

Diff3Dformer: Leveraging Slice Sequence Diffusion for Enhanced 3D CT Classification with Transformer Networks

Zihao Jin^{1*}, Yingying Fang^{2†*}, Jiahao Huang³, Caiwen Xu³, Simon Walsh²,
and Guang Yang^{2,3,4,5†}

¹ Department of Metabolism, Digestion and Reproduction, Imperial College London, London, UK

² National Heart and Lung Institute, Imperial College London, London, UK

³ Bioengineering Department and Imperial-X, Imperial College London, London, UK

⁴ Cardiovascular Research Centre, Royal Brompton Hospital, London, UK

⁵ School of Biomedical Engineering & Imaging Sciences, King's College London, London, UK

{y.fang;g.yang}@imperial.ac.uk

Abstract. The manifestation of symptoms associated with lung diseases can vary in different depths for individual patients, highlighting the significance of 3D information in CT scans for medical image classification. While Vision Transformer has shown superior performance over convolutional neural networks in image classification tasks, their effectiveness is often demonstrated on sufficiently large 2D datasets and they easily encounter overfitting issues on small medical image datasets. To address this limitation, we propose a Diffusion-based 3D Vision Transformer (Diff3Dformer), which utilizes the latent space of the Diffusion model to form the slice sequence for 3D analysis and incorporates clustering attention into ViT to aggregate repetitive information within 3D CT scans, thereby harnessing the power of the advanced transformer in 3D classification tasks on small datasets. Our method exhibits improved performance on two different scales of small datasets of 3D lung CT scans, surpassing the state of the art 3D methods and other transformer-based approaches that emerged during the COVID-19 pandemic, demonstrating its robust and superior performance across different scales of data. Experimental results underscore the superiority of our proposed method, indicating its potential for enhancing medical image classification tasks in real-world scenarios. The code will be publicly available at <https://github.com/ayanglab/Diff3Dformer>.

Keywords: Clustering vision transformer · Diffusion model · 3D CT analysis · Lung disease

*Z. Jin and Y. Fang—Equal contribution.

1 Introduction

3D volume analysis is crucial for the diagnosis or prognosis of lung diseases, as lesions can manifest at various depths in CT scans across different patients [18], such as those with COVID-19 or Interstitial Lung Disease. Compared to 2D analysis, 3D analysis enables a comprehensive examination of abnormal areas throughout the entire volume, offering a more thorough understanding of the patient’s condition. Additionally, models analyzing 3D volumes eliminate the need for slice selection, leading to more efficient and reliable predictions compared to methods that rely on a limited number of preselected slices for patient-level decisions.

Given the imperative nature of these requirements, a quantity of 3D analysis methodologies emerged within the AI community during the COVID-19 pandemic for diagnosing and prognosticating patients [5]. These methodologies can be broadly classified into aggregation (AG) methods [22,14,16], 2.5D methods [15,7], and whole-scan methods (WS) [8,9,6,19]. AG methods analyze 3D scans by aggregating results from all 2D slices [22], inherently limited in capturing intra-slice features. To overcome this limitation, WS methods input the entire scan into the model, allowing for comprehensive feature exploration throughout the 3D volume. Although 3D methods have demonstrated superior performance, they are susceptible to crashing due to overfitting, especially when dealing with a small dataset. As a compromise between diverse training samples and 3D features, 2.5D methods resample a fixed smaller number of slices from the entire scan, treating them as a unified input entity for network-based patient-level decision-making. While the resampling process enables extensive augmentation of the small dataset, the reliance on a randomly sampled subset of slices for patient-level decision-making still raises concerns among doctors who may potentially use these models in high-stakes contexts. Hence, enabling WS analysis within small datasets remains an urgent and unmet challenge.

Transformers have outperformed traditional CNN methods in vision classification tasks but require substantial data and memory resources, posing challenges for small datasets. Recent studies have illuminated the application of advanced Transformer architectures in 3D lung volume classification tasks within limited datasets [7,10,23,22]. To address memory constraints while handling high-dimensional 3D volumes, [7] employed 2.5D techniques, resampling 32 slices as input for the Timesformer model. [22] utilized the AG method, employing a 2D Swin Transformer to process CT volumes with varying slice counts. [10,23] adopted a CNN-based preprocessing step to transform 3D volumes into sequences of low-dimensional CNN-based features, subsequently fed into Transformers for classification. To enhance performance on small datasets, [23] employed Mixup [20] data augmentation, while [7] and [23] explored transfer learning and self-supervised learning to achieve more general representations. Despite improvements brought by these methods, current Transformer-based 3D analysis still faces several limitations: (1) the performance of Transformer-based WS methods is still prone to overfitting and is significantly influenced by data scale; (2) there is a lack of comprehensive comparisons between Transformer-based and CNN-

based models regarding their efficacy and robustness on small 3D datasets; and (3) interpreting the features used for patient-level decisions from these 3D scans remains underexplored in current research.

Motivated by prior work [10,23], our aim is to develop a robust 3D Transformer model that surpasses existing methods by effectively harnessing the global feature learning capabilities of transformers. Simultaneously, we strive to reduce data requirements and improve interpretability in 3D volume decisions. To achieve this, we introduce the novel Diffusion-enhanced 3D Transformer (Diff3Dformer), which combines the advantageous latent space learning of the Diffusion Autoencoder with a Clustering Vision Transformer (ViT). This integration facilitates efficient feature extraction and information reduction during the global feature learning process. The key contributions of this work are summarized as follows: 1) We discover how Clustering ViT can mitigate overfitting and effectively manage small datasets. 2) We introduce the Diffusion Autoencoder for self-supervised learning to extract semantically meaningful representations for enhanced 3D analysis. Additionally, we propose a novel pipeline that enables data-intensive Diffusion to be applied to small-scale 3D analysis using the efficient 3D solver Clustering ViT. 3) We propose an interpretable slice fusion strategy to decode the model’s decisions into contributions from different clusters, enabling the explainability of the final patient-level decision from the Diff3Dformer. 4) We conduct experiments on two different scales of small datasets, showcasing the robustness and consistent superiority of the proposed methods over different types of 3D analysis methods across varying medical image dataset scales.

2 Method

The overview of Diff3Dformer is given in Fig. 1 (A). Prior to Diff3Dformer’s prediction on individuals, we employ the encoder from the pretrained diffusion autoencoder to extract representations of each slice from CT volumes. By aggregating the slice representations from the entire dataset, we can learn slice prototypes (the centre of each cluster) specific to a particular disease using the spherical K-means method. Given the learned prototype, Diff3Dformer starts by transforming a patient’s 3D volumes into a sequence of the representations together with the cluster number it belongs to. The representation together with the assigned prototype number, they are fed into the Clustering ViT for global information learning through the self-attention map in the transformer. The cluster number here aids the model in detecting repetitive and similar patterns in the 3D volume, thereby reducing the number of features and enhancing computational efficiency within the traditional ViT. Following the modification of slice representations, the final layer of Diff3Dformer outputs scores for each slice to make the final decision. Diff3Dformer employs global attention on predefined clusters, which are learned during training, to fuse slice scores from different clusters, thereby generating an explainable patient classification result.




Fig. 1. (A) The overview framework of Diff3Dformer. (B) The diffusion autoencoder is leveraged to learn a semantically meaningful representation by learning to reconstruct the 2D slice from a 512-dimensional representation and being used to represent CT volumes as a sequence of representations as the input of the clustering ViT model. (C) The slice fusion module provides final patient decisions and explanations of Diff3Dformer.

2.1 Representation Learning via Diffusion-based Autoencoder

Motivated by recent strides in feature manipulation and disentanglement within Diffusion’s latent space [13,17,1], we are compelled to exploit these semantically meaningful features as representations for individual slices for downstream tasks. To derive highly meaningful representations for each slice, we utilized a Denoising Diffusion Implicit Model (DDIM)-based autoencoder, as proposed by Preechakul et al. [17], to reconstruct slices from CT scans. This autoencoder architecture consists of an encoder $\mathbf{E}$ and a DDIM model denoted as $\mathbf{D}$. To preserve meaningful information within the encoded representation from $\mathbf{E}$, DDIM is trained to reconstruct the original slice using this representation as a condition.

The models $\mathbf{E}$ and $\mathbf{D}$ are trained concurrently by optimizing the following loss function with respect to θ and ϕ :

$$\min_{\phi, \theta} \mathcal{L} = \|\epsilon - \mathbf{D}_{\theta}(\mathbf{x}_t, t, \mathbf{E}_{\phi}(\mathbf{x}_0))\|_1, \quad (1)$$

where $\mathbf{x}_0$ represents any given slice and $\mathbf{x}_t$ is the noise injected slice (t iterations of Gaussian noise injection). The network $\mathbf{D}$ utilizes a UNet architecture consisting of layers of residual blocks, as described in [2]. Meanwhile, the network $\mathbf{E}$ adopts the encoder architecture from $\mathbf{D}$.

Once the autoencoder achieves the optimal reconstruction quality, the encoder is separately utilized to extract the representations of each slice. These

representations are then aggregated from each patient and clustered into K clusters using Spherical K-means [25]. The clustering step will learn the potential prototypes of the slices within a specific dataset. These prototypes will further enable the quantification of the entire scan into a combination of prototype slices by grouping the slices with similar patterns together. Additionally, it will aid in reducing the features during self-attention learning in the subsequent Clustering ViT model introduced.

2.2 Clustering ViT for 3D Classification

After representation learning, each 3D volume can be transformed into a sequence of meaningful slice representations, each with its corresponding assigned cluster. Inspired by [23,24], we introduce a clustering ViT model for 3D diagnostic and prognostic tasks based on the obtained slice sequence.

As illustrated in Fig. 1 (A), the slice sequences obtained from each patient are padded to a fixed length N , mapped to M dimensions using a linear layer, and then fed into a six-layer Clustering Transformer Encoder. Each layer comprises a clustering attention mechanism with 8 heads and a feed-forward network. Notably, the clustering attention block, proposed in [24], computes the prototype of the queries in each cluster, reducing the number of queries from N to K . This reduces the computational complexity of the attention map from $O(N^2)$ to $O(NK)$ compared to traditional ViT architectures [4]. For our 3D classification task, the clusters within the model correspond to the clusters assigned to each slice, simplifying queries of similar slices into single features of dimension M . The final result of the attention and values consists of K updated vectors, which are then broadcasted back to the N updated slices denoted as s^* by replicating each feature s_k^* into slices assigned to cluster k . Besides computational efficiency, the clustering mechanism also reduces the final updated features in s^* by replicating the prototypes into high-dimensional data, effectively performing dimension reduction and hence mitigating overfitting issues.

Following the Clustering Transformer Encoder, the Diff3Dformer processes the updated features obtained from global learning through a linear layer to obtain the risk score for each slice denoted as r . After the final layer of slice fusion, the model generates a single score as the patient-level score. For our 3D classification task, the clustering ViT model is trained using cross-entropy loss.

2.3 Interpretable Slice-Sequence Fusion

The fusion of slice sequences plays a pivotal role in consolidating information to generate the final patient-level decision. Traditionally, this fusion is accomplished using various pooling methods or linear regression in conventional 3D analysis techniques. In order to avoid potential overfitting resulting from dense layer and the direct averaging of patch levels, which may disregard the varying importance of individual slices in classification tasks, we propose an interpretable 3D decision-making approach, which considers the existence of different prototypes, the quantification of various clusters, and the diverse contributions of slice

patterns to the final task. This can be formulated as:

$$R = \sum_{k=1}^K A_k \bar{r}_k q_k. \quad (2)$$

Here, A_k represents the global cluster attention, emphasizing the significance of the presence of a specific cluster for the final task, which remains consistent across all patients. q_k indicates the ratio of the number of slices in each cluster to the total number of slices from a patient, simulating the lesion extent, while $\bar{r}_k$ denotes the average slice risk within cluster k for each individual.

3 Experiment

3.1 Dataset

To validate the effectiveness of the proposed method on small datasets across different medical tasks, we evaluated the performance of the Diff3Dformer model in both diagnosis and prognosis tasks using two 3D datasets: COVID-19 and fibrotic lung disease (FLD). Specifically, we validated the performance of our model on the CC-CCII [21] dataset to tackle the classification of novel coronavirus pneumonia (NCP) and common pneumonia (CP), and on the FLD dataset for a binary prognostic task to predict the 1-year mortality of FLD patients.

Clean-CC-CCII: The Clean-CC-CCII dataset is a publicly available dataset of CT volumes consisting of three different categories: NCP, CP, and normal patients, which is constructed by preprocessing and restructuring the CC-CCII dataset [21] in [8]. The Clean-CC-CCII dataset contains 3,993 scans from 2,698 patients. In this study, we perform a binary classification task of NCP and CP classes, including 1519 scans from 1047 NCP patients and 1549 scans from 824 CP patients. In our experiments, We randomly divided the scans into training data (2455 scans), validation data (306 scans), and test data (307 scans).

Fibrotic lung disease: The FLD dataset is the public dataset from OSIC¹, comprising 27 patients who died within one year and 704 patients who survived beyond one year during their hospitalisation. We reserve 20% of patients for validation, and use the remaining for training. An in-house external test dataset is obtained from Australia, consisting of 501 CT scans, with 43 patients who died within one year and 458 patients who survived beyond one year.

3.2 Implementation Details

For representation learning, the Adam optimizer [12] with a batch size of 64 is used to optimize the diffusion autoencoder, and the learning rate is set to $1e^{-4}$. The input size of the diffusion autoencoder is 256×256 . It is trained by 93967 slices generated from the 3D OSIC dataset. We trained the model using 8 V100

¹ <https://www.osicild.org/>




Fig. 2. The comparison results of different methods on CC-CCII and FLD datasets

GPUs for 100 epochs. The number of clusters K in the spherical K-means method is set to 64. The clustering ViT model is trained using the Adam optimizer [12] on two RTX3090 GPUs with a batch size of 4 and a learning rate of $1e^{-4}$ for 100 epochs. The dimensional size M is set to 512 and the dropout rate is set to 0.1. The area under the curve (AUC), accuracy, sensitivity, specificity, and F1 score were used as metrics for evaluating the performance of classification.

3.3 Experiment Results

In this study, we compared Diff3Dformer model with other 3D CNN-based methods and transformer-based methods which have open-source code. The comparison results on the two datasets are presented in Fig. 2. The 3D CNN-based methods include WS-DenseNet121 [8], WS-ResNet101 [8], WS-Contrastive 3D [9], 2.5D-ResNet101 [6], and the 3D transformer-based methods include AG-Swin Transformer [22] and ViT-patch [23]. The experiment setting of these methods can be found in Supplementary Table 1.

The proposed method outperformed other transformer-based methods. Compared to the AG-Swin Transformer and ViT-patch methods, our proposed model achieves superior performance in terms of AUC, sensitivity, and F1 score on the CC-CCII dataset, while demonstrating comparable performance on other metrics. On the smaller FLD dataset, the proposed method significantly enhances performance across all metrics, unlike other transformer-based methods, which exhibit sensitivity to dataset size and fail to produce satisfactory results on smaller datasets. These findings suggest that our model effectively mitigates the requirement for large datasets typically needed by transformer-based methods and demonstrates greater robustness with limited data.

The proposed method consistently outperformed different type of 3D classification models. Comparing the F1 score on both datasets, we observe

Table 1. Ablation studies on CC-CCH and FLD datasets.

No.	Ablation Setting	CC-CCH					FLD				
		AUC	Accuracy	Sensitivity	Specificity	F1 Score	AUC	Accuracy	Sensitivity	Specificity	F1 Score
1.	Contrastive + ViT	0.83	0.82	0.77	0.79	0.78	0.75	0.74	0.68	0.77	0.72
2.	Diffusion + ViT	0.84	0.84	0.78	0.81	0.79	0.76	0.76	0.68	0.79	0.73
3.	Contrastive + clustering ViT	0.88	0.84	0.81	0.83	0.82	0.78	0.75	0.75	0.74	0.74
4.	Diffusion + clustering ViT	0.91	0.85	0.86	0.83	0.84	0.79	0.77	0.77	0.75	0.76

that CNN-based methods also tend to perform worse on the extremely small FLD dataset. Specifically, WS-based methods perform the worst due to the susceptibility to overfitting issues. These challenges can be partly addressed by learning a more generalizable representation through contrastive learning techniques, as demonstrated in the WS-Contrastive method and resampling methods in the 2.5D method, where the WS method outperforms the 2.5D method through a more comprehensive analysis. In comparison to these methods on the FLD dataset, both patch-ViT and Diff3Dformer, which utilize the Clustering ViT architecture, outperform those CNN-based methods, indicating their effectiveness in reducing overfitting issues. Moreover, Diff3Dformer outperforms patch-ViT by leveraging slice-sequence analysis on both datasets.

Ablation study. To investigate the effectiveness of components in the proposed model on the small dataset, we conduct some experiments on both two datasets as shown in Table 1. Contrastive learning [11] is another useful self-supervised learning method to learn image representation, and we also remove the clustering attention in the clustering ViT model and compare their performance with the proposed model. The model without clustering attention is identical to the original ViT [3]. The comparison between No.1 and No.2, as well as No.3 and No.4, demonstrates that the diffusion model achieves better representations compared to the contrastive learning method. When comparing No.1 to No.3 and No.2 to No.4, it is evident that clustering attention significantly improves the performance of ViT on both datasets, confirming that clustering attention effectively addresses the overfitting problem in transformer-based methods.

Interpretable results. Based on Eqn. (2), we can identify the most influential cluster contributing to the final score R for each individual by vectorizing the feature $A_k \bar{r}_k$ for each cluster. The heatmap in Supplementary Fig. 1 represents the contribution of the cluster to the final patient-level risk score R on the FLD dataset, where the panels from left to right depict the $A_k \bar{r}_k$ vectors for patients arranged in decreasing order of R value. The rationale behind each patient’s final prediction: the red cube highlights clusters contributing to high-risk scores, while blue indicates a lower risk. From this visualization, we can see that patients with different prediction results are highly disentangled, and the contributing patterns are clearly delineated for each patient. The most influential clusters across the dataset are determined by comparing the average $A_k \bar{r}_k$ values between the two classes with different predictions. The ranking of clusters by contribution to the

‘mortality in one year’ class on the FLD dataset is shown in Supplementary Fig. 2 and the most representative slice patterns are provided in Supplementary Fig. 3, which show that the model can identify common clusters within each class group, enabling us to pinpoint most significant features by visualizing the most frequently contributing clusters among patients.

4 Conclusion

In this paper, we propose Diff3Dformer specifically tailored to overcome the challenges encountered in classifying 3D CT scans using small medical image datasets, outperforming both CNN-based and Transformer-based methods. Leveraging Diffusion-based slice-sequence representations empowers Transformer architecture for high-dimensional 3D volume data, and enhances classification accuracy with its rich and meaningful feature representation. Experimental results demonstrate the superior performance of our proposed method across various scales of small datasets and medical image classification tasks.

Acknowledgments. This study was supported in part by the ERC IMI (101005122), the H2020 (952172), the MRC (MC/PC/21013), the Royal Society (IEC/NSFC/211235), the NVIDIA Academic Hardware Grant Program, the SABER project supported by Boehringer Ingelheim Ltd, Wellcome Leap Dynamic Resilience, NIHR Imperial Biomedical Research Centre, and the UKRI Future Leaders Fellowship(MR/V023799/1).

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Cho, W., Ravi, H., Harikumar, M., Khuc, V., Singh, K.K., Lu, J., Inouye, D.I., Kale, A.: Towards enhanced controllability of diffusion models. arXiv preprint arXiv:2302.14368 (2023)
2. Dhariwal, P., Nichol, A.: Diffusion models beat gans on image synthesis. *Advances in neural information processing systems* **34**, 8780–8794 (2021)
3. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., Houlsby, N.: An image is worth 16x16 words: Transformers for image recognition at scale. *CoRR abs/2010.11929* (2020), <https://arxiv.org/abs/2010.11929>
4. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929 (2020)
5. Fang, Y., Xing, X., Wang, S., Walsh, S., Yang, G.: Post-covid highlights: Challenges and solutions of artificial intelligence techniques for swift identification of covid-19. *Current Opinion in Structural Biology* **85**, 102778 (2024)

6. Harmon, S.A., Sanford, T.H., Xu, S., Turkbey, E.B., Roth, H., Xu, Z., Yang, D., Myronenko, A., Anderson, V., Amalou, A., et al.: Artificial intelligence for the detection of covid-19 pneumonia on chest ct using multinational datasets. *Nature communications* **11**(1), 4080 (2020)
7. Hartmann, K., Hortal, E.: Covid-19 diagnosis in 3d chest ct scans with attention-based models. In: *International Conference on Artificial Intelligence in Medicine*. pp. 229–238. Springer (2023)
8. He, X., Wang, S., Shi, S., Chu, X., Tang, J., Liu, X., Yan, C., Zhang, J., Ding, G.: Benchmarking deep learning models and automated model design for covid-19 detection with chest ct scans. *MedRxiv* pp. 2020–06 (2020)
9. Hou, J., Xu, J., Feng, R., Zhang, Y., Shan, F., Shi, W.: Cmc-cov19d: Contrastive mixup classification for covid-19 diagnosis. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 454–461 (2021)
10. Hsu, C.C., Chen, G.L., Wu, M.H.: Visual transformer with statistical test for covid-19 classification. *arXiv preprint arXiv:2107.05334* (2021)
11. Huang, J., Dong, Q., Gong, S., Zhu, X.: Unsupervised deep learning by neighbourhood discovery. In: Chaudhuri, K., Salakhutdinov, R. (eds.) *Proceedings of the 36th International Conference on Machine Learning*. *Proceedings of Machine Learning Research*, vol. 97, pp. 2849–2858. PMLR (09–15 Jun 2019), <https://proceedings.mlr.press/v97/huang19b.html>
12. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014)
13. Li, A.C., Prabhudesai, M., Duggal, S., Brown, E., Pathak, D.: Your diffusion model is secretly a zero-shot classifier. *arXiv preprint arXiv:2303.16203* (2023)
14. Mei, X., Lee, H.C., Diao, K.y., Huang, M., Lin, B., Liu, C., Xie, Z., Ma, Y., Robson, P.M., Chung, M., et al.: Artificial intelligence-enabled rapid diagnosis of patients with covid-19. *Nature medicine* **26**(8), 1224–1228 (2020)
15. Meng, Y., Bridge, J., Addison, C., Wang, M., Merritt, C., Franks, S., Mackey, M., Messenger, S., Sun, R., Fitzmaurice, T., et al.: Bilateral adaptive graph convolutional network on ct based covid-19 diagnosis with uncertainty-aware consensus-assisted multiple instance learning. *Medical Image Analysis* **84**, 102722 (2023)
16. Miron, R., Moisii, C., Dinu, S., Breaban, M.: Covid detection in chest cts: Improving the baseline on cov19-ct-db. *arXiv preprint arXiv:2107.04808* (2021)
17. Preechakul, K., Chatthee, N., Wizadwongsa, S., Suwajanakorn, S.: Diffusion autoencoders: Toward a meaningful and decodable representation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 10619–10629 (2022)
18. Shamshad, F., Khan, S., Zamir, S.W., Khan, M.H., Hayat, M., Khan, F.S., Fu, H.: Transformers in medical imaging: A survey. *Medical Image Analysis* p. 102802 (2023)
19. Wang, X., Deng, X., Fu, Q., Zhou, Q., Feng, J., Ma, H., Liu, W., Zheng, C.: A weakly-supervised framework for covid-19 classification and lesion localization from chest ct. *IEEE transactions on medical imaging* **39**(8), 2615–2625 (2020)
20. Zhang, H., Cisse, M., Dauphin, Y.N., Lopez-Paz, D.: mixup: Beyond empirical risk minimization. *arXiv preprint arXiv:1710.09412* (2017)
21. Zhang, K., Liu, X., Shen, J., Li, Z., Sang, Y., Wu, X., Zha, Y., Liang, W., Wang, C., Wang, K., et al.: Clinically applicable ai system for accurate diagnosis, quantitative measurements, and prognosis of covid-19 pneumonia using computed tomography. *Cell* **181**(6), 1423–1433 (2020)

22. Zhang, L., Wen, Y.: Mia-cov19d: a transformer-based framework for covid19 classification in chest cts. In: Proceeding of the IEEE/CVF International Conference on Computer Vision Workshops. pp. 513–8 (2021)
23. Zhao, A., Shahin, A.H., Zhou, Y., Gudmundsson, E., Szmul, A., Mogulkoc, N., van Beek, F., Brereton, C.J., van Es, H.W., Pontoppidan, K., et al.: Prognostic imaging biomarker discovery in survival analysis for idiopathic pulmonary fibrosis. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 223–233. Springer (2022)
24. Zheng, M., Gao, P., Zhang, R., Li, K., Wang, X., Li, H., Dong, H.: End-to-end object detection with adaptive clustering transformer. arXiv preprint arXiv:2011.09315 (2020)
25. Zhong, S.: Efficient online spherical k-means clustering. In: Proceedings. 2005 IEEE International Joint Conference on Neural Networks, 2005. vol. 5, pp. 3180–3185. IEEE (2005)