tailed below.

2.1. Synthetic image generation

We utilize denoising diffusion probabilistic models
 (DDPMs) for image generation due to their demon-
 strated ability in learning complex data distributions and
 creating high quality synthetic images [11, 14]. DDPM
 relies on two distinct phases, forward and backward dif-
 fusion. In the forward diffusion process, normally dis-
 tributed random noise is gradually added to the data
 sample with a predetermined variance schedule until time
 index T , then during backward diffusion phase, a deep
 denoiser network, typically a UNet [11], is employed to
 predict the added noise in each time step. The forward
 diffusion function can be described as

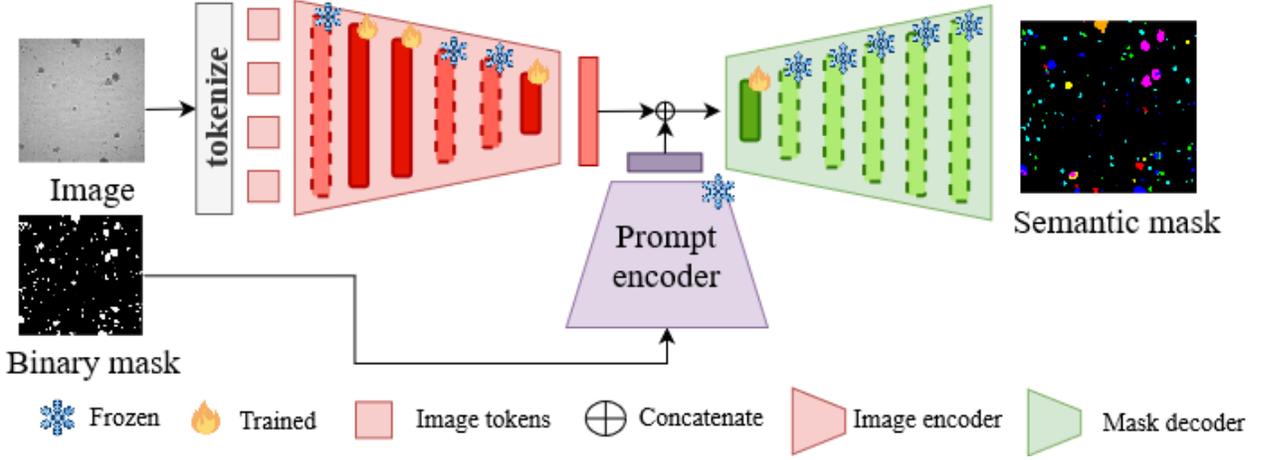


Figure 2. An illustrative overview of the proposed SAM fine-tuning framework, highlighting which components of the architecture are updated during training and which remain frozen. For each block in the network, we indicate its optimization status using visual markers: layers annotated with a fire icon are trainable, whereas those marked with a cold icon are kept fixed. In this setup, we first tokenize the input images and feed them to the Vision Transformer encoder, while binary masks indicating the locations of particles in the image are provided to the prompt encoder; the mask decoder then combines these representations to produce the final semantic segmentation masks.

$$x_t = \sqrt{\alpha_t} x_{t-1} + \sqrt{1 - \alpha_t} \epsilon \quad (1)$$

where α_t denotes the pre-fixed variance schedule, x_t and x_{t-1} are the data samples at time t and $t-1$ respectively and $\epsilon \sim \mathcal{N}(0, I)$. After the forward diffusion process, a deep denoiser network is trained to learn the mapping between data samples from time t to $t-1$ via estimating added noise, the objective function can be written as follows

$$L(\theta) = \mathbb{E}_{t, x_0, \epsilon} \left[\left\| \epsilon - \epsilon_\theta \left(\sqrt{\alpha_t} x_0 + \sqrt{1 - \alpha_t} \epsilon, t \right) \right\|^2 \right], \quad (2)$$

where $\bar{\alpha}_t = \prod_{m=1}^t \alpha_m$, x_0 is the data sample at time zero and $\epsilon_\theta(\cdot)$ denotes the predicted noise with network parameters θ . In our pipeline, we train an unconditional improved-DDPM [14] using the available small annotated set and their corresponding binary masks. Therefore, with the above notation x_0 becomes

$$x_0 \in \mathbb{R}^{H \times W \times 2}, \quad x_0 = \left(x_0^{(1)}, x_0^{(2)} \right), \quad (3)$$

where $x_0^{(1)} \in \mathbb{R}^{H \times W}$ and $x_0^{(2)} \in \mathbb{R}^{H \times W}$ denote clean image and ground truth binary mask with height (H) and width (W) respectively. We generate binary masks to be used in training from the annotated semantic masks. Then, during the inference we synthesize images and corresponding binary masks from random noise with improved-DDPM [14] sub-sampling algorithm. In our experiments we observed the binary mask generation along with the synthetic image is beneficial to guide SAM during the fine-tuning phase by directing the network through prompt encoder (see

Figure 2 for details). A few synthesized samples generated with our diffusion model along with real samples are illustrated in Figure 5 and Figure 4.

2.2. SAM Adaptation

The Segment Anything Model (SAM) is composed of three sub-networks: an image encoder, a prompt encoder, and a mask decoder. The image encoder is a vision transformer [6] trained to extract meaningful features from the input image. The prompt encoder leverages convolutional layers to incorporate informative prompts, such as points, bounding boxes, or image masks, enhancing the accuracy of object segmentation. Finally, the mask decoder combines the encoded image and prompt vectors using a transformer layer and then utilizes transposed convolutions to up-sample and predict precise segmentation masks.

Following [8], we incorporate small adapter networks consisting of multi-layer perceptrons into the attention blocks of transformer layers of both the image encoder and mask decoder networks. During training, only the adapter layers are updated. This decreases the number of trainable parameters and allows the network to retain the knowledge present in the original weights while also fine-tuning the model for our specific task. We utilize a small set of annotated samples and their corresponding binary masks with a combination of cross-entropy and dice losses [18] during fine-tuning. Throughout the process, we constantly used binary masks as the prompts apart from the target images to guide the network more precisely to deal with the complex nature of X-ray phase contrast images and small particle sizes. The fine-tuning procedure is illustrated in Figure 2. We illustrate the quality of the pseudo-annotations per-

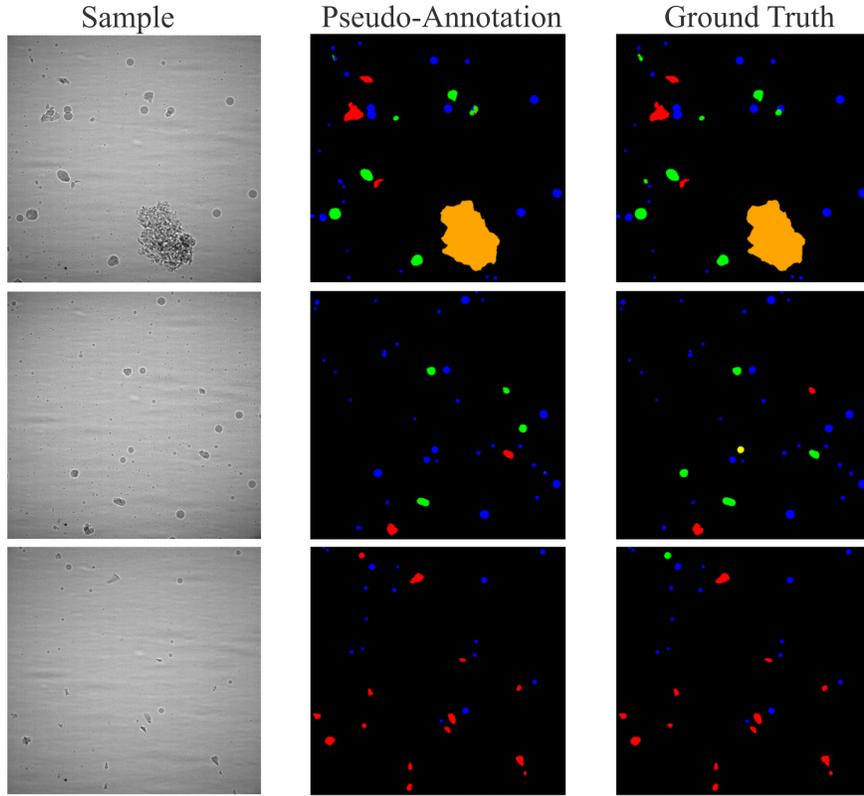


Figure 3. Generated pseudo-annotations for real XPCI images, all produced by our fine-tuned SAM model, compared against the corresponding ground truth segmentation masks. The qualitative alignment between the pseudo-labels and the manually annotated masks highlights the accuracy and reliability of our pseudo-annotation pipeline, indicating that the fine-tuned SAM can effectively capture particle locations and shapes in realistic experimental data.

197 formed by fine-tuned SAM in Figure 3.

198 2.3. Semantic segmentation

199 We combined real and synthetic images in varying propor-
 200 tions, progressively increasing the number of synthetic im-
 201 ages in the final training dataset. While our framework is
 202 compatible with any segmentation network, we specifically
 203 selected three state-of-the-art architectures—TransUNet
 204 [3], a vision transformer-based model; SwinUNETR [10],
 205 a Swin transformer-based network; and EMCAD [16], a
 206 CNN-based model—to demonstrate its efficiency across
 207 different backbone designs.

208 3. Experiments

209 We performed an extensive evaluation to show the benefits
 210 of our proposed approach in generating synthetic datasets.
 211 Our results are grouped with different proportion of the syn-
 212 thetic samples in the final training set. We employ the Dice
 213 similarity coefficient [18], normalized surface Dice (NSD)
 214 [17], and micro F1 scores as the primary metrics for com-
 215 parison in our evaluation. The Dice and NSD scores are

used to assess the overall overlap between the target and
 predicted segmentation masks, providing an overall mea-
 sure of segmentation quality. The micro F1 score accounts
 for the occurrence of false positives and false negatives for
 all classes, considering the imbalanced nature of the dataset,
 contributing to a more robust evaluation.

XPCI Dataset: High-speed XPCI is performed on a vari-
 ety of aluminium-based metal composite and pure titanium,
 hafnium, and zirconium powders using a high energy x-ray
 synchrotron to assess how variations in chemistry impact
 combustion behavior. These experiments were performed in
 a similar manner with a single Al:Zr powder chemistry [19]
 and utilize synchronous hyper-spectral imaging of the com-
 bustion scene to measure combustion temperatures [1]. This
 resulted in the capture of a data set of 500,00 video frames
 containing tens of thousands of particles creating a highly
 imbalanced segmentation dataset due to small particle sizes
 compared to the background, and also due to underrepre-
 sented particle classes. We selected a 200 sample subset of
 the XPCI for manual annotation each with size 624×640

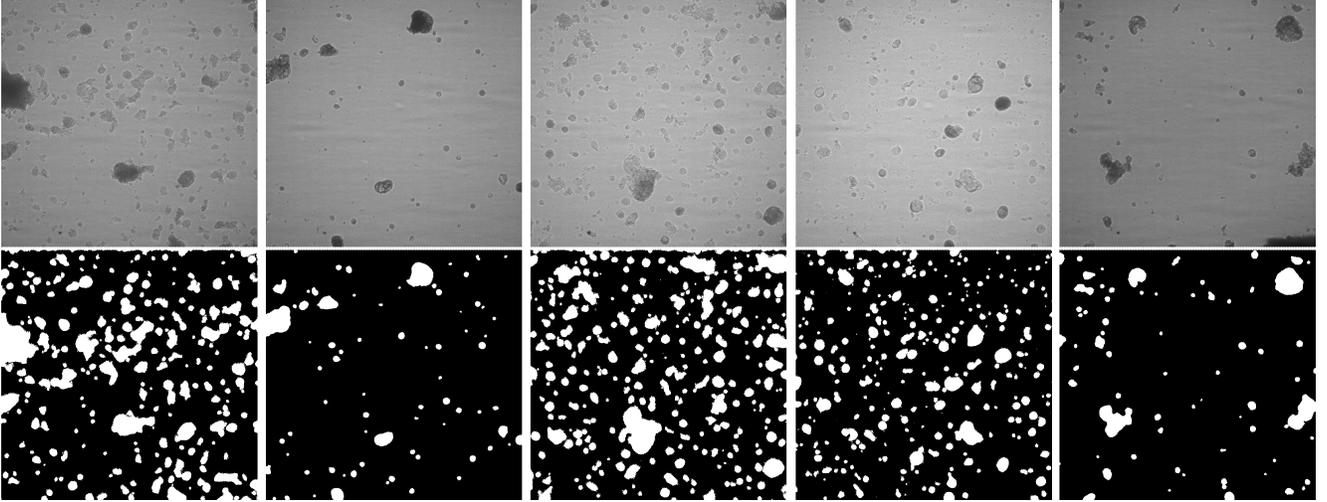


Figure 4. Visualization of synthetic images and their corresponding binary masks, all generated automatically by our proposed pipeline. The examples illustrate the diversity and realism of the synthesized particle patterns, as well as the precise spatial alignment between the images and their binary annotations.

Table 1. Semantic segmentation performance of TransUNet [3] with different numbers of synthetic images generated with SAM-t.

Dataset	Dice Score	NSD	Micro F1
200 real images	0.7791	0.5490	0.6933
200 real images + 500 synthetic images	0.7807	0.5195	0.6854
200 real images + 1,000 synthetic images	0.7985	0.5731	0.7611
200 real images + 2,000 synthetic images	0.7936	0.5605	0.7508
200 real images + 5,000 synthetic images	0.8034	0.5780	0.7531

Table 2. Semantic segmentation performance of Swin UNETR [10] with different numbers of synthetic images generated with SAM-b.

Dataset	Dice Score	NSD	Micro F1
200 real images	0.7628	0.5142	0.7243
200 real images + 500 synthetic images	0.7646	0.5233	0.7373
200 real images + 1,000 synthetic images	0.7755	0.5127	0.7272
200 real images + 2,000 synthetic images	0.7756	0.5393	0.7444
200 real images + 5,000 synthetic images	0.7884	0.5576	0.7531

236 pixels (shown in Figure 5), featuring 9 distinct classes of
 237 particles derived from X-ray phase contrast videos of com-
 238 busting composite and elemental powders. Additionally, we
 239 utilized a separate subset of 59 annotated frames as the test
 240 set to evaluate the final segmentation performance.

241 3.1. Implementation details

242 We utilized a workstation equipped with 8 NVIDIA RTX
 243 A5000 GPUs with 24GB memory for all experiments.
 244 During the pre-training of diffusion model, we utilize the
 245 improved-DDPM [14] pipeline with cosine noise sched-
 246 ule, 4000 forward diffusion steps, 250 sampling steps and
 247 AdamW optimizer with a learning rate of 10^{-4} . DDPM

training is lasted for 50,000 steps with a batch size of 1. 248

We utilize two variants of SAM, SAM-base (SAM-b) 249
 [12] and SAM-tiny (SAM-t) [23] where SAM-tiny utilizes 250
 a much smaller image encoder trained via knowledge dis- 251
 tillation, while SAM-base utilizes a larger image encoder 252
 (ViT-B [6]). During the SAM fine-tuning and final model 253
 training, AdamW optimizer is utilized with beta parameters 254
 (0.99, 0.999) and a weight decay factor of 0.1. SAM is fine- 255
 tuned for 100 epochs with batch size of 1 and 10^{-4} learning 256
 rate with polynomial decay rate with power 0.9. Similarly, 257
 TransUNet, SwinUNETR and EMCAD are trained for 100 258
 epochs with learning rate 10^{-4} with a batch size of 4. 259

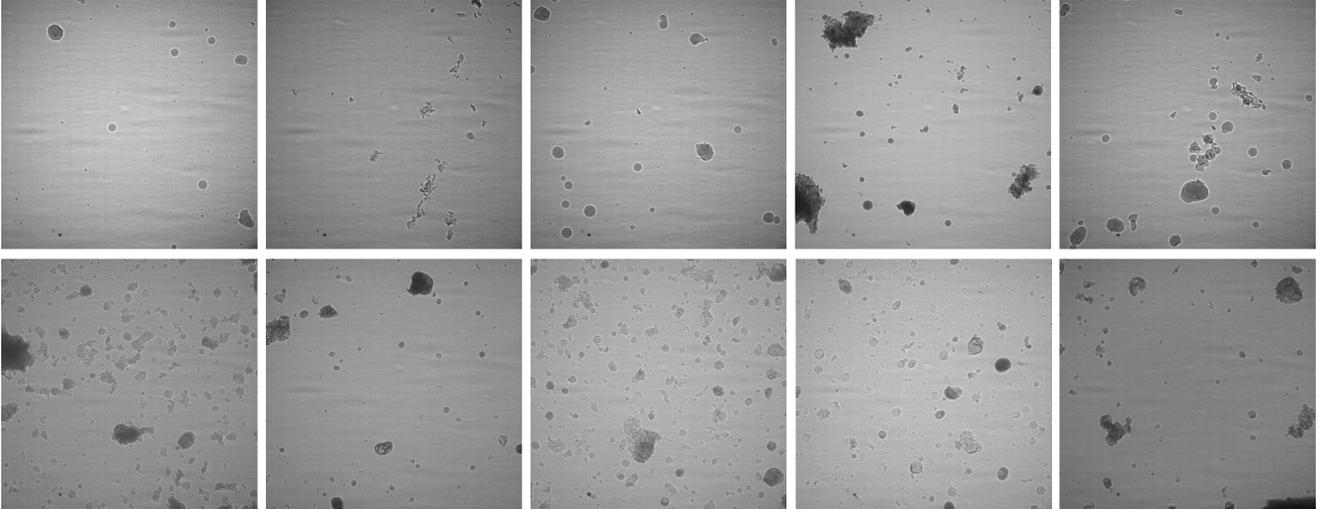


Figure 5. Visualization of real samples from the XPCI particle combustion dataset (first row) and generated synthetic images (second row). Please zoom-in for the best quality.

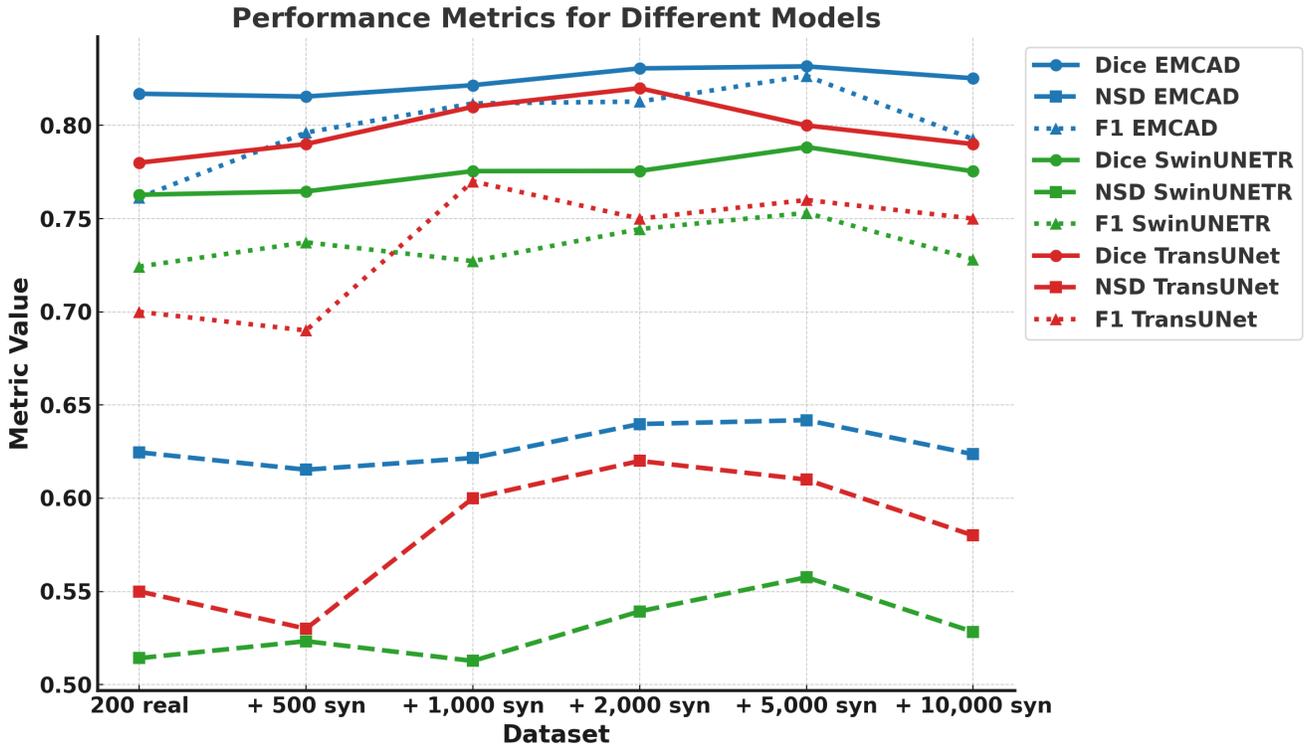


Figure 6. Semantic segmentation performance of different backbones against increasing amounts of synthetic training data.

260 **3.2. Results**

261 We demonstrate the effectiveness of our dataset creation
 262 pipeline by measuring the final segmentation performance
 263 under different proportions of synthetic samples in the training
 264 dataset, starting from a small only real dataset (with 200
 265 samples) and gradually increasing the synthetic sample size.

We illustrate the performance metrics of different seg-
 266 mentation models as a graph in Figure 6. From the graph,
 267 we observe that adding 1,000 to 5,000 synthetic images pro-
 268 vides a noticeable boost in performance, suggesting that a
 269 moderate amount of synthetic data can enhance segmenta-
 270 tion performance by providing additional training diversity.
 271

Table 3. Semantic segmentation performance of EMCAD [16] with different numbers of synthetic images generated with SAM-b.

Dataset	Dice Score	NSD	Micro F1
200 real images	0.8170	0.6246	0.7613
200 real images + 500 synthetic images	0.8155	0.6153	0.7962
200 real images + 1,000 synthetic images	0.8216	0.6216	0.8118
200 real images + 2,000 synthetic images	0.8306	0.6398	0.8128
200 real images + 5,000 synthetic images	0.8317	0.6419	0.8266

Table 4. Semantic segmentation performance of TransUNet [3] with different numbers of synthetic images generated with SAM-b.

Dataset	Dice Score	NSD	Micro F1
200 real images	0.7791	0.5491	0.6933
200 real images + 500 synthetic images	0.7807	0.5387	0.6993
200 real images + 1,000 synthetic images	0.8100	0.6034	0.7752
200 real images + 2,000 synthetic images	0.8196	0.6126	0.7504
200 real images + 5,000 synthetic images	0.8058	0.5963	0.7656

Beyond 5,000 synthetic images, the improvements plateau or slightly decline, indicating diminishing returns. This suggests that while synthetic data is beneficial up to a point, too much synthetic data could lead to overfitting or redundancy for these models.

Detailed metrics with synthetic dataset generated with SAM-b have been shown in Tables 4, 3, and 2 for TransUNet, EMCAD and SwinUNETR respectively. From the tables, we observe that the micro-F1 score shows a greater improvement with the addition of synthetic samples compared to the NSD and Dice scores, suggesting that the model improves with synthetic samples at managing the imbalanced dataset by reducing overall false positives and false negatives across all classes.

Ablation Studies: We analyze the effects of different network components on the final segmentation performance. First, we evaluate the quality of synthetic datasets generated using different variants of SAM, specifically SAM-t and SAM-b shown in Tables 1 and 4 respectively, which correspond to different model sizes. In Table 4, there is a noticeable improvement in performance metrics compared to the baseline up to 2,000 synthetic samples, where the highest performance is observed. However, Table 1 reveals a different pattern, with peak performance achieved at 5,000 synthetic samples. This suggests that the quality of the synthetic dataset significantly impacts the performance metrics, implying that lower-quality synthetic datasets require a larger number of samples to achieve comparable performance. Second, we demonstrated the impact of binary mask prompts on the segmentation performance of SAM, highlighting their role in enhancing the quality of the synthetic datasets. Table 5 demonstrates that the performance

gain from using binary mask prompts is much more pronounced for smaller-sized model. This finding highlights that smaller SAM variants, like SAM-t, benefit significantly from the inclusion of binary masks during synthetic dataset generation likely due to their limited capacity to capture complex features without additional guidance.

Table 5. Semantic segmentation performance of fine-tuned SAM with different configurations.

Model & Configuration	Dice Score
SAM-t without binary mask prompt	0.6610
SAM-t with binary mask prompt	0.7876
SAM-b without binary mask prompt	0.8127
SAM-b with binary mask prompt	0.8181

4. Conclusion

In this work, we addressed the critical challenge of annotated data scarcity in materials science research, with a particular focus on the localization and categorization of particles in XPCI images of combustion events. To mitigate the high cost and difficulty of obtaining expert-labeled data in this domain, we proposed a complete pipeline for generating a large dataset of synthetic images together with their corresponding semantic segmentation annotations. Our approach leverages synthetic data to approximate the visual characteristics and particle distributions observed in real experimental XPCI frames, while automatically producing pixel-level labels that would otherwise require extensive manual effort.

We evaluated the effectiveness of the proposed pipeline

325 by training state-of-the-art semantic segmentation models
326 on different mixtures of real and synthetic data, system-
327 atically varying the proportion of synthetic samples used
328 during training. The experimental results show that mod-
329 els trained with the inclusion of our generated synthetic
330 data consistently achieve improved performance compared
331 to those trained solely on limited real data, indicating that
332 the synthetic dataset provides meaningful and complemen-
333 tary supervision. These findings suggest that our synthetic
334 data generation framework can serve as a practical and scal-
335 able solution for enhancing model performance in scenarios
336 where high-quality annotations are scarce or expensive to
337 obtain.

338 References

- 339 [1] M. Alemohammad, E. R. Wainwright, J. P. Stroud, T. P.
340 Weihs, and M. A. Foster. Kiloherz frame rate snapshot hy-
341 perspectral imaging of metal reactive materials. *Applied Op-*
342 *tics*, 59, 2020. 4
- 343 [2] R. Chang, Y.-X. Wang, and E. Ertekin. Towards overcom-
344 ing data scarcity in materials science: unifying models and
345 datasets with a mixture of experts framework. *npj Computa-*
346 *tional Materials*, 8(1):242, 2022. 1
- 347 [3] J. Chen, Y. Lu, Q. Yu, X. Luo, E. Adeli, Y. Wang, L. Lu, A. L.
348 Yuille, and Y. Zhou. Transunet: Transformers make strong
349 encoders for medical image segmentation. *arXiv preprint*
350 *arXiv:2102.04306*, 2021. 2, 4, 5, 7
- 351 [4] T. Deshpande, E. Prakash, E. G. Ross, C. Langlotz, A. Ng,
352 and J. M. J. Valanarasu. Auto-generating weak labels for
353 real & synthetic data to improve label-scarce medical image
354 segmentation. *arXiv preprint arXiv:2404.17033*, 2024. 1
- 355 [5] F. A. Di Bello, S. Ganguly, E. Gross, M. Kado, M. Pitt, L.
356 Santi, and J. Shlomi. Towards a computer vision particle
357 flow. *The European Physical Journal C*, 81:1–14, 2021. 1
- 358 [6] A. Dosovitskiy. An image is worth 16x16 words: Trans-
359 formers for image recognition at scale. *arXiv preprint*
360 *arXiv:2010.11929*, 2020. 3, 5
- 361 [7] M. Eisenstein. Ai under the microscope: the algorithms pow-
362 ering the search for cells. *Nature*, 623(7989):1095–1097,
363 2023. 1
- 364 [8] H. Gu, H. Dong, J. Yang, and M. A. Mazurowski. How to
365 build the best medical image segmentation algorithm using
366 foundation models: a comprehensive empirical study with
367 segment anything model. *arXiv preprint arXiv:2404.09957*,
368 2024. 3
- 369 [9] A. Hallou, H. G. Yevick, B. Dumitrascu, and V. Uhlmann.
370 Deep learning for bioimage analysis in developmental biol-
371 ogy. *Development*, 148(18):dev199616, 2021. 1
- 372 [10] A. Hatamizadeh, V. Nath, Y. Tang, D. Yang, H. R. Roth, and
373 D. Xu. Swin unetr: Swin transformers for semantic segmen-
374 tation of brain tumors in mri images. In *International MIC-*
375 *CAI Brainlesion Workshop*, pages 272–284. Springer, 2021.
376 2, 4, 5
- 377 [11] J. Ho, A. Jain, and P. Abbeel. Denoising diffusion proba-
378 bilistic models. *Advances in Neural Information Processing*
379 *Systems*, 33:6840–6851, 2020. 1, 2
- [12] A. Kirillov, E. Mintun, N. Ravi, H. Mao, C. Rolland, L.
Gustafson, T. Xiao, S. Whitehead, A. C. Berg, W.-Y. Lo, P.
Dollar, and R. Girshick. Segment anything. In *Proceedings*
of the IEEE/CVF International Conference on Computer Vi-
sion, pages 4015–4026, 2023. 2, 5
- [13] Q. Nguyen, T. Vu, A. Tran, and K. Nguyen. Dataset dif-
fusion: Diffusion-based synthetic data generation for pixel-
level semantic segmentation. *Advances in Neural Informa-*
tion Processing Systems, 36, 2024. 1
- [14] A. Q. Nichol and P. Dhariwal. Improved denoising diffu-
sion probabilistic models. In *International Conference on*
Machine Learning, pages 8162–8171. PMLR, 2021. 2, 3, 5
- [15] S. M. Peiris and N. Bolden-Frazier. Applications of reactive
materials in munitions. In *Shock Phenomena in Granular*
and Porous Materials, pages 165–191. Springer, 2019. 1
- [16] M. M. Rahman, M. Munir, and R. Marculescu. Emcad: Effi-
cient multi-scale convolutional attention decoding for medi-
cal image segmentation. In *Proceedings of the IEEE/CVF*
Conference on Computer Vision and Pattern Recognition,
pages 11769–11779, 2024. 2, 4, 7
- [17] S. Seidlitz, J. Sellner, J. Odenthal, B. Özdemir, A. Studier-
Fischer, S. Knödler, L. Ayala, T. J. Adler, H. G. Kenngott, M.
Tizabi, M. Wagnet, F. Nickel, B. Müller-Stich, and L. Maier-
Hain. Robust deep learning-based semantic organ segmen-
tation in hyperspectral images. *Medical Image Analysis*, 80:
102488, 2022. 4
- [18] C. H. Sudre, W. Li, T. Vercauteren, S. Ourselin, and M. J.
Cardoso. Generalised dice overlap as a deep learning loss
function for highly unbalanced segmentations. In *Deep*
Learning in Medical Image Analysis and Multimodal Learn-
ing for Clinical Decision Support: Third International Work-
shop, DLMIA 2017, and 7th International Workshop, ML-
CDS 2017, Held in Conjunction with MICCAI 2017, Québec
City, QC, Canada, September 14, Proceedings 3, pages 240–
248. Springer, 2017. 3, 4
- [19] E. R. Wainwright, S. V. Lakshman, A. F. T. Leong, A. H.
Kinsey, J. D. Gibbins, S. Q. Arlington, T. Sun, K. Fezzaa,
T. C. Hufnagel, and T. P. Weihs. Viewing internal bubbling
and microexplosions in combusting metal particles via x-ray
phase contrast imaging. *Combustion and Flame*, 199:194–
203, 2019. 4
- [20] H. Wang, T. Fu, Y. Du, W. Gao, K. Huang, Z. Liu, P.
Chandak, S. Liu, P. Van Katwyk, A. Deac, A. Anandkumar,
K. Bergen, C. P. Gomes, S. Ho, P. Kohli, J. Lasenby,
J. Leskovec, T.-Y. Liu, A. Manrai, D. Marks, B. Ramsundar,
L. Song, J. Sun, J. Tang, P. Veličković, M. Welling, L. Zhang,
C. W. Coley, Y. Bengio, and M. Zitnik. Scientific discovery
in the age of artificial intelligence. *Nature*, 620(7972):47–60,
2023. 1
- [21] S.W. Wilkins, T.E. Gureyev, D. Gao, and A.W. Steven-
son. Phase-contrast imaging using polychromatic hard x-
rays. *Nature*, 384:335–338, 1996. 1
- [22] W. Wu, Y. Zhao, H. Chen, Y. Gu, R. Zhao, Y. He, H. Zhou,
M. Z. Shou, and C. Shen. Datasetdm: Synthesizing data with
perception annotations using diffusion models. *Advances in*
Neural Information Processing Systems, 36:54683–54695,
2023. 1

- 437 [23] C. Zhang, D. Han, Y. Qiao, J. U. Kim, S.-H. Bae, S.
438 Lee, and C. S. Hong. Faster segment anything: Towards
439 lightweight sam for mobile applications. *arXiv preprint*
440 *arXiv:2306.14289*, 2023. 5
- 441 [24] S. Zhou, B. Chen, E. S. Fu, and H. Yan. Computer vision
442 meets microfluidics: a label-free method for high-throughput
443 cell analysis. *Microsystems & Nanoengineering*, 9(1):116,
444 2023. 1