

Multi-Modal Super Resolution for Dense Microscopic Particle Size Estimation

Sarvesh Patil*, Chava Y P D Phani Rajanish*, and Naveen Margankunte

Abstract—Particle Size Analysis (PSA) is an important process carried out in a number of industries, which can significantly influence the properties of the final product. A ubiquitous instrument for this purpose is the Optical Microscope (OM). However, OMs are often prone to drawbacks like low resolution, small focal depth, and edge features being masked due to diffraction. We propose a powerful application of a combination of two Conditional Generative Adversarial Networks (cGANs) that Super Resolve OM images to look like Scanning Electron Microscope (SEM) images. We further demonstrate the use of a custom object detection module that can perform efficient PSA of the super-resolved particles on both, densely and sparsely packed images. The PSA results obtained from the super resolved images have been benchmarked against human annotators, and results obtained from the corresponding SEM images. The proposed models show a generalizable way of multi-modal image translation and super-resolution for accurate particle size estimation.

Index Terms—Generative adversarial networks, object detection, microscopic particle size analysis, neural networks



1 INTRODUCTION

Particle Size Analysis (PSA) enables accurate estimation of particle sizes of microspheres, offering insights into the material or the powder being formed, and is often a crucial process in manufacturing pipelines. Estimating the size of particles in a sample can help in determining many physical and chemical properties of the material like packing density, flow characteristics, reactivity, dissolution rate, texture, and appearance of the finished product [1] [2].

The most common techniques for PSA are Laser Diffraction (LD), Dynamic Light scattering (DLS) and Imaging Particle Analysis (IPA). IPA is extensively used For microspheres where particle shape and size is of extreme importance. A significant advantage of using OM image-based analysis is that, it is a non-intrusive method with reasonable costs and decent performance. However, more advanced microscopy techniques like Scanning Electron Microscopes (SEMs) that are quite expensive can be used for granular magnification to capture images with sharp boundaries and depth information.

The use of traditional Image processing techniques like edge extraction, thresholding, and segmentation, for PSA, has been shown in [3] [4] [5]. The segmentation maps obtained are used for particle sizing. A conversion factor is used to obtain real values using the size of the particles in pixels.

Methods like Hough Circle Transform (HCT) have been used for localization and counting of cell nuclei in phase-contrast microscopy images [6]. In [7], a modified HCT was used for nanoparticle diameter estimation. Although HCT performs well for high-resolution electron microscopy images, results on OM images need parameter fine-tuning due to variable lighting and background conditions across various images. This reduces the accuracy of predictions of the algorithm for unseen images.

- Sarvesh Patil and Chava Y P D Phani Rajanish contributed equally to this work.
- Sarvesh Patil, Chava Y P D Phani Rajanish, Naveen Margankunte are with HyperWorks Imaging, Bangalore, India.
E-mail: {sarvesh, rajanish, naveen}@hyperworksimaging.com




Fig. 1: Ground truth images. (a) shows bleeding of edges in OM images. (b) shows sharp edges from SEM images.

Deep neural network based methods have proven to be robust, and outperform the traditional image processing based techniques [8]. We use a Convolutional Neural Network (CNN) in the form of a custom object detection model to perform particle size analysis. Use of CNNs for cell nuclei classification, detection, and segmentation of nanoparticles has been shown in [9] [10] [11] [12]. These approaches have proven useful for particle sizing on high-quality SEM images. However, when applied on OM images, these methods fail to detect dense agglomerates of blurred particles effectively.

Although use of deep neural network based approaches seems straightforward, performance of these models is heavily dependent on the ground truths being perfect. However, accurate annotation of optical microscopy particles is a non-trivial task, even for human annotators. Due to bleeding of the edges of the particles caused by diffraction of light, accurately sizing them is an uphill task. On the other hand, from Fig. 1, we can see that annotations for SEM images can be easily created due to their sharp edges. Thus we propose that the quality of the ground truth data can be improved significantly by converting OM images to the SEM domain and then applying an object detection algorithm to perform particle size analysis.

To convert OM images to look like SEM images, image-based conditional Generative Adversarial Networks (GANs) like Pix2Pix [13] can be used. This process is known as image translation, and it converts images from OM domain to the SEM domain. However, the Pix2Pix model fails to add sharp edges and detailed texture information of the SEM particles. Hence, we use the ESRGAN model [14] for the super-resolution of the Pix2Pix output. In this regard, the core contributions of this work are as follows:

- A combination of two GANs that can be used to translate and subsequently Super Resolve OM images to look like SEM images.
- A dense object detection model to perform PSA on super-resolved images.

2 BACKGROUND

2.1 Super Resolution GANs

There has been great improvement in the quality of images generated by GANs since their inception. However, for the most part, GANs are considered as black boxes with a utopian Nash Equilibrium that may never be reached [15]. Despite being difficult to train, the synthesizing capability of GANs has proven useful in myriads of applications [16] [17]. Conditional GANs (cGANs) [18] demonstrate the capability to condition the Generator to produce desired output. cGANs have made great progress in image generation from text [19], image in-painting [20], style transfer [13] and super-resolution [21].

There have been many works presenting the use of GANs [22] to super resolve microscopy images [23] [24] [25]. In [24], cGANs have been used for cross-modal super-resolution of Stimulated Emission Depletion (STED) microscopy images. In [23], a U-Net based GAN was used for the super-resolution of low-resolution SEM images. These methods involve the use of images from similar distribution to perform super-resolution.

In [25] segmentation maps of OM images of pollen grains were super resolved to look like SEM images at a higher magnification. This approach shows use of multi-modal super resolution. However, the amount of individual preprocessing involved in segmenting each particle is not practical for our application, due to the presence of a large number of clustered particles. Hence, we propose use of a Pix2Pix model for the initial image translation and an ESRGAN model for the subsequent super resolution. Our method is reproducible and works without the need of manual preprocessing on the particles.

2.2 Particle Size Analysis

Some novel techniques in PSA involve the use of neural networks for regression and object detection. In [26], neural networks have been used to estimate sieve size for particles on a conveyor belt, based on edge detection metrics like Feret distance, best-fit rectangle, area, etc. In [27], the use of Mask R-CNN has been demonstrated for object detection, segmentation, and particle size analysis. These approaches have proven useful for high-quality electron microscopy images. However, their use on optical microscope images becomes a difficult task, primarily due to the thick blurred edges.

From the region-based object detection networks to monolithic object detection networks like YOLO [28], RetinaNet [29] and SSD [30], the field has witnessed an astonishing growth in algorithms in the last few years. Single shot detection algorithms

have dominated the field with high mean Average Precision (mAP) values while being able to run in real-time. However, even with dense object detection algorithms like RetinaNet, these algorithms often find it difficult to detect and size densely packed particles with high precision.

Our object detection algorithm also uses a single feed-forward operation but has a four-dimensional patch at its output as opposed to a fully connected layer. Also, since our microscopic particles are circular in nature, we regress their circle coordinates (x, y, r) and confidence p as opposed to box coordinates. This method has been inspired by [31], where the authors present a multi-scale approach for object detection in high-density images.

We devise an object detection algorithm – Super Resolution Particle Sizing Algorithm (SR-PSA) – which is used to detect particles in sparse as well as highly dense clusters of particles. Along with our "primary" dataset, which is used for preliminary training and validation, we use two more datasets – "secondary" and "tertiary", to measure the performance of the model on particles with varied distributions as shown in Fig. 7. We compare the results from our algorithm to those generated by human annotators.

3 METHOD




Fig. 2: (a), (b) are the ground truth images, (c) is the super resolved image, and (d) shows PSA output on (c).

3.1 Image Pair Registration

Image Pair Registration is an important step for Super-Resolution using GANs, as it aligns the image in one mode (OM) with respect to another (SEM). This enables the model to learn the one-to-one mapping between the low-resolution OM images and the high-resolution SEM images. As shown in Fig. 3 OM and SEM image pairs were obtained by using fiducial markers on the captured samples. The image pairs were registered using affine transforms in a landmark-based registration algorithm. The semi-automatic algorithm used landmark points from the user on the input image and the target image to perform registration.




Fig. 3: The pipeline describing the flow of the entire process

3.2 Multimodal Image Translation

We use the Pix2Pix GAN [32] for multimodal image translation as our baseline model. In the Pix2Pix architecture, a conditional image is provided as input to the U-Net generator. The down-sampling layers in the U-Net model form a latent representation of the input image. By using advanced operations like Spectral Normalization [33], Self-Attention [34], and Anti-Aliasing pooling [35] on the convolutional layers, we improve the latent representations of our input data. This improves the performance of the generator by simplifying the task for the up-sampling layers, thus enabling faster convergence of the model. We also implement spectral normalization and anti-aliasing pooling for the PatchGAN discriminator [36].

Loss Function For image-conditional GANs, the generator loss is obtained by calculating the error between the generated output and the real image. L1 loss function minimizes the sum of the absolute difference between the generated and the target images, while the L2 loss function minimizes the sum of squared differences between the generated and the target images.

L2 loss has been shown to induce blurriness in convolutional generative models due to its averaging effect [37] [20]. Hence, the default Pix2Pix model implements an L1 loss for the generator. L1 loss promotes sharp boundaries in natural scenes. However, we observe that use of L1 loss results in elongated boundaries for circular particles. Since our main objective is particle sizing for circular particles, these distortions were not acceptable. We propose the use of Root Mean Squared Error (RMSE) loss to mitigate these distortions. RMSE loss increased the circularity of the particles, whilst keeping blurring at a minimum. A simple comparison between the results obtained by these loss functions is demonstrated in Fig. 8. Hence, for a n input images x and the generator $G(x)$, our loss is given by:

$$L_{RMSE}(G) = \mathbb{E}_{x,y} \left[\sqrt{\frac{\sum_{i=1}^n (y - G(x))^2}{n}} \right] \quad (1)$$

Where y is the target SEM image. During the initial epochs, the generator output results in large loss values, which is limited by the RMSE loss function. As training progresses, the losses decrease to small values, which are slightly amplified by the RMSE loss function. This results in a more dynamic weight update as compared to L1 and L2 losses.

The final optimization over the network is given by:

$$\arg \min_G \max_D L_{cGAN}(G, D) + \lambda_{RMSE} L_{RMSE}(G) \quad (2)$$

Where λ controls the weight of RMSE loss function and can be considered as a hyperparameter.

3.3 Image Super Resolution

Since Pix2Pix is an image translation model, the output lacks the intricate details of the SEM particles. To add sharp edges and diverse texture information we use another cGAN for super-resolution of the Pix2Pix output. The ESRGAN [14] model has proven useful in many applications like the super-resolution of old games, low-resolution photographs, old movies, etc. Hence, we train an ESRGAN for our application due to its generalizability across various domains.

The use of Relativistic average discriminator [38] (D_{Ra}) has been shown to improve the generator by making the discriminator predict relativistic probabilities. Rather than classifying images as real or fake, D_{Ra} tries to predict the probability that a real image x_r is more realistic than a fake image x_f as shown in [14].

Loss Function For the generator firstly we use L1 loss between the fake and the real images, for maximizing the sharpness of the particles. The relativistic discriminator loss L_{GAN}^R is taken as the average value of loss generated by the real image, and the loss generated by the fake image [14]. We also incorporate perceptual loss from the default implementation. A pretrained VGG-19 [39] model is used to extract pre-activation feature maps from the generated output as well as the real image. The L1 distance between the two feature maps is used to calculate the perceptual loss $L_{perceptual}$. Thus, the total generator loss L_G is given as,

$$L_G = L_{perceptual} + \lambda L_{GAN}^R + \eta L_1 \quad (3)$$

λ and η are parameters that can be tuned to give weightage to relativistic loss and L1 loss respectively. As the perceptual loss values are very small, we use $\lambda = 0.05$ and $\eta = 0.01$. We demonstrate good quality of images being generated by the ESRGAN after certain improvements in Fig. 2c. Generated images preserve the geometrical and structural properties of particles of various sizes through the super-resolution process.

3.4 Particle Size Analysis

For the particle sizing application, instead of bounding boxes, we regress circle coordinates (x , y , and r). For single class object detection, the need for a separate classification layer is mitigated. Hence, for our application, we use a single feed-forward network to detect circles in the super-resolved images. The ground truth is created using manual annotation over the particles, and the network is trained to predict annotations with high precision and accuracy.

The SR-PSA detection module is a simple Convolutional Neural Network (CNN) that has 6 convolutional layers. The first five layers have a kernel size of 3×3 to extract features and the final layer has four feature maps where classification and regression are done in a single iteration. The layers are chosen such that a




Fig. 4: (a) is the input OM image and (e) is the corresponding SEM image. (b) is the result of default implementation of the pipeline. (c) is the result after adding Spectral Normalization. (d) is the best result obtained after adding all the proposed improvements.

maximum of two max-pooling operations are performed, i.e the input image is downscaled by 4. The last layer is replaced by a 2D convolution which converts the input feature maps to a 4-dimensional output $- [p, x, y, r]$, where p is the confidence map and x, y , and r are the particle coordinates and radius respectively. Since the size of input images is very large (4752×3168), we save the images as $4\times$ and $8\times$ down-scaled images (1188×792 , 594×396).

Loss Functions At each pixel level in the feature map, the model is trained to assign high probabilities if the center of a particle exists at that location and low probabilities for any other point in p . The probability mask is applied on output to suppress predictions lesser than a predefined threshold (p_{thresh}), given as $p_{mask} = p > p_{thresh}$. Similarly, for x, y , and r , the model regresses their respective values at every location in the output. We use L2 distance to compute loss for the model and mask only those regions where the radius values exist. For the i^{th} output, the predicted coordinates are l and ground-truth coordinates are g , then the objective is given as,

$$\begin{aligned} x_{ierr} &= \mathbb{E}_{x_i} [MSE(l_x, g_x) \odot p_{mask}] \\ y_{ierr} &= \mathbb{E}_{y_i} [MSE(l_y, g_y) \odot p_{mask}] \\ r_{ierr} &= \mathbb{E}_{r_i} [MSE(l_r, g_r) \odot p_{mask}] \\ p_{ierr} &= \mathbb{E}_{p_i} [\lambda MSE(l_p, g_p) \odot p_{mask}], \end{aligned} \quad (4)$$

where λ is a tunable weight parameter. The final loss is given as,

$$L_{total} = \frac{x_{ierr} + y_{ierr} + r_{ierr} + p_{ierr}}{4} \quad (5)$$

4 GAN IMPROVEMENTS

The default implementation (Pix2Pix with the ESRGAN) generated particles with boundaries blended into other particles in the cluster. As a result, the particle sizing algorithm found it difficult to regress circle coordinates, giving incorrect size estimations. Hence, we added the state-of-the-art improvements in CNNs and GANs to our models for better-looking images with sharp and clear boundaries. The changes can be tracked in Fig. 4.

4.1 Spectral Normalization

We use Spectral Normalization (SN) [33] on all the down-sampling 2D convolution layers of our Pix2Pix as well as ESRGAN models. This helps us bound the latent space generation of the deep feature maps. In the original implementation of ESRGAN, batch normalization was eliminated due to less generalizability and the presence of unpleasant artifacts. But

without any form of normalization, our experiments often resulted in either mode-collapse or overfitting. The ESRGAN training stability improved by the use of SN over the generator as well as the discriminator.

Even with large variations in the input images, the variation in the distribution of generated images remains low due to the addition of SN. In all the experiments done using our model, observed convergence of the losses after incorporating SN. Spectral normalization only requires one hyperparameter tuning, the Lipschitz constant, which works well at its default value of 1.

A function f is said to be Lipschitz continuous if, for a given metric (L2 Loss), it satisfies the following condition:

$$\|f(y) - f(x)\| \leq K \|y - x\| \quad (6)$$

If K is minimum, then it is the Lipschitz constant of a function, and Lipschitz continuity limits the rate of change of the function with respect to K . The weights in the convolutional layers can be bound by a Lipschitz constant which is controlled by applying spectral norm on each layer.

In [33] is shown that for a weight matrix W , the spectral norm is given by the square root of the largest eigenvalue of $W^T W$. The *Power Iteration* method is used to compute the spectral norm of W , given by $\sigma(W)$, efficiently. Spectral Normalization of a layer then becomes simply replacing W by $W/\sigma(W)$.

4.2 Anti-Aliasing

The down-sampling network in U-Net uses strided-convolutions in the default Pix2Pix implementation. With strided-convolutions, we observe deformations of particle boundaries for images with size distributions different than training images. Hence, we adopt anti-aliasing pooling [34] in our Pix2Pix model for the down-sampling layers in the generator as well as the discriminator. We observe increased circularity of the particles with the addition of anti-aliasing pooling, especially in dense particle clusters. Although the anti-aliasing operations require more compute than simple max-pooling operations, the performance increase makes up for it.

4.3 Self-Attention

With the final addition of Self-Attention (SA), the output from our Pix2Pix model shows uniform illumination over the entire generated image, whilst maintaining high recall of boundaries and edges of the particles. The modified generator architecture is shown in Fig. 5. The number of parameters in SA layers increase exponentially with increase in the size of the feature map. For our

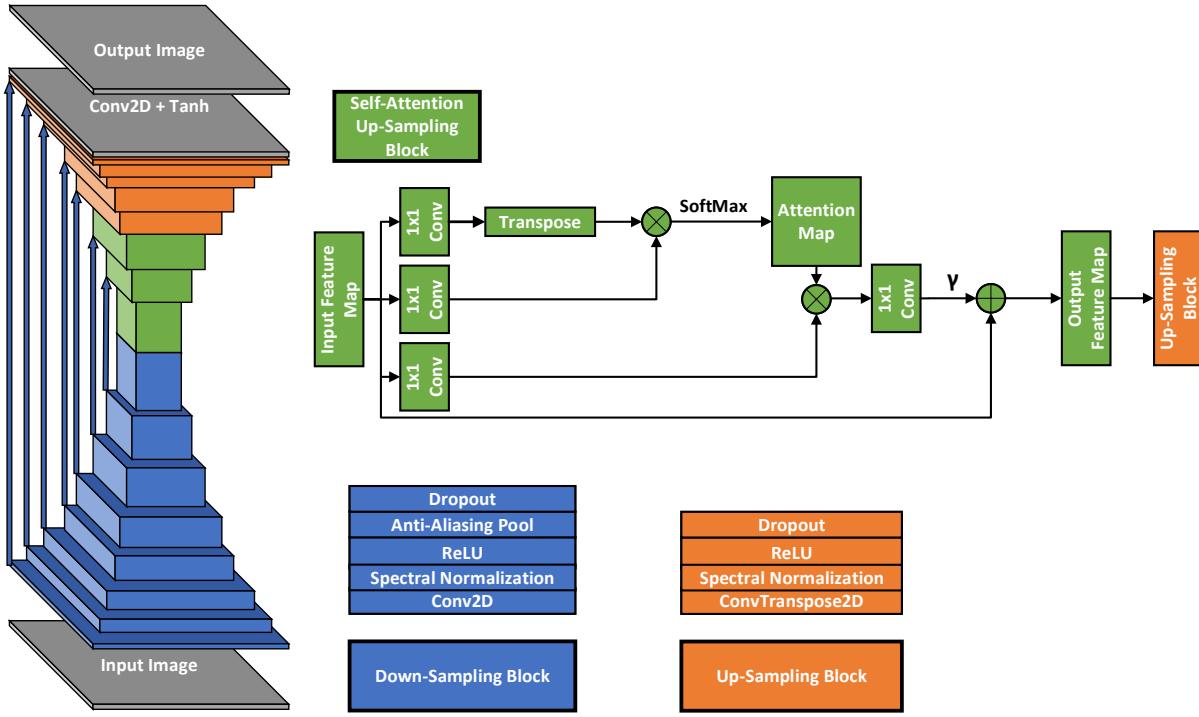


Fig. 5: Final Pix2Pix generator architecture after improvements. The self-attention block has been inspired from [35]




Fig. 6: SR-PSA Object Detection model. The CNN has two max-pooling layers. Thus, the ground truth maps are scaled down by 4. For each iteration, the image is down-sampled by 4 and 8 for the same model.

application, since the size of input images is 512×512 , we restrict SA operations over the smallest three layers in the generator.

For a hidden layer, the self-attention [40] algorithm computes three 1×1 convolutions, $f(x)$, $g(x)$ and $h(x)$. Matrix multiplication is performed on $f(x)^T$ and $g(x)$ and a SoftMax function is applied on the output. This results in generation of the attention map. The attention map is again multiplied with $h(x)$. A final 1×1 convolution operation is performed on the output, to generate $v(x)$. The output of the self-attention algorithm is then added to the original hidden layer through a multiplication factor of γ . Thus, the final output is given as,

$$y_i = \gamma o_i + x_i \quad (7)$$

γ is a learnable variable which is initialized to 0. In the initial stages of training, the network focuses on local cues. As training progresses, more weight is given to non-local features. Thus, the network learns easier tasks first, and gradually learns complex connections with distant spatial features [35]. The final Pix2Pix model with improvements is shown in Fig. 5

4.4 Training Stabilization

During GAN training, we observe that the discriminator usually gets trained fairly quickly. This results in it overpowering the generator, by accurately classifying generated images as fake, thus disabling the generator from improving. The Two Time-Scale Update Rule (TTUR) presented in [41], advocates the use

of separate learning rates for the generator and the discriminator. Implementation of TTUR in SAGAN [35] was shown to speed up the training of the discriminator.

We propose a switching TTUR technique for efficiently training the network. Initially, we set the learning rate for the generator higher than the discriminator. This induces divergence in the GAN losses after low-level features are captured. We then switch the learning rates for the generator and the discriminator and continue training while periodically decreasing both the learning rates. We also use label-smoothing to reduce "Over Confidence" of the discriminator during loss generation [42] [43].

5 IMPLEMENTATION

The following section describes the implementation details of our pipeline. We use ADAM [44] Optimizer for all the models and set the Lipschitz constant to 1 for spectral normalization.

5.1 Data

For image super-resolution, we use registered optical microscope and scanning electron microscope images for training the GANs. Since data extraction is a difficult and time-consuming process, a total of 578 original image pairs were available. Since these images were of 4K resolution, we extracted small 1024x1024 patches from the original images. These patches were resized to 512x512 resolution for computation efficiency.

We perform data augmentation to increase the size of the dataset, add more variability, and to prevent overfitting to the small sample size. For particle size analysis (PSA), once the images are super-resolved, we stitch the super-resolved images together to get the original sized image (4752x3168), which are then trained to detect the particles. Ground truth data for PSA is generated by human annotators. We compare our model output annotations against ground truth human annotations and other algorithms that have been proposed for particle size analysis using image processing.

5.2 Training

5.2.1 AASNA-Pix2Pix

For multi-modal image translation, we used 512x512 patches of OM and SEM pairs. Initially, the input images were read in RGB. But we found higher generalizability by the model for grayscale OM input images. We attribute this to the variability in illumination in the images captured by optical microscopes. For the Pix2Pix model with Anti-Aliasing, Spectral Normalization, and Self Attention (AASNA-Pix2Pix) learning rates of 0.005 and 0.001 were used for the generator and the discriminator respectively. We switch the learning rates at around half the epochs and then train the network with periodically decaying learning rates. For momentum, the parameters are set as $\beta_1 = 0.5$, $\beta_2 = 0.999$ and the RMSE loss parameter $\lambda_{RMSE} = 5$.

5.2.2 Improved-ESRGAN

For super-resolution using our ESRGAN with SN (Improved ESRGAN), input images from the AASNA-Pix2Pix model of size 512x512 were resized to 128x128. The 128x128 images were then super-resolved to get a 512x512 size image outputs using ESRGAN-SN. For training, we set learning rates to 0.003 and 0.001 for the generator and the discriminator respectively. Similar to previous training, we switch the learning rates, when

the losses start to diverge and decay them periodically after the switch. We use a set of 23 Residual-in-Residual Blocks (RRDBs) [14] in the generator. For momentum, the parameters are set as $\beta_1 = 0.9$, $\beta_2 = 0.999$.

5.2.3 Particle Size Analysis

The data was collected such that all practical modalities were covered. We divided the entire set of images into primary, secondary, and tertiary datasets. Ground truths (GT) were created by manually annotating the particles. 80% of the primary dataset was used for training and the remaining for validation and testing. The distributions of all datasets are shown below.




Fig. 7: Primary, Secondary and Tertiary data set distributions

The training data and corresponding ground-truth, before being fed to the SR-PSA network, were converted to grayscale and scaled-down by 4 and 8 to overcome scale imbalance. Feeding the network at different scales helps in extracting features of larger as well as smaller sized particles, which in turn helps in detecting particles from various size distributions. The ground-truth radii were normalized by a factor called max-radius, which is considered as a hyperparameter. The output at different scales was finally merged to get the final predictions. The objective was to optimize the Loss function given in Eq. 5. Throughout the training, we used a batch size of 4 with ADAM optimizer with learning rate, weight decay, λ , and max-radius as 0.0001, 0.000001, 5, 5 respectively.

6 EXPERIMENTS

We use separate evaluation metrics for qualitative measurement of GAN outputs and quantitative measurement of the PSA outputs.

6.1 Super Resolution

Evaluation of the quality of synthetic images is a highly domain-specific question. One of the most intuitive ways for qualitative measurement of generated images is manual GAN evaluation. Manual evaluation is suitable for the development and quick prototyping of models by changing hyperparameters. However, due to inherent bias in human supervision and large amount of time involved, it is not suitable. Hence, we rely on qualitative measures to evaluate the performance of our GANs. For all comparison metrics, we use three sets of implementations. For each implementation, the best model obtained after hyperparameter tuning was used to get the evaluation metrics.

Experiment	SSIM (primary)	SSIM (secondary)	PSNR (primary)	PSNR (secondary)
Default	0.492	0.503	28.4780 dB	28.9571 dB
SN-Pix2Pix + SN-ESRGAN	0.5	0.463	28.7491 dB	28.2375 dB
AASNA-Pix2Pix + SN-ESRGAN	0.566	0.534	29.7563 dB	29.68949 dB

TABLE 1: SSIM and PSNR scores for best models from different experiments

Mode	Primary		Secondary		Tertiary	
	Density	Coverage	Density	Coverage	Density	Coverage
Default	0.9353	0.7809	0.6966	0.6592	0.4346	0.6029
SN-Pix2Pix + SN-ESRGAN	0.9162	0.5432	0.5865	0.4851	0.4084	0.4157
AASNA-Pix2Pix + SN-ESRGAN	0.9527	0.8840	0.9325	0.8765	0.7822	0.6947

TABLE 2: Comparison b/w the three implementations on basis of Density and Coverage values for Primary, Secondary and Tertiary datasets.

Experiment	Mean	D50	Mode	StDev	CV
Ground Truths	9.34	8.89	8.60	2.33	24.96
SR-PSA	9.19	8.86	8.55	2.42	26.32
YOLOv5	9.33	8.89	8.64	2.35	25.24

TABLE 3: Comparison b/w Ground Truth, SR-PSA and YOLOv5 Predictions on Primary data set

Experiment	Mean	D50	Mode	StDev	CV
Ground Truths	6.74	6.47	6.46	1.48	22.00
SR-PSA	6.94	6.71	6.47	1.48	21.38
YOLOv5	7.06	7.11	6.77	1.46	20.70

TABLE 4: Comparison b/w Ground Truth, SR-PSA and YOLOv5 Predictions on Secondary data set

Experiment	Mean	D50	Mode	StDev	CV
Ground Truths	6.73	6.86	7.01	0.92	13.62
SR-PSA	6.83	6.88	6.97	0.65	9.49
YOLOv5	7.6	7.11	7.35	1.3	17.08

TABLE 5: Comparison b/w Ground Truth, SR-PSA and YOLOv5 Predictions on Tertiary data set

Particle Count	GT	SR-PSA (% detected)	YOLOv5 (% detected)
Primary	47166	46282 (98.13 %)	45765 (97.03 %)
Secondary	14806	14772 (99.77 %)	11689 (78.95 %)
Tertiary	14697	13397 (91.15 %)	8040 (54.71 %)
Combined Detections	-	97.10 (%)	85.42 (%)

TABLE 6: Comparison of Particle detection’s b/w SR-PSA and YOLOv5 Predictions with respect to GT (ground truths) on combined data set (primary+secondary+tertiary)

6.1.1 Default Implementation

Initially default implementations of Pix2Pix and ESRGAN were replicated individually for our use-case. While ESRGAN is powerful, it fails to perform end-to-end multi-modal super-resolution effectively. Pix2Pix model is an image translation cGAN which uses U-NET architecture for multi-modal transfer of peculiar image features. A standalone Pix2Pix model trained on OM-SEM pairs performs image translation of the particles and does not solve the problem of thick and unsharp edges in OM images. Hence, we combined Pix2Pix with ESRGAN to obtain truly super-resolved particles from the input OM images.

6.1.2 GANs with Spectral Normalization

We observed frequent non-convergence of the models in the default implementations. We attribute the non-convergence problem to the discriminator learning faster than the generator, resulting in a sub-optimal generator residing in local minima, which keeps generating identical particles without any diversity. To bring stability to the training of both the models, SN was used in their discriminators [33]. With the addition of SN and RMSE loss for

the Pix2Pix model, the circularity of particles increased, however, these changes resulted in the images looking less realistic.

6.1.3 AASNA-Pix2Pix with ESRGAN-SN

Inspired by the Self-Attention GAN (SA-GAN) paper [35], we added self-attention to our generator to capture spatially relevant information. SA-GAN implements spectral normalization in generator and discriminator and shows significant improvement in image generation capabilities of GANs. For a cGAN use-case, we implement self-attention in the up-sampling part of the network so that the generator doesn’t focus only on the local features. This enables the generated images to have consistent texture information even in the spatially distant parts.

To further bring stability to our models, we use anti-aliasing pooling instead of max-pooling. We use a normalized Binomial-5 filter which has shown to increase the consistency in CNNs [34]. Since ESRGAN only looks at the Pix2Pix output, we add self-attention and anti-aliasing to the Pix2Pix model. This enables Pix2Pix to generate high quality images consistently irrespective of the aberrations in the input, thus maintaining the quality of the final results. Hence, the bulk of the responsibility is handled by




Fig. 8: (a) L2 loss result shows significant blurring at the boundaries of the particles. (b) L1 loss result shows sharp, but non-circular boundaries. (c) RMSE loss result shows slightly thicker, but highly circular boundaries.

AASNA-Pix2Pix, so that the task becomes easier for the improved ESRGAN to accomplish. Results of experiments with different generator loss functions are shown in Fig. 8.

6.2 Evaluation

In the community, standard GAN performance is usually benchmarked using metrics like Inception Score [34], Fréchet Inception Distance (FID) [41], Structural Similarity Index Measure (SSIM) [45] and Peak Signal-to-Noise Ratio (PSNR). For standard GANs, since the real data is not exactly identical to the generated data, one-to-one metrics like SSIM and PSNR do not work well. However, for conditional GANs, if real images corresponding to the input are available, these metrics are often used to measure the quality of generated images.

6.2.1 Structural Quality Metrics

PSNR was originally used as a quality measurement tool between compressed and reconstructed images. We use PSNR to show marginal improvements in the image quality when compared to the ground truth, as our implementation got better. We also compare our implementations using SSIM. An empirical comparison between the two methods is given in [46]. We show that our AASNA-Pix2Pix + improved ESRGAN results are better than the default implementation as shown in Table 1.

6.2.2 Fidelity and Diversity Metrics

Besides structural similarity, a good metric for GAN outputs is FID. However, FID scores do not make sense for our application because the generated images are quite similar looking to the real images. Hence, singular FID scores of less than 2 are obtained, which does not prove a reliable metric to compare. Recently, precision and recall based metrics have been proposed by [47], [48]. These metrics are based on an implicit quality of the data called "manifolds". For complex high-dimensional data, that has some underlying geometric differences, the differences can be captured by its manifold in a 2D space. To evaluate our model, we use the recently developed metrics of density and coverage [49], which are an improvement over the aforementioned precision and recall.

For calculation of manifolds, representative embeddings are needed for the ground-truth and the generated images. Many evaluation metrics rely on the use of models pre-trained on ImageNet. This may affect the evaluation as the pre-trained models include a dataset bias as reported in [50], [49]. Thus, we use a randomly initialized Inception_V3 model to obtain the representative embeddings, as random embeddings are free from any kind of bias


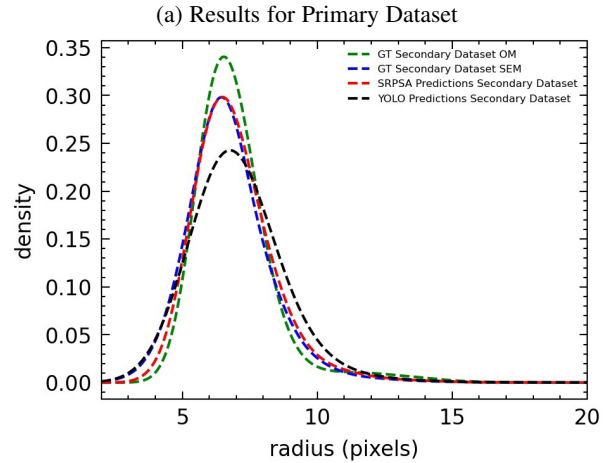
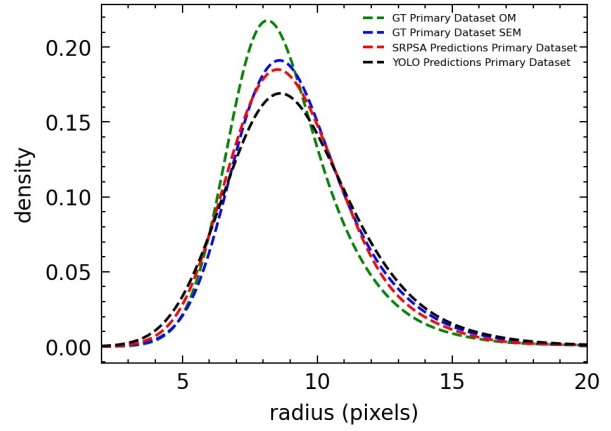


Fig. 9: Comparison between ground truths for OM and SEM datasets, and predictions by the SR-PSA and YOLOv5 models for (a) Primary Dataset, (b) Secondary Dataset, and (c) Tertiary Dataset.

and hence can provide sensible evaluation metrics as reported in [49].

To test the density and evaluation metrics for our use case, we generated two sets of synthetic images. One set of images had random backgrounds with particles intact, and the other set had washed out particles with the background intact. For the 1st set, with particles intact, we observed a high coverage value. For the

2nd set, with the background intact, we observed a high density value. Thus, for comparison between different models, we use the density and coverage values as indicators of the quality of the background and the generated particles respectively. Values for the three datasets are presented in Table 2.

We observe a decline in the results of the second implementation with just spectral normalization, as the images generated had many artifacts in the background for secondary and tertiary testing datasets. With the implementation of Anti-Aliasing and Self-Attention, we see a significant improvement across all the datasets in both, density as well as coverage.

6.3 Particle Size Analysis

The important parameters in PSA are mean, d50, mode, coefficient-of-variation (CV) and standard deviation (StDev) of the particle size distribution. YOLOv5 object detection model was trained on the same primary dataset to have a fair comparison with SR-PSA. Since the testing dataset distributions were completely different from the training, our evaluation metrics determine the generalizability of SR-PSA. To quantify the robustness of our model, we compare the kernel density estimation (KDE) plots of SR-PSA and the YOLOv5 model with the human-generated annotations ("Ground Truth") on OM and SEM images. This comparison was done for the primary, secondary and tertiary datasets.

The distributions were made in a non-parametric way using the Improved Sheather-Jones (ISJ) algorithm. ISJ estimates bandwidth by minimizing the asymptotic mean integrated square error (AMISE) [51]. From the distribution the statistics such as mean, median (D50), mode, StDev, CV were estimated. The comparison of distributions for primary, secondary and tertiary data sets are shown in Fig. 9a, Fig. 9b, and Fig. 9c respectively. As can be seen from the plots, SR-PSA predictions have consistently followed the same distribution as the manual annotations on the SEM images.




Fig. 10: (a) & (b) Predictions on Secondary Dataset, (c) & (d) Predictions on Tertiary Dataset

The comparison of statistics for primary, secondary and tertiary data sets are shown in Table 3, Table 4 and Table 5 respectively. Table 6 shows detection capabilities of SR-PSA and YOLOv5 opposed to the ground truth annotations. Along with the statistical accuracy, the amount of particles detected using SR-PSA is significantly higher than the YOLOv5 model. YOLOv5 performed well for the primary dataset but failed to work on

secondary and tertiary data sets with a different particle size distribution. Fig. 10 shows the inability of YOLOv5 to detect particles that are smaller and densely packed. On the other hand SR-PSA, consistently detects particles on distributions that are drastically different from the primary dataset.

7 CONCLUSION AND FUTURE WORK

In this paper, we propose a powerful reproducible method for super-resolution of optical microscopic particles and consequent particle size analysis. In addition, we improve on the state-of-the-art image-based cGAN methods to present the AASNA-Pix2Pix architecture for image translation. The combination of GANs in our method, which delivers high fidelity multi-modal super-resolution can be extended to other domains with similar requirements too. We demonstrate the application of our particle sizing algorithm on top of the super-resolved images that achieves accurate near-human performance for particle size analysis on circular particles.

Further investigation is still needed to understand the generalizability of the proposed pipeline across various industries. Moreover, we restricted the use of self-attention to the three smallest layers in our Pix2Pix model. More experiments are needed with other spatial attention techniques to reduce the computational complexity of the model, whilst maintaining the quality of the performance.

REFERENCES

- [1] S.-Y. Fu, X.-Q. Feng, B. Lauke, and Y.-W. Mai, "Effects of particle size, particle/matrix interface adhesion and particle loading on mechanical properties of particulate-polymer composites," *Composites Part B: Engineering*, vol. 39, pp. 933–961, Sept. 2008.
- [2] J. P. Lacey, O. Evrard, H. G. Smith, W. H. Blake, J. M. Olley, J. P. G. Minella, and P. N. Owens, "The challenges and opportunities of addressing particle size effects in sediment source fingerprinting: A review," *Earth-Science Reviews*, vol. 169, pp. 85–103, June 2017.
- [3] W. Wang, "Size and Shape Measure of Particles by Image Analysis," in *Combinatorial Image Analysis* (R. Reulke, U. Eckardt, B. Flach, U. Knauer, and K. Polthier, eds.), Lecture Notes in Computer Science, (Berlin, Heidelberg), pp. 253–262, Springer, 2006.
- [4] F. Vega, W. Lasso, and C. Torres, "Image processing applied to measurement of particle size," *Journal of Physics: Conference Series*, vol. 582, p. 012061, Jan. 2015.
- [5] H. Chen, H.-w. Tang, Y. Liu, H. Wang, and G.-p. Liu, "Measurement of particle size based on digital imaging technique," *Journal of Hydrodynamics, Ser. B*, vol. 25, pp. 242–248, Apr. 2013.
- [6] M. Smereka and I. Duřeba, "Circular Object Detection Using a Modified Hough Transform," *International Journal of Applied Mathematics and Computer Science*, vol. 18, pp. 85–91, Mar. 2008.
- [7] M. Mirzaei and H. K. Rafsanjani, "An automatic algorithm for determination of the nanoparticles from TEM images using circular hough transform," *Micron*, vol. 96, pp. 86–95, May 2017.
- [8] Z. Zhao, P. Zheng, S. Xu, and X. Wu, "Object detection with deep learning: A review," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 30, no. 11, pp. 3212–3232, 2019.
- [9] Y. Zhu, Q. Ouyang, and Y. Mao, "A deep convolutional neural network approach to single-particle recognition in cryo-electron microscopy," *BMC Bioinformatics*, vol. 18, p. 348, July 2017.
- [10] F. Wang, H. Gong, G. Liu, M. Li, C. Yan, T. Xia, X. Li, and J. Zeng, "DeepPicker: A deep learning approach for fully automated particle picking in cryo-EM," *Journal of Structural Biology*, vol. 195, pp. 325–336, Sept. 2016.
- [11] Y. Xiao and G. Yang, "A fast method for particle picking in cryo-electron micrographs based on fast r-CNN," in *AIP Conference Proceedings*, Author(s), 2017.
- [12] G. Güven and A. B. Oktay, "Nanoparticle detection from TEM images with deep learning," in *2018 26th Signal Processing and Communications Applications Conference (SIU)*, pp. 1–4, May 2018.

- [13] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, "Unpaired Image-to-Image Translation using Cycle-Consistent Adversarial Networks," *arXiv:1703.10593 [cs]*, Nov. 2018. arXiv: 1703.10593.
- [14] B. Lim, S. Son, H. Kim, S. Nah, and K. M. Lee, "Enhanced Deep Residual Networks for Single Image Super-Resolution," in *2017 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pp. 1132–1140, July 2017. ISSN: 2160-7516.
- [15] F. Farnia and A. Ozdaglar, "GANs May Have No Nash Equilibria," *arXiv:2002.09124 [cs, stat]*, Feb. 2020. arXiv: 2002.09124.
- [16] T. Karras, T. Aila, S. Laine, and J. Lehtinen, "Progressive Growing of GANs for Improved Quality, Stability, and Variation," *arXiv:1710.10196 [cs, stat]*, Feb. 2018. arXiv: 1710.10196.
- [17] Z. Wang, Q. She, and T. E. Ward, "Generative Adversarial Networks in Computer Vision: A Survey and Taxonomy," *arXiv:1906.01529 [cs]*, June 2020. arXiv: 1906.01529.
- [18] M. Mirza and S. Osindero, "Conditional generative adversarial nets," *arXiv preprint arXiv:1411.1784*, 2014.
- [19] S. Reed, Z. Akata, X. Yan, L. Logeswaran, B. Schiele, and H. Lee, "Generative Adversarial Text to Image Synthesis," *arXiv:1605.05396 [cs]*, June 2016. arXiv: 1605.05396.
- [20] D. Pathak, P. Krähenbühl, J. Donahue, T. Darrell, and A. A. Efros, "Context Encoders: Feature Learning by Inpainting," in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2536–2544, June 2016. ISSN: 1063-6919.
- [21] C. Ledig, L. Theis, F. Huszar, J. Caballero, A. Cunningham, A. Acosta, A. Aitken, A. Tejani, J. Totz, Z. Wang, and W. Shi, "Photo-Realistic Single Image Super-Resolution Using a Generative Adversarial Network," *arXiv:1609.04802 [cs, stat]*, May 2017. arXiv: 1609.04802.
- [22] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative Adversarial Nets," in *Advances in Neural Information Processing Systems 27* (Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, eds.), pp. 2672–2680, Curran Associates, Inc., 2014.
- [23] K. de Haan, Z. S. Ballard, Y. Rivenson, Y. Wu, and A. Ozcan, "Resolution enhancement in scanning electron microscopy using deep learning," *Scientific Reports*, vol. 9, p. 12050, Aug. 2019.
- [24] H. Wang, Y. Rivenson, Y. Jin, Z. Wei, R. Gao, H. Günaydin, L. A. Bentolila, C. Kural, and A. Ozcan, "Deep learning enables cross-modality super-resolution in fluorescence microscopy," *Nature Methods*, vol. 16, pp. 103–110, Jan. 2019.
- [25] J. A. Grant-Jacob, B. S. Mackay, J. A. G. Baker, Y. Xie, D. J. Heath, M. Loxham, R. W. Eason, and B. Mills, "A neural lens for super-resolution biological imaging," *Journal of Physics Communications*, vol. 3, p. 065004, June 2019.
- [26] E. Hamzeloo, M. Massinaei, and N. Mehrshad, "Estimation of particle size distribution on an industrial conveyor belt using image analysis and neural networks," *Powder Technology*, vol. 261, pp. 185–190, July 2014.
- [27] M. Frei and F. E. Kruis, "Image-based size analysis of agglomerated and partially sintered particles via convolutional neural networks," *Powder Technology*, vol. 360, pp. 324–336, Jan. 2020.
- [28] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You Only Look Once: Unified, Real-Time Object Detection," in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 779–788, June 2016. ISSN: 1063-6919.
- [29] T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollár, "Focal Loss for Dense Object Detection," in *2017 IEEE International Conference on Computer Vision (ICCV)*, pp. 2999–3007, Oct. 2017. ISSN: 2380-7504.
- [30] W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. Reed, C.-Y. Fu, and A. C. Berg, "SSD: Single Shot MultiBox Detector," in *Computer Vision – ECCV 2016* (B. Leibe, J. Matas, N. Sebe, and M. Welling, eds.), Lecture Notes in Computer Science, (Cham), pp. 21–37, Springer International Publishing, 2016.
- [31] D. B. Sam, S. V. Peri, M. N. Sundararaman, A. Kamath, and R. V. Babu, "Locate, Size and Count: Accurately Resolving People in Dense Crowds via Detection," *arXiv:1906.07538 [cs]*, Feb. 2020. arXiv: 1906.07538.
- [32] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros, "Image-to-Image Translation with Conditional Adversarial Networks," *arXiv:1611.07004 [cs]*, Nov. 2018. arXiv: 1611.07004.
- [33] T. Miyato, T. Kataoka, M. Koyama, and Y. Yoshida, "Spectral Normalization for Generative Adversarial Networks," *arXiv:1802.05957 [cs, stat]*, Feb. 2018. arXiv: 1802.05957.
- [34] R. Zhang, "Making Convolutional Networks Shift-Invariant Again," *arXiv:1904.11486 [cs]*, June 2019. arXiv: 1904.11486.
- [35] H. Zhang, I. Goodfellow, D. Metaxas, and A. Odena, "Self-Attention Generative Adversarial Networks," *arXiv:1805.08318 [cs, stat]*, June 2019. arXiv: 1805.08318.
- [36] C. Li and M. Wand, "Precomputed Real-Time Texture Synthesis with Markovian Generative Adversarial Networks," in *Computer Vision – ECCV 2016* (B. Leibe, J. Matas, N. Sebe, and M. Welling, eds.), Lecture Notes in Computer Science, (Cham), pp. 702–716, Springer International Publishing, 2016.
- [37] W. Lotter, G. Kreiman, and D. Cox, "Deep Predictive Coding Networks for Video Prediction and Unsupervised Learning," *arXiv:1605.08104 [cs, q-bio]*, Feb. 2017. arXiv: 1605.08104.
- [38] A. Jolicœur-Martineau, "The relativistic discriminator: a key element missing from standard GAN," *arXiv:1807.00734 [cs, stat]*, Sept. 2018. arXiv: 1807.00734.
- [39] K. Simonyan and A. Zisserman, "Very Deep Convolutional Networks for Large-Scale Image Recognition," *arXiv:1409.1556 [cs]*, Apr. 2015. arXiv: 1409.1556.
- [40] X. Wang, R. Girshick, A. Gupta, and K. He, "Non-local Neural Networks," *arXiv:1711.07971 [cs]*, Apr. 2018. arXiv: 1711.07971.
- [41] M. Heusel, H. Ramsauer, T. Unterthiner, B. Nessler, and S. Hochreiter, "GANs Trained by a Two Time-Scale Update Rule Converge to a Local Nash Equilibrium," *arXiv:1706.08500 [cs, stat]*, Jan. 2018. arXiv: 1706.08500.
- [42] T. Salimans, I. Goodfellow, W. Zaremba, V. Cheung, A. Radford, and X. Chen, "Improved Techniques for Training GANs," *arXiv:1606.03498 [cs]*, June 2016. arXiv: 1606.03498.
- [43] G. Pereyra, G. Tucker, J. Chorowski, L. Kaiser, and G. Hinton, "Regularizing Neural Networks by Penalizing Confident Output Distributions," *arXiv:1701.06548 [cs]*, Jan. 2017. arXiv: 1701.06548.
- [44] D. P. Kingma and J. Ba, "Adam: A Method for Stochastic Optimization," *arXiv:1412.6980 [cs]*, Jan. 2017. arXiv: 1412.6980.
- [45] Z. Wang, A. Bovik, H. Sheikh, and E. Simoncelli, "Image quality assessment: from error visibility to structural similarity," *IEEE Transactions on Image Processing*, vol. 13, pp. 600–612, Apr. 2004.
- [46] A. Horé and D. Ziou, "Image Quality Metrics: PSNR vs. SSIM," in *2010 20th International Conference on Pattern Recognition*, pp. 2366–2369, Aug. 2010. ISSN: 1051-4651.
- [47] M. S. M. Sajjadi, O. Bachem, M. Lucic, O. Bousquet, and S. Gelly, "Assessing Generative Models via Precision and Recall," *arXiv:1806.00035 [cs, stat]*, Oct. 2018. arXiv: 1806.00035.
- [48] T. Kynkäänniemi, T. Karras, S. Laine, J. Lehtinen, and T. Aila, "Improved Precision and Recall Metric for Assessing Generative Models," in *Advances in Neural Information Processing Systems 32* (H. Wallach, H. Larochelle, A. Beygelzimer, F. d. Alché-Buc, E. Fox, and R. Garnett, eds.), pp. 3927–3936, Curran Associates, Inc., 2019.
- [49] M. F. Naeem, S. J. Oh, Y. Uh, Y. Choi, and J. Yoo, "Reliable Fidelity and Diversity Metrics for Generative Models," *arXiv:2002.09797 [cs, stat]*, June 2020. arXiv: 2002.09797.
- [50] R. Geirhos, P. Rubisch, C. Michaelis, M. Bethge, F. A. Wichmann, and W. Brendel, "ImageNet-trained CNNs are biased towards texture; increasing shape bias improves accuracy and robustness," *arXiv:1811.12231 [cs, q-bio, stat]*, Jan. 2019. arXiv: 1811.12231.
- [51] Z. I. Botev, J. F. Grotowski, and D. P. Kroese, "Kernel density estimation via diffusion," *Annals of Statistics*, vol. 38, pp. 2916–2957, Oct. 2010.

Sarvesh Patil is a data scientist at HyperWorks Imaging. He received his Bachelor's Degree in Electronics and Telecommunication from the VES Institute of Technology, Mumbai. His research interests lie in computer vision, with a focus on object detection, generative adversarial networks, and scene understanding.





Chava Y P D Phani Rajanish received his B.Tech in Aerospace Engineering from UPES, M.Tech in Thermal Engineering from KL University, and is pursuing his Masters in Data Science from BITS Pilani. He is currently working at Hyperworks Imaging Pvt Ltd, Bangalore, India as a Data Scientist with a focus on Computer Vision and Data Engineering.



Naveen Margankunte received his PhD in physics from the University of Florida (UF). He is currently CEO of HyperWorks Imaging enabling creation and deployment of bespoke AI and ML solutions for various domains including advanced R&D, Agtech, and Retail.