

ADASSM: Adversarial Data Augmentation in Statistical Shape Models From Images

Mokshagna Sai Teja Karanam^{1,2}, Tushar Kataria^{1,2}, and Shireen Elhabian^{1,2}

¹ Kahlert School of Computing, University Of Utah

² Scientific Computing and Imaging Institute, University of Utah
{mkaranam, tushar.kataria, shireen}@sci.utah.edu

Abstract. Statistical shape models (SSM) have been well-established as an excellent tool for identifying variations in the morphology of anatomy across the underlying population. Shape models use consistent shape representation across all the samples in a given cohort, which helps to compare shapes and identify the variations that can detect pathologies and help in formulating treatment plans. In medical imaging, computing these shape representations from CT/MRI scans requires time-intensive preprocessing operations, including but not limited to anatomy segmentation annotations, registration, and texture denoising. Deep learning models have demonstrated exceptional capabilities in learning shape representations directly from volumetric images, giving rise to highly effective and efficient Image-to-SSM. Nevertheless, these models are data-hungry and due to the limited availability of medical data, deep learning models tend to overfit. Offline data augmentation techniques, that use kernel density estimation based (KDE) methods for generating shape-augmented samples, have successfully aided Image-to-SSM networks in achieving comparable accuracy to traditional SSM methods. However, these augmentation methods focus on shape augmentation, whereas deep learning models exhibit image-based texture bias results in sub-optimal models. This paper introduces a novel strategy for on-the-fly data augmentation for the Image-to-SSM framework by leveraging data-dependent noise generation or texture augmentation. The proposed framework is trained as an adversary to the Image-to-SSM network, augmenting diverse and challenging noisy samples. Our approach achieves improved accuracy by encouraging the model to focus on the underlying geometry rather than relying solely on pixel values.

Keywords: Statistical Shape Model · Data Augmentation · Adversarial Training

1 Introduction

Statistical shape modeling (SSM) is widely used in the fields of medical image analysis and biological sciences for studying anatomical structures and conducting morphological analysis. It enables shape analysis by facilitating the understanding of the geometrical properties of shapes that are statistically consistent

across a population. SSM has diverse applications in neuroscience [27,17], cardiology [8], orthopedics [18,5], and radiology [9,15].

Optimization-based SSM [10] methods typically involve anatomy segmentation, data preprocessing (e.g., image resampling, denoising, rigid registration), and optimizing population-level shape representation i.e., correspondence points (or particles), all of which require substantial expertise-driven workflow, involving intensive preprocessing that can be time-consuming. Deep learning approaches for SSM train networks to learn the functional mapping from unsegmented images to statistical representations of anatomical structures [8,6,2,4,3,22]. This shift towards deep learning-based methods offers a more efficient and automated approach to SSM, bypassing the need for extensive manual preprocessing and leveraging the power of neural networks to learn directly from raw imagery data [4,23]. However, deep learning models are notorious for requiring enormous quantities of data to achieve acceptable performance [21], necessitating the use of data augmentation to supplement the available training data [8].

In the field of deep learning for medical image analysis, data augmentation plays a crucial role [1,11,20]. Nevertheless, unlike computer vision applications, acquiring a substantial number of segmented medical images is difficult due to privacy concerns, the substantial human effort and expertise required, and the intensive preprocessing involved [14]. Off-the-shelf data augmentation methods may not generate augmented samples that promote invariances and improve the task-specific generalizability of the model [16]. Therefore, having a large amount of data supplemented with challenging task-specific variations would be extremely beneficial for training deep neural networks and improving model performance. In the field of medical imaging, attempts have been made to employ task-driven automatic data augmentation techniques [12] for tasks such as image segmentation and classification [24,14,11]. For regression tasks, various strategies have been proposed for handling data augmentation that includes, data-dependent shape augmentation for Image-to-SSM networks [7,6,2] and a mix-up [26] based augmentation by interpolating input samples using the similarity of labels [25].

DeepSSM [6,7], and other variants [2,3,4], learn to map unsegmented images to shape models, exhibiting comparable performance to traditional SSM [10] methods, as well as in downstream tasks such as atrial fibrillation recurrence [8,18]. DeepSSM relies heavily on offline shape-based data augmentation via kernel density estimation (KDE) in PCA subspace [7,6]. This shape augmentation approach entails using generative modeling to sample augmented shapes from probability distribution estimated via KDE. These offline methods have three deficiencies: (1) the generation of augmented samples is independent of the task (shape modeling) at hand, (2) the augmentation process focuses on shape augmentation rather than incorporating noise/texture augmentation, neglecting the inherent texture bias often present in deep learning models [19], which can lead to sub-optimal models in shape analysis, and (3) It requires extensive offline data compilation, which is time- and resource-consuming.

We draw inspiration from adversarial domain adaptation [13] and adversarial data augmentation for classification tasks [14]. Here, we adapted these ideas to regression tasks with application to Image-to-SSM networks. The regression task poses a greater challenge as it is not straightforward to generate challenging adversarial samples for the learning task at hand due to the absence of label-separating hyperplanes. Consequently, we focus our methodology on generating adversarial noise as augmentation. The proposed method results in noise-augmented samples as an alternative to KDE based shape augmentation [7,6] and implicitly drives the model to attend to the shape of the underlying object of interest instead of explicit shape augmentation [7,6]. We also demonstrate that data- and task-dependent noise augmentation is better than off-the-shelf noise augmentation with varied variance levels. The proposed augmentation approach is generic enough to be used for any Image-to-SSM network, but here we focus on DeepSSM [7] to showcase the efficacy of the on-the-fly noise augmentation vs offline shape and noise augmentation.

As we focus on image noise augmentation for this work, the shape representation of the augmented images should not be affected. As a result, we employ a contrastive loss to inform the deep learning model that noisy and their corresponding original images should be projected to the same latent representation. This contrastive loss acts as a regularizer to both the augmentation framework and the Image-to-SSM network.

The contributions of this paper are as follows:-

- A computationally efficient, automated, on-the-fly adversarial data augmentation method for regression tasks with better generalizability.
- A contrastive loss based regularization that enables enhanced noise generation that is more task- and data-dependent.
- Image-to-SSM and downstream results on left atrium and femur datasets show the efficacy of the proposed approach.

2 Methodology

This section explains the details of the proposed method and the regularization losses. The block diagram for the proposed approach is shown in Figure 1.

2.1 Adversarial Data Augmentation Block

The proposed framework for augmentation aims to enhance the performance of the Image-to-SSM network by generating data-dependent noise. This architecture can be especially useful in the context of SSM tasks, where the input data size is typically limited. Due to the paucity of data samples, deep learning models may suffer from overfitting, leading to poor generalization performance on unseen data. To address this, we integrate an adversarial [13,14] data augmentation approach for regression tasks with applications in shape modeling. The generator produces adversarial data samples that are difficult for the shape model to project into statistical shape representation.

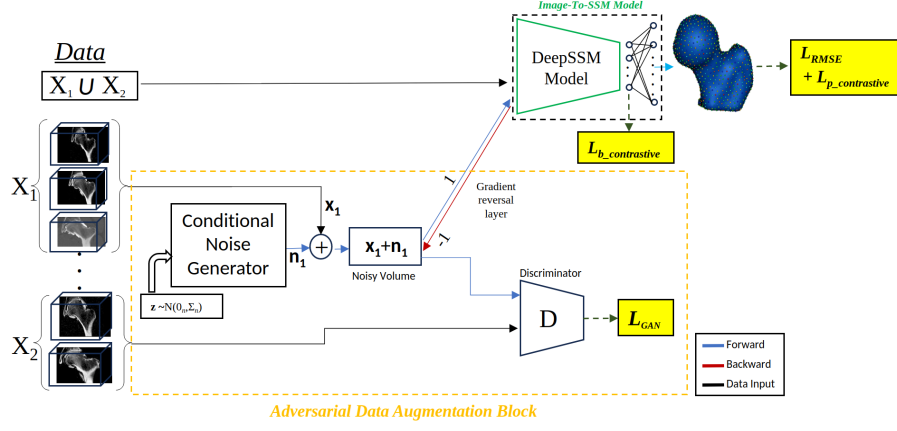


Fig. 1. ADASSM Block Diagram. Conditional Noise Generator uses the input volumes (\mathbf{x}_1) to generate noise (\mathbf{n}_1), which is added back to the input volume. Discriminator used the remaining volumes (\mathbf{X}_2) for GAN loss (\mathbf{L}_{GAN}) to ensure the generated data is in distribution. The noise generator and DeepSSM are connected by a gradient reversal layer, which sets up the second adversarial training paradigm. Both the noisy and original volumes are used to train the Image-to-SSM framework (DeepSSM). \mathbf{L}_{RMSE} is the correspondences RMSE loss, $\mathbf{L}_{p_contrastive}$ & $\mathbf{L}_{b_contrastive}$ are correspondence and bottleneck contrastive loss.

Conditional Noise Generator: Conditional noise generator receives an Image(\mathbf{x}_1) from input set X_1 and a randomly sampled Gaussian vector, \mathbf{z} as inputs. It then generates noise \mathbf{n}_1 , which is subsequently added to the original input volume $\hat{\mathbf{x}}_1 = \mathbf{x}_1 + \mathbf{n}_1$, resulting in the generation of a noisy augmented sample.

$$\hat{\mathbf{x}}_1 = G(\mathbf{z}, \mathbf{x}_1) \oplus \mathbf{x}_1 \quad (1)$$

Here, G denotes the conditional generator and \oplus represents voxelwise addition. To maintain controlled augmentations and minimize excessive perturbations, a regularization loss based on total variation (TV) is incorporated [14]. TV loss also enables the generation of noise variations that exhibit smooth transitions, which is crucial for capturing the inherent features of real-world images.

$$\mathbf{L}_{TV} = \|G(\mathbf{z}, \mathbf{x}_1)\|_2 \quad (2)$$

To further regularize the noise generation a discriminator is used. This discriminator uses the remaining input samples from set \mathbf{X}_2 , as reference distribution, and noisy samples are treated as samples of input distributions. The objective of the discriminator is to assist the generator in producing realistic noise while ensuring that the noisy augmented sample comes from the same distribution as the original data. The generative adversarial network (GAN) aims to strike a balance between meaningful data-specific augmentations and excessive pertur-

bations. The loss of the block is given in the equation below:-

$$\mathbf{L}_{\text{GAN}} = \min_G \max_D \mathbf{E}_{X \in \Omega} [\log D(\mathbf{x}_2)] + \mathbf{E}_{X \in \Omega} [\log(1 - D(G(\mathbf{z}, \mathbf{x}_1) \oplus \mathbf{x}_1))] + \beta \mathbf{L}_{TV} \quad (3)$$

where β is hyperparameter.

2.2 Adversary to Image-To-SSM network

The noisy augmented sample is fed into the Image-To-SSM network to obtain the predicted shape representation($\hat{\mathbf{y}}$), which is compared to the original shape representation via an RMSE loss.

$$\mathbf{L}(\hat{\mathbf{y}}, \mathbf{y}) = \text{RMSE}(\mathbf{y}, \text{DeepSSM}(\hat{\mathbf{x}}_1)) \quad (4)$$

The Image-to-SSM network and GAN framework are put in an adversarial relationship with the help of a gradient reversal layer, as shown in the block diagram above. The following objective function aims to minimize the Image-To-SSM network error while maximizing the conditional generator G perturbations, setting up a second adversarial objective:

$$\mathbf{L}_{\text{RMSE}} = \mathbf{E}_{X, Y \in \Omega} [\min_M \max_G \mathbf{L}(\hat{\mathbf{y}}, \mathbf{y})] \quad (5)$$

The framework above allows the augmentation model to search along the adversarial direction, leading to the generation of challenging noise augmentations that facilitate the learning of more robust shape features.

2.3 Image-To-SSM Network

The DeepSSM Model employs a deterministic encoder and a deterministic linear decoder. The reconstructed correspondences are obtained as the output of the Image-to-SSM network, providing both shape representation as well as low dimensional latent features for each input volume.

2.4 Shape Regularization Loss

Noise augmentation affects only the texture of the input volume and does not affect the underlying shape. We hypothesize that the shape representation (after DeepSSM) of the Augmented noisy volume ($\hat{\mathbf{x}}_1$) and original volume (\mathbf{x}_1), should be closer in the shape space. We use a contrastive loss as an additional regularizer to ensure that both the Image-to-SSM network and GAN account for this.

$$\mathbf{L}_{\text{contrastive}} = -\log\left(\frac{\exp(\text{sim}(\text{DeepSSM}(\mathbf{x}_1), \text{DeepSSM}(\hat{\mathbf{x}}_1)))}{\sum^N \exp(\text{sim}(\text{DeepSSM}(\mathbf{x}_1), \text{DeepSSM}(\hat{\mathbf{x}}_1)))}\right) \quad (6)$$

We propose to use contrastive loss at two different latent representations as shown in Figure 1: 1) *Correspondences* ($\mathbf{L}_{\text{p_contrastive}}$, PC), which are the

predicted SSM representation by the Image-to-SSM network. This loss ensures that the augmentation does not effect the final statistical shape representation of the augmented image. 2) *Bottleneck* ($\mathbf{L}_{b_contrastive}$, BC), which is the low-dimensional space representation obtained from the Image-to-SSM network. This loss helps the model learn the same latent representation despite noise augmentation. These regularization losses encourage both the generator and shape model to focus more on shape-related information during the learning process and factor out texture variations.

The overall objective function for the proposed model :

$$\mathbf{L} = \alpha \mathbf{L}_{GAN} + \mathbf{L}_{RMSE} + \lambda_1 \mathbf{L}_{b_contrastive} + \lambda_2 \mathbf{L}_{p_contrastive} \quad (7)$$

where $\alpha, \lambda_1, \lambda_2$ are hyperparameters.

3 Results

We use the same Image-to-SSM (DeepSSM [6]) model architecture across all experiments to ensure that variations in model performance can only be attributed to different augmentation techniques. As a baseline for comparison, we train an Image-to-SSM architecture without any augmentations (*NoAug*). Additionally, we train another model using KDE [6] augmentation (*KDE [6]*). We also compute 2 other baselines with off-the-shelf Gaussian noise augmentation with different variances ($\sigma = 1$ and 10).

3.1 Metrics

Root Mean Squared Error (RMSE): To measure the error, we calculate the average relative mean squared error (RMSE) between the predicted 3D correspondences and this is achieved by computing the RMSE for the x, y, and z coordinates and averaging them as shown in (7)

$$RMSE = \frac{1}{3}(RMSE_x + RMSE_y + RMSE_z) \quad (8)$$

For N 3D correspondences, $RMSE_x = \sqrt{\frac{\|C_x - C'_x\|_2^2}{N}}$. The same calculation is applied to $RMSE_y$ and $RMSE_z$ for the respective coordinates. Additionally, we calculate the RMSE error for each correspondence point as, $RMSE_i = \sqrt{\frac{\|C_x^i - C_x'^i\|_2^2 + \|C_y^i - C_y'^i\|_2^2 + \|C_z^i - C_z'^i\|_2^2}{3}}$. The per-point RMSE helps us assess the accuracy of DeepSSM in modeling various local anatomical features. For all experiments, the shape representations were calculated on same test data (held-out data) using the trained DeepSSM model and were only used for inference.

Surface-to-Surface Distance (mm): The surface-to-surface distance is measured between the ground truth mesh and the mesh reconstructed from the predicted correspondences by DeepSSM. This distance provides a more precise

measure of how well the correspondences adhere to the shape and indicates their suitability for anatomy segmentation.

Furthermore, we validate the effectiveness of DeepSSM with the learned shape representations by utilizing its correspondences for various downstream analysis applications. The specific downstream applications vary for each dataset and are described in separate subsections below.

3.2 Femur

Data Description and Processing: The femur dataset consists of 49 CT images of the femur bone, of which 42 are considered healthy with no morphological abnormalities. DeepSSM [7,6] requires generating point distribution models (or correspondences) for the training images. we use ShapeWorks [10] to optimize a shape model with 1024 correspondences. Along with the training and validation data, we also randomly selected 7 controls and 2 CAM-FAI scans for testing the DeepSSM Model. To meet GPU memory requirements, each image has been downsampled by a factor of 2 from $260 \times 184 \times 235$ (0.5mm isotropic voxel spacing). The resulting downsampled 3D volumes have isotropic voxel spacing of 1mm, with dimensions of $130 \times 92 \times 117$ voxels. The training images are divided into training and validation sets with 80% -20% split.

Training Specifics: Empirically we set $\alpha = 1$ and $\beta = 0.1$. The default parameters and configuration are used to get the result for KDE [6] augmentation for Femur dataset. The training process involves optimizing the loss on correspondences for 1500 epochs, employing a data augmentation framework based on validation loss. For all ADASSM experiments, involving the proposed regularization losses, a learning rate of $5e-5$ is utilized while ADASSM itself is trained with a learning rate of $1e-5$. The generator and discriminator learning rates are set to $5e-3$, except for ADASSM+BC+PC, which employs a learning rate of $1e-3$. A batch size of 4 is used for training all the proposed models and baseline models. In the generator, the noise perturbation range is set to 500 for ADASSM experiments.

Evaluation and Analysis: Figure 2 visualizes the RMSE results alongside the surface-to-surface distance. Augmenting the DeepSSM model with Gaussian noise with different variances without any adversarial training improves the RMSE when compared to the KDE augmentation[6]. However, an increase in the surface distance of the predicted correspondences indicates misalignment with the ground truth femur bone segmentation. This result shows that standard noise augmentation can provide better results for Image-to-SSM networks.

We can observe that for the proposed ADASSM (data-dependent noise augmentation), RMSE results are better compared to both KDE [6] and Gaussian noise. The addition of contrastive regularization losses further improves RMSE error and surface distance, which proves that data-dependent noise-augmented samples are better compared to Gaussian augmented samples and shape augmentation [6].

Visualizations of surface-to-surface distance for the best, median, and worst cases for the test set are shown in Figure 3. Upon careful examination of these




Fig. 2. Femur Test Results. The *RMSE* is the Euclidean distance between the ground truth and predicted correspondences for all the test samples. The Surface-to-Surface distance (mm) is computed by comparing the reconstructed mesh using the ground truth correspondences and the predicted correspondences for all the test samples. λ_1 for ADASSM+BC is 0.5, λ_2 for ADASSM+PC is 0.1 and λ_1, λ_2 for ADASSM+BC+PC is 0.5. Y-axis is the magnitude of the errors displayed.

visualizations, we can observe that the proposed models demonstrate a remarkable reduction in errors in critical regions such as the greater trochanter, growth plate, femoral neck, and epiphyseal lines in the best-case scenario. In the median case, a detailed analysis of both views reveals that in View (1), the KDE baseline [6] exhibits some errors around the trochanter region that are substantially reduced by the proposed models. In View (2), the error around the trochanter region is completely reduced with the ADASSM+BC+PC model. In the worst-case scenario, View (2) displays the majority of errors in the KDE baseline [6], but these errors are significantly diminished, particularly in the lower trochanter region, with the employment of the ADASSM variants.

Downstream Task - Group Differences: To evaluate the effectiveness of the learned shape representations using the proposed models, we conducted a downstream analysis aimed at evaluating whether the models can accurately capture group differences [18] in medically relevant regions. To achieve this, we formed two groups: a control group and a pathology group consisting of CAM-FAI cases. We calculated the mean differences (μ_{normal} and $\mu_{cam-FAI}$) between these groups and visualized the differences on a mesh and compared the group differences obtained using the predicted correspondences from the proposed models and the ShapeWorks PDM model. We utilized the entire dataset, including both training and testing samples, for these group differences, and the results are presented in Figure 4.

Each group difference illustrates the transition from the mean shape of the pathological group to that of the control group, overlaid on the mean patho-

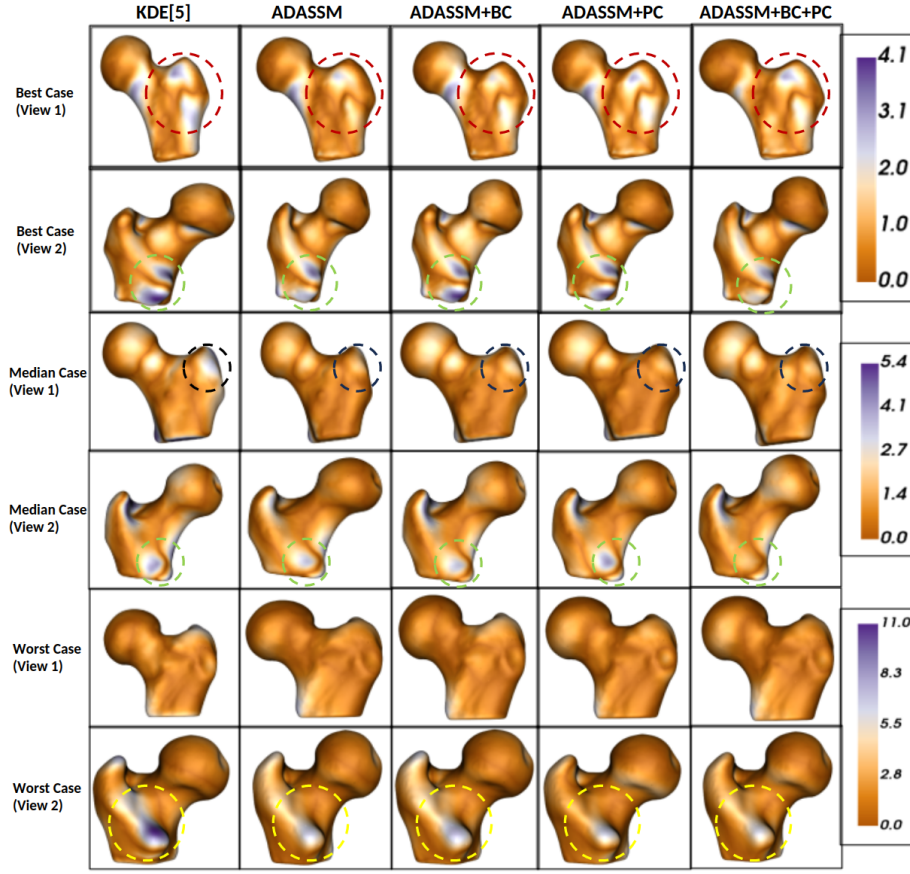


Fig. 3. Femur Surface-to-Surface distance (mm) visualization: Shape reconstruction error is displayed as a heatmap on the ground truth reconstructed meshes. The results for different proposed models are presented, showcasing their performance in the best, median, and worst-case scenarios.

logical scan. Interestingly, we observed that the group differences between the state-of-the-art PDM model and ADASSM+PC were quite similar compared to baseline KDE [6]. In some cases, the proposed models exhibited similar differences in medically relevant regions of the femur, whereas in other areas, the models identified additional variations. These findings suggest that the established correspondences can be employed to characterize CAM deformity effectively.

3.3 Left Atrium

Data Description and Processing: The left atrium dataset consists of 176 late gadolinium enhancement (LGE) MRI images from patients that have been diagnosed with atrial fibrillation (AF). These scans are acquired after the first




Fig. 4. Femur shape group difference between CAM-FAI and controls. The difference $\mu_{cam} - \mu_{normal}$ and is projected on μ_{cam} with both training and test samples, the arrows denote the direction of the correspondence (particle) movement, and the heatmap showcases the normalized magnitude.

ablation. 80-20% split is used to split the data where 146 volumes are used to train the model and 30 scans which are used to test the DeepSSM network. To generate point distribution models (or correspondences) for training images, we utilize ShapeWorks [10] to optimize a shape model with 1024 correspondences. For training purposes, the MRIs are downsampled from $235 \times 138 \times 175$ (0.625mm isotropic voxel spacing) to $117 \times 69 \times 87$ (1.25mm voxel spacing) by a factor of 2.

Training Specifics : The results for KDE augmentation [6] on the Left Atrium dataset are obtained using the default parameters and configuration. With the proposed data augmentation framework DeepSSM model is trained for 1000 epochs. For the various ADASSM models with the aforementioned regularization losses, a learning rate of $1e-4$ is utilized, while ADASSM itself is trained with a learning rate of $5e-3$. A batch size of 6 is employed during the training of both the proposed and baseline models. In all ADASSM experiments, the noise perturbation range in the generator is set to 100. Following the publication of the manuscript, we plan to make the training models, implementation code, and relevant hyperparameters publicly available.

Evaluation and Analysis: Figure 5 displays the RMSE results alongside the surface-to-surface distance. We can make the following observations from the bar graph:- 1) When augmenting the DeepSSM model with Gaussian noise of varying variances without any adversarial training, we find that it does not improve performance, which may be because the original dataset already has more intensity variation when compared to CT volumes. 2) By enhancing the data- and task-dependency of the noise and integrating the proposed data augmentation framework with various regularization losses, the methodology surpasses the baseline Gaussian noise and KDE shape augmentation framework [6]. In the Left Atrium, the performance disparity between the ADASSM variants is more pronounced compared to the Femur. This can be attributed to the significant




Fig. 5. Left Atrium Test Set Results. The *RMSE* is the Euclidean distance between the ground truth and predicted correspondences for all the test samples. The Surface-to-Surface distance (mm) is computed by comparing the reconstructed mesh using the ground truth correspondences and the predicted correspondences for all the test samples. λ_1 for ADASSM+BC is 0.001, λ_2 for ADASSM+PC is 0.05 and λ_1, λ_2 for ADASSM+BC+PC is 0.05. Y-axis is the magnitude of the errors displayed.

variations observed in the Left Atrium, where the proposed method excels in effectively regulating the Image-to-SSM task.

In Figure 6, visualizations of the surface-to-surface distance are presented for the best, median, and worst cases in the test set. For best-case and median-case views, the proposed model plainly outperforms other methods, while worst-case views are comparable.

Downstream Task - AF Recurrence Prediction: The shape of the left atrium can provide insights into the recurrence of atrial fibrillation (AF) [18]. The dataset has binary outcome labels indicating whether patients experienced AF recurrence after ablation. The goal is to estimate the probability of AF recurrence based on the learned shape representations. We use PCA projections of the shape representations as features for a Multi-Layer Perceptron (MLP) for classification.

The results are summarized in Table 1. Compared with the traditional SSM [10] and KDE, we observe similar performance. All ADASSM variants perform on par with the PDM, with the ADASSM+PC model outperforming the baselines by capturing better shape descriptors for the left atrium than the PCA scores learned in other models. Due to the fact that the classification model is based on the PCA scores of correspondences, ADASSM+PC has the highest accuracy, as the contrastive loss will bring the correspondence’s latent space representation closer. But if we train a classifier with non-linear features (other than PCA), ADASSM+BC+PC might result in the best accuracy.




Fig. 6. Left Atrium Surface-to-Surface distance (mm) visualization: Shape reconstruction error is displayed as a heatmap on the ground truth reconstructed meshes. The results for the best, median, and worst-case scenarios are shown for various proposed models. We can observe that proposed methods are outperforming the baseline results in best and median-case.

Table 1. AF Recurrence Prediction: Accuracy of AF recurrence that uses PCA scores as shape descriptors from different models.

Model	Accuracy
ShapeWorks [10]	53.99% \pm 6.20
KDE[6]	52.66% \pm 1.77
Gaussian(sigma=1)	50.66% \pm 6.22
Gaussian(sigma=10)	51.99% \pm 7.11
ADASSM	51.33% \pm 2.66
ADASSM+BC	55.99% \pm 15.11
ADASSM+PC	67.99% \pm 2.67
ADASSM+BC+PC	51.33% \pm 2.66

3.4 Training Time

The proposed augmentation method not only improved model performance for both the left atrium and femur datasets but also significantly reduces the training time by approximately 60% compared to the baseline method [6] as shown in Table 2.

Table 2. Resources Required: Comparison for time taken to run *Augmentation & Model* training required on Nvidia RTX 5000 system.

Dataset	KDE[5]	ADASSM
Left Atrium	690.35+122.35 minutes	336.15 minutes
Femur	725.35+ 93.5 minutes	399.6 minutes

4 Conclusion and Future Work

In this study, we introduced a novel methodology by proposing an adversarial data augmentation framework for generic regression tasks with applicability to Image-to-SSM networks. Using data-dependent noise augmentation, the proposed method seeks to discover effective shape representations for three-dimensional volumes. By generating challenging augmentations during model training, the proposed method eliminates the need for offline data augmentation, effectively training a more accurate Image-to-SSM network. The proposed noise augmentation framework outperforms the shape augmentation framework [8] and standard noise augmentation, demonstrating that data-dependent noise aids the model by implicitly attending to shape. Through downstream task analysis, we confirmed that the proposed method effectively taught models robust shape descriptors that capture pertinent pathology information. In addition, compared to earlier frameworks for shape augmentation, the proposed methodology is not only more robust but also faster. The limitation of the proposed framework is that it trains only on data-dependent intensity/noise augmentations and does not take shape augmentation into account. We plan to extend this framework to data-dependent shape augmentation as well.

References

1. Abdollahi, B., Tomita, N., Hassanpour, S.: Data augmentation in training deep learning models for medical image analysis. Deep learners and deep learner descriptors for medical applications pp. 167–180 (2020)
2. Adams, J., Bhalodia, R., Elhabian, S.: Uncertain-deepssm: From images to probabilistic shape models. In: Shape in Medical Imaging: International Workshop, ShapeMI 2020, Held in Conjunction with MICCAI 2020, Lima, Peru, October 4, 2020, Proceedings. pp. 57–72. Springer (2020)

3. Adams, J., Elhabian, S.: From images to probabilistic anatomical shapes: A deep variational bottleneck approach. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 474–484. Springer (2022)
4. Adams, J., Elhabian, S.: Fully bayesian vib-deepssm. arXiv preprint arXiv:2305.05797 (2023)
5. Bhalodia, R., Dvoracek, L.A., Ayyash, A.M., Kavan, L., Whitaker, R., Goldstein, J.A.: Quantifying the severity of metopic craniosynostosis: a pilot study application of machine learning in craniofacial surgery. *The Journal of craniofacial surgery* **31**(3), 697 (2020)
6. Bhalodia, R., Elhabian, S., Adams, J., Tao, W., Kavan, L., Whitaker, R.: Deepssm: A blueprint for image-to-shape deep learning models. arXiv preprint arXiv:2110.07152 (2021)
7. Bhalodia, R., Elhabian, S.Y., Kavan, L., Whitaker, R.T.: Deepssm: a deep learning framework for statistical shape modeling from raw images. In: Shape in Medical Imaging: International Workshop, ShapeMI 2018, Held in Conjunction with MIC-CAI 2018, Granada, Spain, September 20, 2018, Proceedings. pp. 244–257. Springer (2018)
8. Bhalodia, R., Goparaju, A., Sodergren, T., Morris, A., Kholmovski, E., Marrouche, N., Cates, J., Whitaker, R., Elhabian, S.: Deep learning for end-to-end atrial fibrillation recurrence estimation. In: 2018 Computing in Cardiology Conference (CinC). vol. 45, pp. 1–4. IEEE (2018)
9. Bharath, K., Kurtek, S., Rao, A., Baladandayuthapani, V.: Radiologic image-based statistical shape analysis of brain tumours. *Journal of the Royal Statistical Society. Series C, Applied statistics* **67**(5), 1357 (2018)
10. Cates, J., Elhabian, S., Whitaker, R.: Shapeworks: particle-based shape correspondence and visualization software. In: Statistical shape and deformation analysis, pp. 257–298. Elsevier (2017)
11. Chlap, P., Min, H., Vandenberg, N., Dowling, J., Holloway, L., Haworth, A.: A review of medical image data augmentation techniques for deep learning applications. *Journal of Medical Imaging and Radiation Oncology* **65**(5), 545–563 (2021)
12. Cubuk, E.D., Zoph, B., Mane, D., Vasudevan, V., Le, Q.V.: Autoaugment: Learning augmentation policies from data. arXiv preprint arXiv:1805.09501 (2018)
13. Ganin, Y., Ustinova, E., Ajakan, H., Germain, P., Larochelle, H., Laviolette, F., Marchand, M., Lempitsky, V.: Domain-adversarial training of neural networks. *The journal of machine learning research* **17**(1), 2096–2030 (2016)
14. Gao, Y., Tang, Z., Zhou, M., Metaxas, D.: Enabling data diversity: efficient automatic augmentation via regularized adversarial training. In: Information Processing in Medical Imaging: 27th International Conference, IPMI 2021, Virtual Event, June 28–June 30, 2021, Proceedings. pp. 85–97. Springer (2021)
15. Gardner, G., Morris, A., Higuchi, K., MacLeod, R., Cates, J.: A point-correspondence approach to describing the distribution of image features on anatomical surfaces, with application to atrial fibrillation. In: 2013 IEEE 10th International Symposium on Biomedical Imaging. pp. 226–229. IEEE (2013)
16. Geiping, J., Goldblum, M., Somepalli, G., Shwartz-Ziv, R., Goldstein, T., Wilson, A.G.: How much data are augmentations worth? an investigation into scaling laws, invariance, and implicit regularization. arXiv preprint arXiv:2210.06441 (2022)
17. Gerig, G., Styner, M., Jones, D., Weinberger, D., Lieberman, J.: Shape analysis of brain ventricles using spharm. In: Proceedings IEEE workshop on mathematical methods in biomedical image analysis (MMBIA 2001). pp. 171–178. IEEE (2001)

18. Harris, M.D., Datar, M., Whitaker, R.T., Jurrus, E.R., Peters, C.L., Anderson, A.E.: Statistical shape modeling of cam femoroacetabular impingement. *Journal of Orthopaedic Research* **31**(10), 1620–1626 (2013)
19. Hermann, K., Chen, T., Kornblith, S.: The origins and prevalence of texture bias in convolutional neural networks. *Advances in Neural Information Processing Systems* **33**, 19000–19015 (2020)
20. Hussain, Z., Gimenez, F., Yi, D., Rubin, D.: Differential data augmentation techniques for medical imaging classification tasks. In: *AMIA annual symposium proceedings*. vol. 2017, p. 979. American Medical Informatics Association (2017)
21. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. *Communications of the ACM* **60**(6), 84–90 (2017)
22. Tóthová, K., Parisot, S., Lee, M.C., Puyol-Antón, E., Koch, L.M., King, A.P., Konukoglu, E., Pollefeys, M.: Uncertainty quantification in cnn-based surface prediction using shape priors. In: *Shape in Medical Imaging: International Workshop, ShapeMI 2018, Held in Conjunction with MICCAI 2018, Granada, Spain, September 20, 2018, Proceedings*. pp. 300–310. Springer (2018)
23. Xu, H., Elhabian, S.Y.: Image2ssm: Reimagining statistical shape models from images with radial basis functions. *arXiv preprint arXiv:2305.11946* (2023)
24. Xu, J., Li, M., Zhu, Z.: Automatic data augmentation for 3d medical image segmentation. In: *Medical Image Computing and Computer Assisted Intervention—MICCAI 2020: 23rd International Conference, Lima, Peru, October 4–8, 2020, Proceedings, Part I* 23. pp. 378–387. Springer (2020)
25. Yao, H., Wang, Y., Zhang, L., Zou, J.Y., Finn, C.: C-mixup: Improving generalization in regression. *Advances in Neural Information Processing Systems* **35**, 3361–3376 (2022)
26. Zhang, H., Cisse, M., Dauphin, Y.N., Lopez-Paz, D.: mixup: Beyond empirical risk minimization. *arXiv preprint arXiv:1710.09412* (2017)
27. Zhao, Z., Taylor, W.D., Styner, M., Steffens, D.C., Krishnan, K.R.R., MacFall, J.R.: Hippocampus shape analysis and late-life depression. *PLoS One* **3**(3), e1837 (2008)