Decoding the Encoder

Aritra Ray[1], Xiang Li[1], Laura Barisoni[4,5], Krishnendu Chakrabarty[1], and Kyle Lafata[1,2,3]

[1]Department of Electrical and Computer Engineering, Duke University
[2]Department of Radiation Oncology, Duke University
[3]Department of Radiology, Duke University
[4]Department of Pathology, Duke University
[5]Department of Medicine, Division of Nephrology, Duke University

SoutheastCon 2023 | 978-1-6654-7611-9/23/$31.00 ©2023 IEEE | DOI: 10.1109/SoutheastCon51012.2023.10115100

***Abstract*—Autoencoders are used in a variety of safety-critical applications. Uncertainty quantification is a key component to bolster the trustworthiness of such models. With the growing complexity of the autoencoder design and the dataset they are trained on, there is a dwindling correlation between the input and feature space representation. To address this latent space degeneracy, we propose a novel method of monotonically perturbing the encoded latent space to increase the entropy in the learned representations for every corresponding input. For every perturbation, we obtain a unique decoded signature corresponding to an evaluation metric in the continuous domain, which can be clustered to build a knowledge base and subsequently analyzed for outlier analysis. For the test cases, in the absence of ground truth, we can perturb the latent space representation and find the closest match of the test cases' unique signatures to the existing knowledge base for uncertainty quantification and outlier detection. We evaluate our proposed novel method on glomeruli segmentation for frozen kidney donor section on whole slide imaging, a safety-critical application in digital pathology which serves as a precursor to kidney transplantation. We prove the proposed method's effectiveness for outlier detection by ranking the test cases according to their associated uncertainties to leverage the attention of medical experts on boundary cases.**

***Index Terms*—Autoencoders, Perturbation, Latent Space, Outlier, Uncertainty quantification**

I. Introduction

Autoencoders (AE) are used in various applications, ranging from classification, regression, and segmentation tasks to denoising ones. AEs are trained to map input space to an encoded latent space, which can be meaningfully decoded. However, with the growing complexities of encoder-decoder designs and the dataset they are trained on, there is little to no correlation between the input and the encoded latent space [1]. This is referred to as latent space degeneracy. Also, AEs find their way into diverse safety-critical applications like digital pathology, radiology, high-energy physics, to name a few. The trustworthiness of AEs is a bottleneck to their wide adaption. Moreover, in the absence of ground truth for test cases, the AE model does not assign an uncertainty score to the prediction results, indicating a constant model test accuracy for all test cases. This paper introduces a novel perturbation-based method for uncertainty quantification and outlier detection in test cases for AE models. The method is applicable to a wide variety of applications. We extensively evaluate our proposed approach on a digital pathology application.

The encoder is trained to learn the correlation between the input space and latent space embeddings. The decoder then reconstructs meaningful representations from the learned feature space. In our work, we propose monotonically perturbing the encoded latent space to increase the entropy in the learned representations for every corresponding input. We perturb the non-zero elements of the corresponding latent feature space matrix in the continuous domain for every input. For every perturbation corresponding to every input, we decode the response through the trained decoder, evaluate it, and generate a response curve in the continuous domain. This response curve or signature would be unique for every input. Consequently, we generate unique signatures for every input query and build a knowledge base. Next, we cluster the responses and analyze every cluster. Response curves farther away from their cluster center indicates the presence of a possible outlier. In the testing phase (queries without ground truth), we perform similar perturbations to obtain a unique signature of all the test cases. We assign the test cases' unique signature to a cluster from the existing knowledge base with closest signature matching. Depending on the proximity to its assigned cluster center, an uncertainty score can be assigned to every test case. Finally, the test cases can be ranked according to their associated model uncertainties.

In our work, we evaluate the proposed method on medical image segmentation application pertaining to digital pathology. Identifying and segmenting frozen kidney donor sections for kidney transplantation is a safety-critical healthcare application. We opt for a trained AE model that can segment the sclerotic and non-sclerotic glomeruli in frozen kidney whole slide imaging tissues to deem the tissue suited for transplantation or otherwise. Quantifying uncertainties in the segmentation decision of the AE model is desired by the medical experts. Also, ranking the slides in order of their associated uncertainties would assist the medical experts in identifying cases that require human intervention. The key contribution of our work is that we developed a novel uncertainty quantification tool that monotonically perturbs the feature space in AEs to generate unique signatures and rank the test cases in order of their uncertainty score to detect outliers. The tool can identify all sorts of outliers for the test cases (absence of ground truth) that can arise from poor staining of the whole slide imaging, noise in the imaging technology, or genuinely out-of-distribution cases. Our result shows that we can successfully detect outliers and rank test cases to let the medical experts focus on boundary cases.

Our paper is organized as follows. We delve into the related works in Section II, describe our proposed method in Section III, followed by an evaluation of the proposed method on medical image segmentation task for glomeruli segmentation in the frozen kidney donor section in Section IV. Section V concludes our paper.

II. Related Works

In brief, encoders of AE models are trained to extract and encode information in the latent feature space such that the decoder can interpret meaningful reconstructions of the same. However, with complex designs of AE models, we observe a latent space degeneracy, meaning a lack of correlation between the input and its corresponding feature space. In this section, we delve into the representation learning of AE models, latent space degeneracy of AEs, and establish the relevance of outlier detection for the trustworthiness of AE models.

A. Learning in Autoencoders

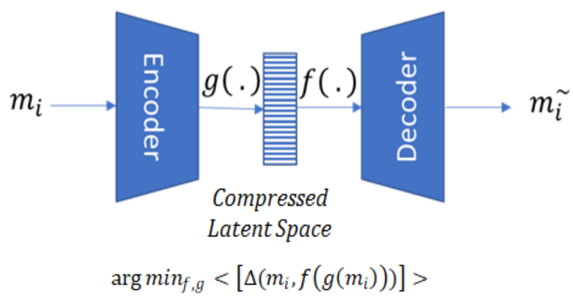

Fig. 1. Training of Autoencoders to learn meaningful representations in encoded latent space.

Autoencoders [2] consists of an encoder and a decoder block. The encoder block, through dimensionality reduction, learns meaning representations from every training sample, m_i. As illustrated in Fig. 1, the compressed latent space is obtained through the encoder's learned function $g(.)$. The encoded feature space is then decoded through the decoder's learned function $f(.)$. These functions, $g(.)$ and $f(.)$ are trained through the optimization function as given in Equation 1, where $<>$ indicates all the training samples, and Δ is the reconstruction loss between input m_i and re-constructed sample $m_i^{\sim}$.

$$argmin_{f,g} < [(\Delta(m_i, f(g(m_i)))] > \quad (1)$$

There are other varieties of AE models [3] like denoising AE which subtracts the noise from the input queries and generates a meaningful noiseless reconstruction of the query, and variational-AE (VAE), wherein the AE is regularized while training to avoid overfitting. Overall, AE models aim to learn representations that can be decoded meaningfully.

B. Latent Space Degeneracy of Autoencoders

Encoders are trained to establish a correlation from the input to latent space. The learned representations of the feature space can be thereby interpreted by the subsequent decoder network, completing the design of autoencoders. To the contrary, studies [1] show that increasing complexity of the encoder or decoder units dwindles the correlation significantly.

Addressing the degeneracy in variational AEs, H. Zheng et al. [1] proposed to analyze the transmitted information across the VAE layers while accounting for the information loss. It was observed that a deeper decoder compared to its encoder, yielded a meaningless representation of the latent space to the inputs. A deeper encoder has a poor decoder reconstruction yield compared to the encoder. Finally, with a very deep encoder and decoder, AE models fail to learn any meaningful representation. Another key finding was that skip connections help minimize information loss the best. Given this observation, we adopted an AE model with an equal yet moderately deep encoder and decoder units, with skip connections in our evaluation design. Another work by Helena et al. [4] leverages theoretical physics modeling the latent space as energy spectrums and using Hamiltonian operators to study the embedding degeneracy in AEs.

C. Outlier Detection and Trustworthiness of Autoencoders

With the growing degeneracy of latent space representations in AEs, quantifying uncertainties and identifying outliers in AE models have become increasingly challenging. Ofir et al. proposed Probabilistic Robust Autoencoders [5] aimed to split between the inlier and outlier samples through a lower dimensional representation of the inliers. It aimed to exclude outliers through a regularized design of AEs that does not conform to the low-dimensional inlier sample space. Finke et al. [6] studied the reconstruction loss of AEs for unsupervised anomaly detection. It proposed techniques for performance improvements in model-independent anomaly detection settings for high-energy physics applications. Kieu et al. [7] proposed designs of robust and explainable AE models for outlier detection in time-series models through post-hoc explainability measures. Yong et al. [8] studied the relevance of bottlenecks in AE design and proposed ways to do away with the bottleneck. It also examined that bottleneck removal techniques can help outperform the bottlenecked AEs and delved deeper into studying the effect of anomaly detection in similar settings. Other works include outlier detection using de-biasing VAE likelihoods [9] and anomaly detection in semi-supervised settings [10]. The application settings in which this has been studied remain vast, from high energy physics to classification, regression, and high-performance computing systems [11].

Anomaly detection has thus far been a long-standing problem. We, in this paper, address this from the degeneracy of latent space representations point of view, which finds no other mention in the existing literature. Also, we perform evaluations on the safety-critical application of medical image segmentation application which is a precursor to kidney transplantation.

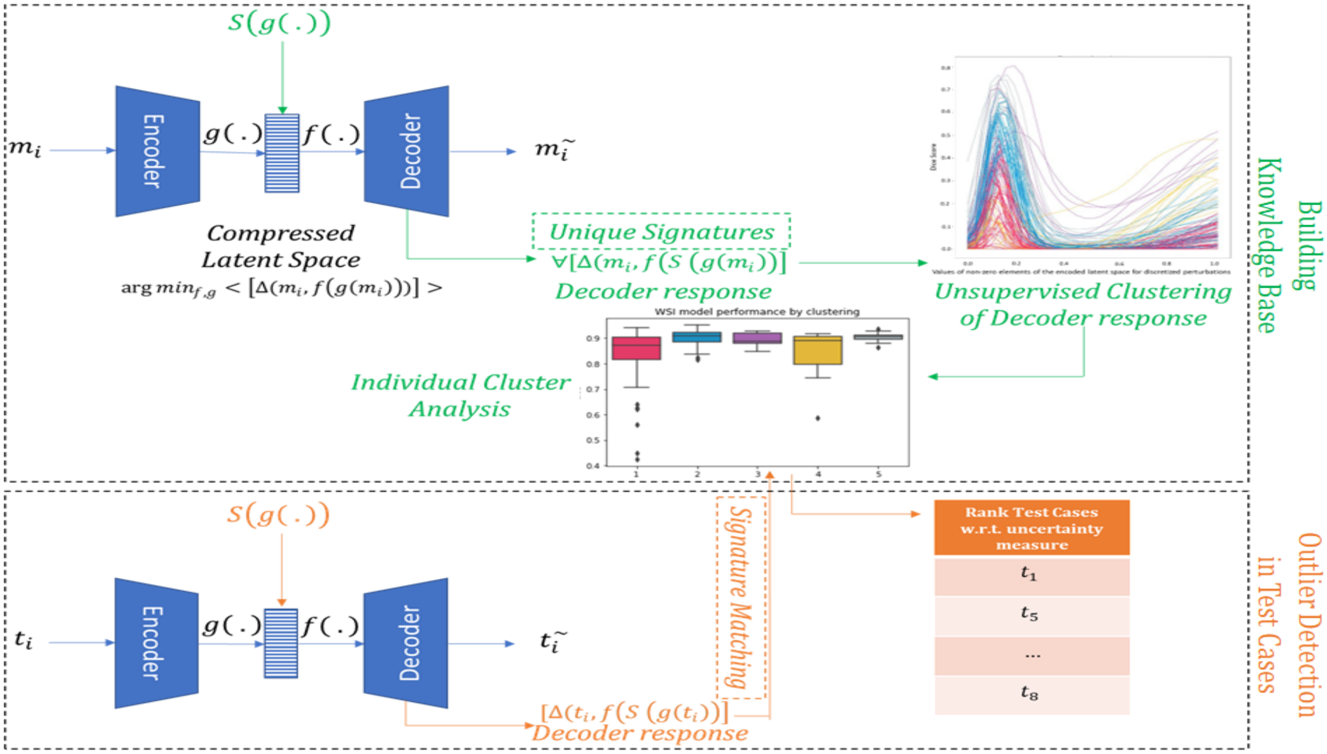

Fig. 2. This figure demonstrates our proposed method. In *Building Knowledge Base* phase, for every input, we introduce monotonic perturbations $S(g(.))$ in the latent space, cluster the decoded unique signatures, and perform individual cluster analysis. In *Outlier Detection in Test Cases* phase, for every test case (absence of ground truth), through monotonic perturbations, we obtain the decoder signature, perform closest signature matching to the existing knowledge base, quantify the uncertainties, and rank all test case in order of the measured uncertainties. This helps medical experts to address cases involving higher uncertainties with high priority.

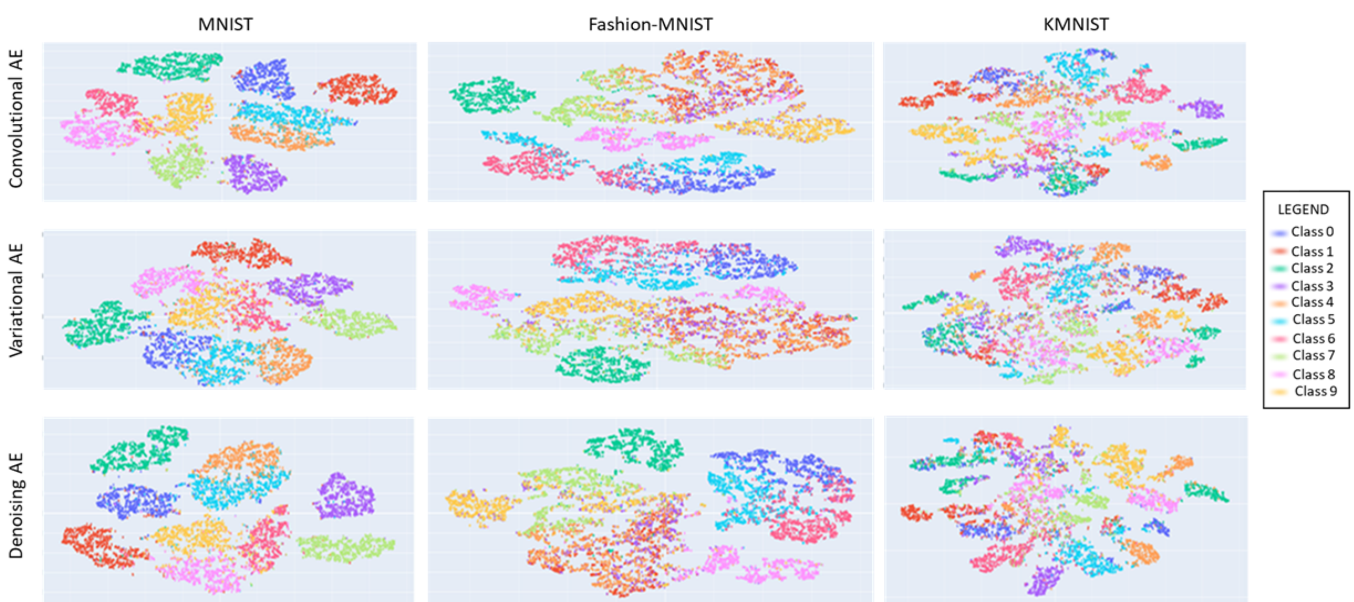

Fig. 3. This figure visualizes the latent space for increasing autoencoder model complexity coupled with dataset complexity. The first column visualizes the latent space for the MNIST dataset, the second for F-MNIST, and the third for K-MNIST. The first row uses a convolutional autoencoder, the second row uses a variational autoencoder, and the third uses a denoising autoencoder to learn the feature space for the corresponding datasets. With growing model and dataset complexity, the correlation from input to latent space degenerates, which emphasizes our problem statement.

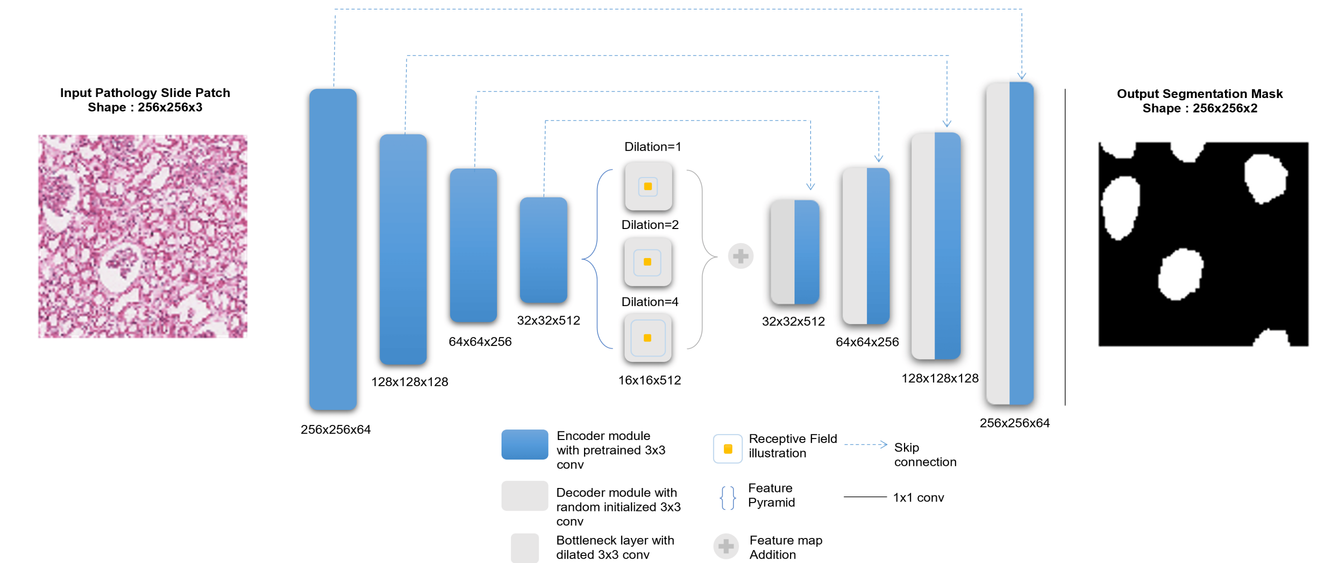

Fig. 4. The adopted model architecture of the autoencoder model we use to extract the latent space embeddings for our image segmentation task. [20]

III. Proposed Method: Decoding the Encoder

In this section, we discuss our proposed novel method of increasing the entropy of the compressed latent space of an autoencoder through monotonic perturbations. For every input, the perturbations in the feature space yield unique signature responses when decoded. Next, we delve deeper into analyzing the decoder responses through unsupervised ML techniques to build a knowledge base. In the testing phase, for a given unknown test case, we can generate the unique signature and find its closest match in the knowledge base. Analyzing the same would help in outlier detection to boost the trustworthiness of the autoencoder model, as illustrated in Fig. 2.

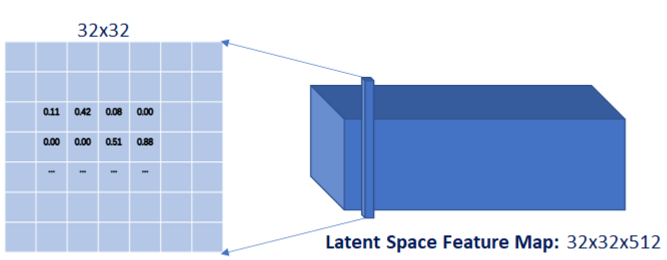

Fig. 5. For all non-zero elements in the feature space matrix corresponding to every input, we perform monotonic discrete perturbations in the latent space.

- **Latent Space Degeneracy:** The design of AEs can be broadly classified as AEs with a deeper encoder than its decoder, a deeper decoder than its encoder, and an encoder-decoder of equal depth. In all scenarios, the goal is to have the latent feature space learn meaning representations from the training dataset. Therefore, the architecture's depth (number of hidden layers) is favorable for extracting meaningful representations. To its contrary, depth in the network also weakens the correlation between the input and the compressed latent space. To emphasize the degeneracy of latent space in AEs, we trained three variants of AEs, namely convolutional AE, variational AE, and de-noising AE on MNIST [12], F-MNIST [13], and K-MNIST [14] datasets. We visualize the latent space through principal component analysis of various AE models for the rising complexity of the model and the dataset they are trained on. We observe that correlation significantly drops between the inputs and the learned representations with increasing model complexity and the dataset it is trained on, as visualised in Fig. 3. This further strengthens our problem statement.
- **Perturbations in the Latent Feature Space:** We monotonically perturb the encoder learned function, $g(.)$, through $S(g(.))$ that increases the entropy in the learned representations. For a given input to a trained encoder, the meaningful representations in the latent space are denoted by the non-zero values. Our method monotonically perturbs the non-zero elements of the latent space matrix from values tending to zero to the maximum possible representative value, which, when normalized, scales from 0 to 1.

- **Decoder Response:** For every input m_i, we perturb the meaningful latent feature space matrix with $S(g(.))$ and ask the decoder to generate the response. The decoder generates unique signatures for every m_i. Ideally, the best response signature is obtained through perturbation in the monotonically continuous domain (which is discretized at a high sampling frequency).
- **Unsupervised Clustering of Decoder Response:** For the perturbed decoder response from every input m_i, we perform quality assessment of the decoded response (as per the application-oriented metric). The responses can then be clustered with unsupervised clustering methods like the birch [15] clustering technique to determine the number of clusters.
- **Individual Cluster Analysis:** For every cluster, which is a function of $S(g(.))$ with quality of the perturbed latent space in a continuous domain, we analyze every cluster. For every cluster, we identify the median response curve and the response curves that are farther from the median. Our hypothesis is that response curves farthest from the cluster's median response curve would be outliers or the ones (classified/segmented/de-noised) with the least confidence.
- **Outlier Detection with Signature Matching:** In the testing scenario, in the absence of ground truth, we can perturb the latent feature space of the test instance to generate the unique signature. We can assign this response to a cluster from the knowledge base with the nearest neighbor measure. Consequently, with the distance metric, we can find the signature's relative distance from its assigned cluster's median response and quantify the confidence in the prediction result of the test case. For several such test cases, we can rank them in order of their relative distance from their assigned cluster's median response curve, in ascending order of their quantified uncertainties as a probable outlier. Ranking all test cases in order would help increase the model's trustworthiness. This would provide medical experts insights into which test cases must be closely observed in safety-critical applications. The one classified/segmented/de-noised by the AE with high confidence can be overlooked.

IV. Evaluation of Proposed Method for Glomeruli Segmentation on Kidney Donor Frozen Sections

We evaluate our proposed novel perturbation method for outlier detection on a very sensitive medical image segmentation application. In this section, we first establish the sensitivity associated with our target application. Following on, we provide the experimental setup with discussions on the dataset and the AE model we use for our study. After that, we delve deeper into analyzing our proposed technique with discussions on every aspect of our novel proposed method and how it boosts the model's trustworthiness.

A. Application overview

Advancements in nephrology research have weighed in on renal allograft transplantion as a substitute for dialysis for subjects with severe kidney disorders [16]. Scarcity of allografts led to the rise of extended criteria donor allografting program wherein transplantation was legalized from cadaveric donors too [17]. To further bolster the program dynamics, Kidney Donor Profiling Index [18] [19] was developed. From an histological point of view, kidney biopsies remain the widely accepted norm of estimating the sclerotic and non-sclerotic glomeruli with hematoxylin and eosin (H&E) staining. Kidneys with non-sclerotic glomeruli above the accepted clinical standard are deemed acceptable for transplantation.

Whole Slide Imaging (WSI) enables digitization of the biopsy samples which in turn enables computational analysis on them. AEs has predominantly established itself in this space. In our evaluation, we study our proposed method on a developed encoder-decoder network [20] which detects and segments sclerotic and non-sclerotic glomeruli on frozen sections of donor kidney biopsies.

B. Experimental Setup

This study was performed at Duke University wherein a bottleneck dialated U-Net model with skip connections [20] was designed to detect and segment sclerotic and non-sclerotic glomeruli on a WSI dataset. Some details outlining the dataset and model is outlined in the following subsections.

1) Dataset: This study involved 268 frozen sections of H&E stained slides from 211 subjects. Out of these 211 subjects, 75 kidney biopsies were performed at Duke University which contributed to 128 H&E stained slides. Rest were performed at other institutions, and post extensive review by the university's medical experts, were approved for the study. A team of three well-trained medical experts, as approved by the university, manually annotated (segmented) the scelotic and non-sclerotic glomeruli. To ensure high quality manual annotation (segmentation), two renal pathologists were further involved in the study. Following a rigorous quality control, these samples were used to train a AE model for detection and segmentation of sclerotic and non-sclerotic glomeruli in frozen kidney donor sections.

2) Autoencoder Model: The 75 WSI kidney biopsies performed at Duke University were split into training and validation dataset in 80-20 ratio, while the 135 WSI kidney biopsies performed at other universities were harnessed as testing dataset. For the detection and segmentation of sclerotic vs. non-sclerotic glomeruli, a nine-module U-Net architecture with dilations at bottleneck layers was adopted from our previous work [20]. It consists of a symmetric encoder-decoder convolutional neural network architecture to ensure pixel-level segmentation. The adopted model to extract latent space embeddings is as depicted in Fig. 4. The encoder uses VGG pre-trained weights which are fine tuned through training epochs. The bottleneck layer dilates features based on convolutional operations. Skip connections help restore encoded information from specific layers of the encoder in the

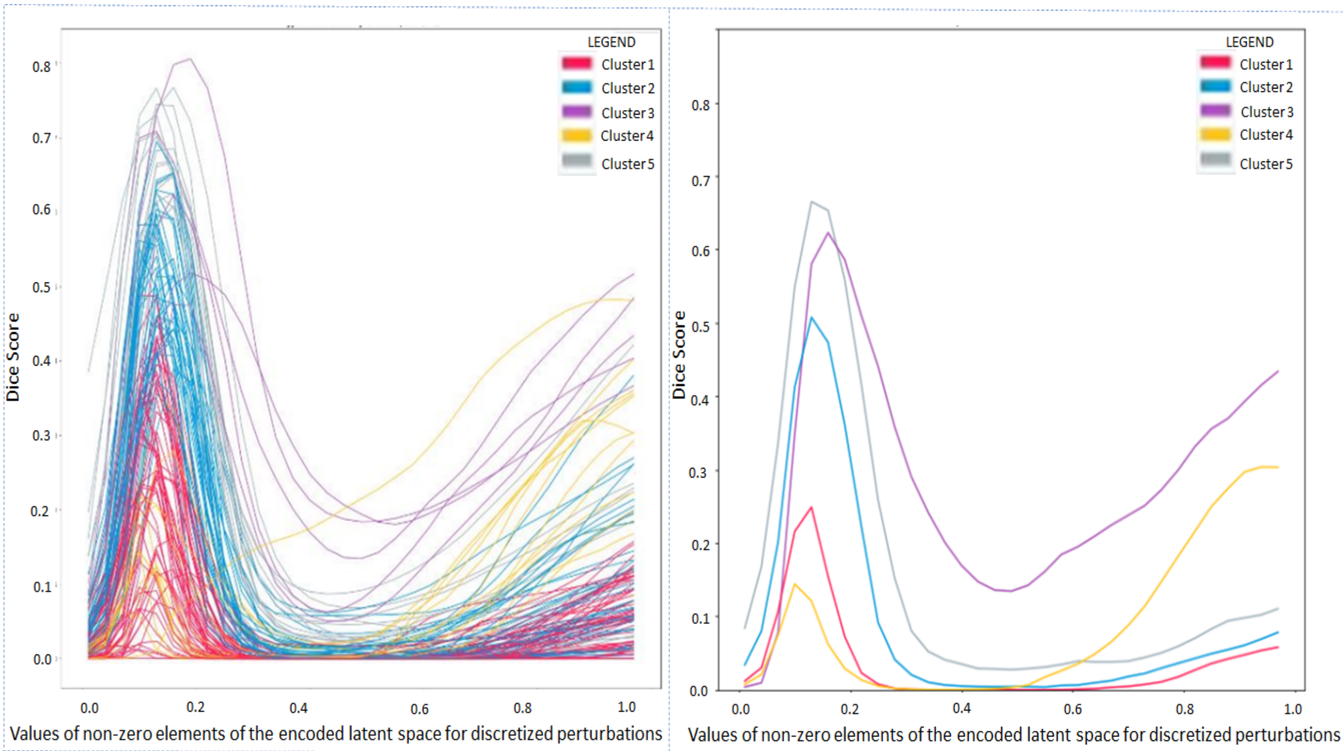

Fig. 6. For all latent space perturbation corresponding to every input, we decode the unique signatures (decoder response). We obtain the number of clusters through unsupervised clustering technique to be 5. Figure on left shows the clustered response curves (agglomerative clustering). Figure on right shows the anchor curve (median response of every cluster).

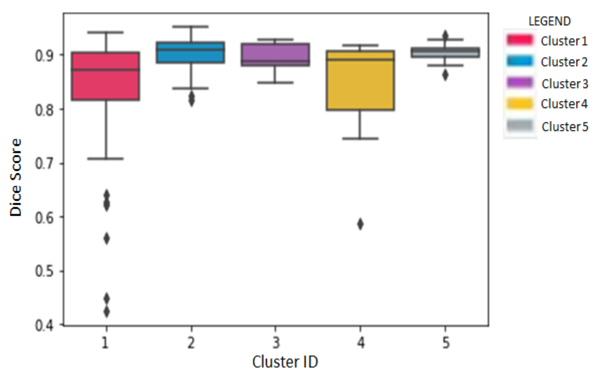

Fig. 7. Outlier analysis in every cluster.

decoding process. The hyper-parameters and implementation details are elaborately illustrated in our previous work [20].

In this work, we focus on quantifying the uncertainties, ranking the test cases based on their uncertainty score, and detect outliers in the test cases (absence of ground truth). For every WSI kidney biopsy input, we perturb the non-zero values of the encoder layer in steps of 0.02, from 0 to the maximum value, to obtain the response curves as a measure of dice score. Next, we cluster the response curves through birch [15] and agglomerative [21] clustering technique to obtain the cluster number. Thereafter, we calculate the median response curve of every cluster and find the outliers. For new test cases, we propose to perturb the learned feature space of the encoder to generate the response curve. The generated response curve can then be attributed to belong to a cluster from the existing knowledge base. This can help to quantify model uncertainties depending on the response curve's relative distance to its assigned cluster center.

C. Results

For every image in the WSI dataset, we introduce monotonically increasing discretized perturbations in learned feature space of the AE model to generate the unique signature corresponding to every input. The signatures are then clustered in an unsupervised fashion for individual cluster analysis. In the testing phase, the responses can be asserted to an existing cluster through signature matching. Our goal remains to rank the test cases in order to their uncertainties in the segmentation task through signature matching to prioritize attention of medical experts on critical cases.

1) Perturbation in latent space: The last block of the trained encoder network consists of a 32x32 feature map, stacked up into 512 layers, as depicted in Fig. 5. For every WSI input, we monotonically perturb the corresponding non-zero elements of the 32x32x512 matrix from zero to its maximum value of 1, discretized in steps of 0.2, pass it through the decoder to obtain the response curve or signature.

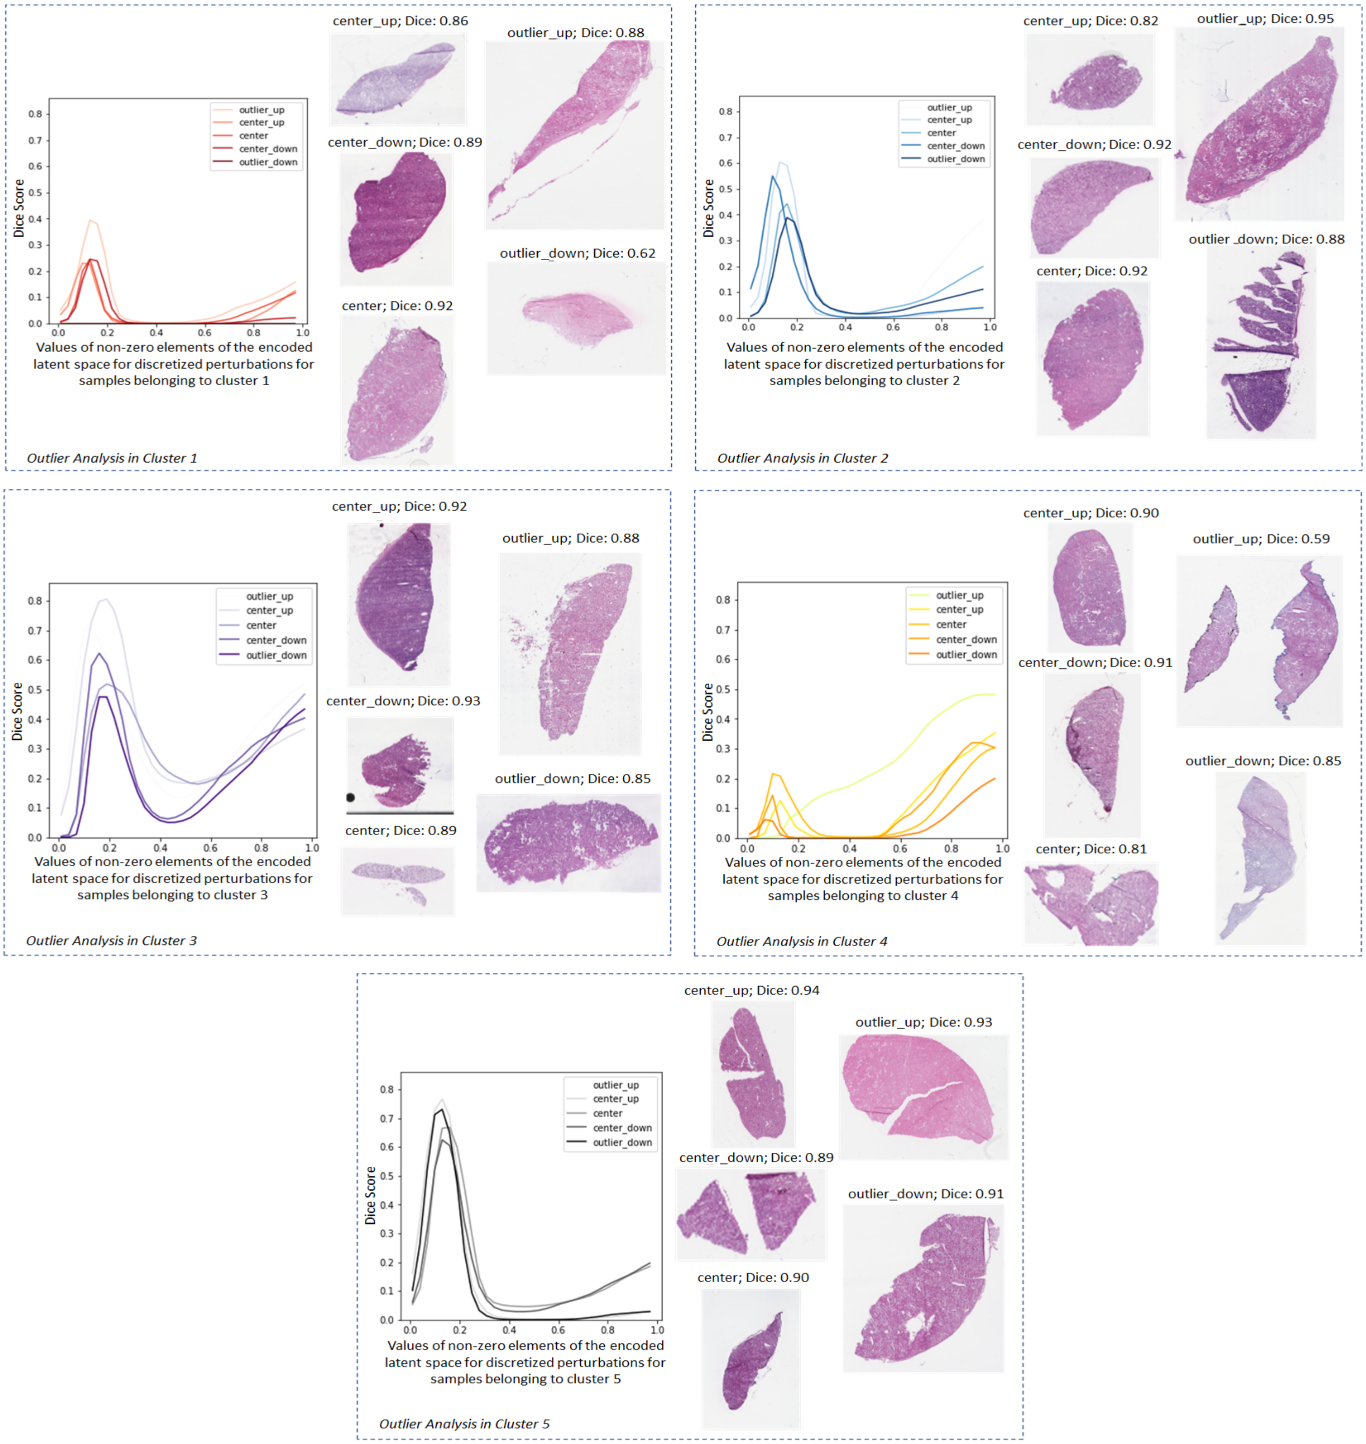

Fig. 8. Evaluation of our proposed method for uncertainty quantification and outlier analysis. For every cluster, we analyze the median response curve, noted by *center*, the signature one unit away in either direction of the median response curve, *center_up*, *center_down*, and the signatures farthest away from the median response curve, *outlier_up*, and *outlier_down*. We are able to detect obvious outliers in clusters 1 and 4. For test cases (absence of ground truth), for every response curve assigned to a cluster, we are able to assign an uncertainty score based on the relative distance from the median response and the absolute difference of dice score between the two extreme cases for the cluster.

2) Clustering the response curves: Quantification of the number of clusters is done in an unsupervised fashion through birch clustering [15] technique. The number of clusters is noted to be 5. The response curves, grouped through agglomerative clustering [21] with the requisite number of clusters, are visualized with its assigned cluster number in Fig. 6.

3) Dice Performance by Cluster: We investigate into the dice performance for every cluster through a box and whisker plot, as visualized in Fig. 7. We are able to detect the outliers from every cluster.

4) Cluster-wise Anomaly Detection: Given the clusters and the median response curve, we analyze signatures those are farthest away from the cluster median. In Fig. 8, for every cluster, we visualize and estimate the model performance on the frozen kidney donor section for the cluster median, signature ordered one unit away from the median response in either direction, and the signatures those are farthest away from the median response in either direction. Median response curve has the highest dice for every cluster, indicating our proposed technique is able to position the best case from every cluster as its median. For outlier analysis, in cluster 1 and 4 we are able to successfully detect an outlier with dice scores of 0.62 and 0.59 respectively. In regards to uncertainty quantification, the absolute difference of dice score between the two extreme cases for each cluster yields the range of uncertainty of the cluster. In testing scenario, for every response curve assigned to a cluster, we are able to assign an uncertainty score based on the relative distance from the median response and the absolute difference of dice score between the two extreme cases for the cluster. Ranking the test cases according to their uncertainty helps medical experts focus on boundary cases while also boosting the trustworthiness of the AE model.

V. Conclusion

In this paper, we proposed a novel method of monotonically perturbing the encoded latent space to increase the entropy in the learned representations for every corresponding input. Through our evaluation use case of glomeruli segmentation in the frozen kidney donor section, we could extract unique signatures for every input, cluster them and analyze the clusters for outliers. Subsequently, in the testing phase, in the absence of ground truth, we were able to rank the cases by first assigning them to an existing cluster in the knowledge base and quantifying the uncertainty according to its relative distance from the cluster's median response curve. This helps boost the trustworthiness of AE models. Therefore, we can demonstrate the effectiveness of our proposed novel uncertainty quantification method and outlier detection.

References

[1] Zheng, H., Yao, J., Zhang, Y., & Tsang, I. W. (2018). Degeneration in vae: in the light of fisher information loss. arXiv preprint arXiv:1802.06677.

[2] Michelucci, U. (2022). An Introduction to Autoencoders. arXiv preprint arXiv:2201.03898.

[3] Bank, D., Koenigstein, N., & Giryes, R. (2020). Autoencoders. arXiv preprint arXiv:2003.05991.

[4] Andrés-Terré, Helena, and Pietro Lió. "Perturbation theory approach to study the latent space degeneracy of Variational Autoencoders." arXiv preprint arXiv:1907.05267 (2019).

[5] Aizenbud, Y., Lindenbaum, O., & Kluger, Y. (2021). Probabilistic Robust Autoencoders for Anomaly Detection. arXiv preprint arXiv:2110.00494.

[6] Finke, T., Krämer, M., Morandini, A., Mück, A., Oleksiyuk, I. (2021). Autoencoders for unsupervised anomaly detection in high energy physics. Journal of High Energy Physics, 2021(6), 1-32.

[7] Kieu, T., Yang, B., Guo, C., Jensen, C. S., Zhao, Y., Huang, F., Zheng, K. (2022). Robust and Explainable Autoencoders for Unsupervised Time Series Outlier Detection. In Proceeding of the 38th IEEE International Conference on Data Engineering, ICDE 2022 (pp. 3038-3050)

[8] Yong, B. X., Brintrup, A. (2022). Do autoencoders need a bottleneck for anomaly detection?. arXiv preprint arXiv:2202.12637.

[9] Chauhan, K., Shenoy, P., Gupta, M., Sridharan, D. (2022). Robust outlier detection by de-biasing VAE likelihoods. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (pp. 9881-9890).

[10] Minhas, M. S., Zelek, J. (2020). Semi-supervised anomaly detection using autoencoders. arXiv preprint arXiv:2001.03674.

[11] Borghesi, A., Bartolini, A., Lombardi, M., Milano, M., Benini, L. (2019, July). Anomaly detection using autoencoders in high performance computing systems. In Proceedings of the AAAI Conference on Artificial Intelligence (Vol. 33, No. 01, pp. 9428-9433).

[12] Y. LeCun, C. Cortes, and C. Burges, "Mnist handwritten digit database," Available: http://yann.lecun.com/exdb/mnist, 1998.

[13] H. Xiao, K. Rasul, and R. Vollgraf, "Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms," arXiv preprint arXiv:1708.07747, 2017.

[14] T. Clanuwat, M. Bober-Irizar, A. Kitamoto, A. Lamb, K. Yamamoto, D. Ha, "Deep Learning for Classical Japanese Literature", arXiv:1812.01718.

[15] Zhang, T., Ramakrishnan, R., Livny, M. (1997). BIRCH: A new data clustering algorithm and its applications. Data mining and knowledge discovery, 1(2), 141-182.

[16] Wolfe, R. A., Ashby, V. B., Milford, E. L., Ojo, A. O., Ettenger, R. E., Agodoa, L. Y., Held, P. J., Port, F. K. (1999). Comparison of mortality in all patients on dialysis, patients on dialysis awaiting transplantation, and recipients of a first cadaveric transplant. New England journal of medicine, 341(23), 1725-1730.

[17] Perico, N., Ruggenenti, P., Scalamogna, M., Remuzzi, G. (2003). Tackling the shortage of donor kidneys: how to use the best that we have. American journal of nephrology, 23(4), 245-259.

[18] Dahmen, M., Becker, F., Pavenstädt, H., Suwelack, B., Schütte-Nütgen, K., Reuter, S. (2019). Validation of the Kidney Donor Profile Index (KDPI) to assess a deceased donor's kidneys' outcome in a European cohort. Scientific reports, 9(1), 1-11.

[19] Pruett, T. L., Vece, G. R., Carrico, R. J., & Klassen, D. K. (2021). US deceased kidney transplantation: Estimated GFR, donor age and KDPI association with graft survival. EClinicalMedicine, 37, 100980.

[20] Davis RC, Li X, Xu Y, Wang Z, Souma N, Sotolongo G, Bell J, Ellis M, Howell D, Shen X, Lafata K. Deep learning segmentation of glomeruli on kidney donor frozen sections. medRxiv. 2021, 8(6), 067501.

[21] Tokuda, E. K., Comin, C. H., & Costa, L. D. F. (2022). Revisiting agglomerative clustering. Physica A: Statistical Mechanics and its Applications, 585, 126433.