

# Leveraging Human Selective Attention for Medical Image Analysis with Limited Training Data

Yifei Huang<sup>1</sup>

hyf@iis.u-tokyo.ac.jp

Xiaoxiao Li<sup>2</sup>

xiaoxiao.li@ece.ubc.ca

Lijin Yang<sup>1</sup>

yang-lj@iis.u-tokyo.ac.jp

Lin Gu<sup>3</sup>

lin.gu@riken.jp

Yingying Zhu<sup>4</sup>

yingying.zhu@uta.edu

Hirofumi Seo, Qiuming Meng<sup>1</sup>

hseotky@gmail.com, mengqiuming1993@gmail.com

Tatsuya Harada, Yoichi Sato<sup>1</sup>

harada@mi.t.u-tokyo.ac.jp, ysato@iis.u-tokyo.ac.jp

<sup>1</sup> The University of Tokyo  
Japan

<sup>2</sup> The University of British Columbia  
Canada

<sup>3</sup> RIKEN, AIP  
Japan

<sup>4</sup> University of Texas Arlington  
USA

---

## Abstract

Human gaze is a cost-efficient physiological data that reveals human underlying attentional patterns. The selective attention mechanism helps the cognition system focus on task-relevant visual clues by ignoring the presence of distractors. Thanks to this ability, human beings can efficiently learn from a very limited number of training samples. Inspired by this mechanism, we aim to leverage gaze for medical image analysis tasks with small training data. Our proposed framework includes a backbone encoder and a Selective Attention Network (SAN) that simulates the underlying attention. The SAN implicitly encodes information such as suspicious regions that is relevant to the medical diagnose tasks by estimating the actual human gaze. Then we design a novel Auxiliary Attention Block (AAB) to allow information from SAN to be utilized by the backbone encoder to focus on selective areas. Specifically, this block uses a modified version of a multi-head attention layer to simulate the human visual search procedure. Note that the SAN and AAB can be plugged into different backbones, and the framework can be used for multiple medical image analysis tasks when equipped with task-specific heads. Our method is demonstrated to achieve superior performance on both 3D tumor segmentation and 2D chest X-ray classification tasks. We also show that the estimated gaze probability map of the SAN is consistent with an actual gaze fixation map obtained by board-certified doctors.

# 1 Introduction

Human gaze behavior is ecologically linked to the spatial underlying attention [36, 50]. With numerous objects and rich information in view, human learns to develop a selective attentional mechanism that instinctively moves the fovea centralis to process the important task-relevant visual clues. This strategy is crucial to protect the brain, which has limited sensory and cognitive capacities, from being bombarded with an excess of information [8]. “Without selective interest, experience is an utter chaos” [20]. Thanks to the selective attentive mechanisms, humans can efficiently learn from experience with limited samples.

The selective attentional mechanism of the human gaze has been exploited to provide additional cues to improve the performance of multiple tasks, such as action recognition [16, 29, 63], object referring [42], and image captioning [10]. A very recent work [60] demonstrates that the estimation of gaze could effectively improve the classification even without directly accessing the target data (zero-shot learning).

Inspired by the existing success, we explore the possibility of using human gaze supervised attention model to promote medical image analysis tasks under limited training data. We propose a novel framework that exploits the selective attention revealed by human gaze to allow efficient training with limited data. As illustrated in Fig. 1, we simulate the human’s selective attention with a Selective Attention Network (SAN, bounded in blue), which is trained with actual human gaze collected from medical experts when they were performing the screening. We design a plug-and-play module that enables the information from SAN to be leveraged for directly enhancing the backbone performance on multiple tasks. This is achieved by a novel Auxiliary Attention Block (AAB) that integrates the information from SAN into the backbone. Noting the scaled dot-product attention in transformer [43] is very similar to the attention exhibited by human gaze, we design the AAB based on a transformer architecture. This block uses features from the backbone encoder as the *key* vectors and *value* vectors, and the representation from SAN as the *query* vectors. For each query vector, the transformer computes its relation with all locations of the key vector to generate an attention map, which is very similar to the process that human medical experts’ search suspicious regions on the image based on their spatial underlying attention. As a result, the embedded attention refines the features from the backbone encoder, making it better focus on important regions when conducting segmentation or classification tasks.

To evaluate the proposed framework, we conduct experiments on two public datasets. The first experiment is brain tumor segmentation on the BraTS 2020 dataset [2]. We show that our method can achieve superior performance when compared to other state-of-the-art methods when using limited training data. We also collect gaze data of board-certified doctors for the BraTS 2020 dataset which could further contribute to the community. For a more complete study of our framework, we also experiment with chest X-ray classification tasks on the MIMIC-CXR-gaze dataset [25]. Experiments demonstrate that our framework can exploit gaze information to benefit multiple backbone networks with only limited training data. The experiments with two kinds of tasks show the generalization ability of our framework.

Our contributions are summarized as follows:

- We propose a novel framework that leverages human gaze to enhance the medical image tasks of classification and segmentation on both 2D and 3D data. This could promote medical AI’s real-world clinical applications where labels are hard to collect.
- We collect new gaze data of board-certified doctors is for the BraTS 2020 dataset.
- Experiments on the BraTS 2020 dataset [2] and the MIMIC-CXR-gaze [25] dataset

demonstrate that our framework can significantly improve the performance of the backbone network on limited training data, with the help of human gaze in training.

## 2 Related Works

**Gaze assisted deep learning** Since gaze is a cost-efficient physiological data that could be relatively easy to collect, there is a rising trend of collecting and utilizing the gaze data [17, 25, 29] to assist deep learning. Noticing gaze are directly relevant to human attention or intention from the first-person viewpoint [15, 27], early researches mainly leverage gaze for first-person action recognition [10, 16, 28]. This is usually done by multi-task learning [24] or using gaze as soft attention maps [29]. However, the nice performance of these action classification methods could not be replicated on classification tasks such as chest X-ray classification [25]. More recently, human gaze are used to improve model performance in other tasks such as object detection [42], knowledge distillation [65] and image captioning [42]. A very recent work leveraged human gaze to improve image classification when no training sample of target class is given [60]. However, this method relies on additional attribute information, and is not flexible for other tasks such as segmentation. In this work, we propose a novel framework that explores the use of gaze information to enhance multiple medical image analysis tasks including brain tumor segmentation and chest X-ray classification with limited data.

**Brain tumor segmentation** Tumor segmentation is one of the most challenging task in medical imaging, because lesions are relative small compared to the whole image and may have various deformations, sizes or subtypes. DeepMedic [23] is a pioneering work showing superior performance for brain tumor segmentation using 3D patches with 3D CNN. Following DeepMedic, Casamitjana *et al.* [4] proposed a 3D CNN to take the entire 3D volume for prediction. Brain tumor segmentation performances on BraTS challenges [11] are significantly boosted with 3D UNets [6, 8, 19, 26, 37]. 3D CNN typically contains much more parameters than 2D CNN, thus its performance is constrained by the limited labeled training samples in the medical imaging domain. Although there are existing works directly using gaze data for brain tumor segmentation [68, 69], how to benefit model design on general medical tasks from gaze mechanism is under-explored.

**Chest X-Ray** Wang *et al.* firstly created a multi-label large-scale Chest X-ray (CXR) diagnosis dataset by extracting 14 disease labels from clinical notes [44]. Following works either using different backbones [45] or leveraging semantic disease relationship to improve the CXR diagnosis performance [61]. Although these works show ideal performance with large training dataset, the performance severely decreases if trained on a small CXR dataset, even on the disease with fewer training samples. Collecting a large-scale and balanced CXR dataset is almost impossible in real-world clinical settings due to the expensive expert labelling cost and heterogeneous disease distribution. Alternatively, we aim at training a CXR classification model using smaller dataset by leveraging the radiologists' gaze information since the radiologist can learn to diagnose CXR with small size training samples.

## 3 Our Proposed Framework

Here, we propose a framework that leverages gaze data as an auxiliary supervisory signal to improve the performance of different tasks, particularly under limited training data. As illustrated in Fig 1, our proposed framework is composed of three parts: (1) a backbone




Figure 1: Overview of the proposed framework.

encoder that encodes the input images into features, and different heads for performing corresponding tasks such as segmentation and classification. (2) A Selective Attention Network (SAN) that simulates the human selective attention mechanism and maintains a representation of overall cognition, and (3) in between the backbone encoder and the task-related heads, we insert a novel Auxiliary Attention Block (AAB) that leverages the cognition representation from SAN to guide the framework to attend to the disease-relevant clues.

**Backbone Encoder & task-specific heads** By design, our framework can be applied to many existing medical imaging networks by integrating the SAN and AAB. Here, we demonstrate the 3D brain tumor segmentation task based on [13] with 3D Unet as the backbone [4] and use a segmentation head. For the 2D chest X-ray classification, we select [25] with EfficientNet [40] as the backbone and use a classification head.

In the following part of this section, we showcase our framework on the task of 3D brain tumor segmentation [4], where the inputs are 3D scans of brain images  $I \in \mathbb{R}^{D,W,H}$ , where  $D, W, H$  are the input depth, width, and height, respectively. 2D task of chest X-ray classification could be regarded as a special case of 3D with  $D = 1$ , which also involves replacing 3D kernels with 2D ones.

### 3.1 Selective Attention Network

SAN is inspired by the behavior of medical experts during screening, where they first selectively go through the whole slices of images quickly to acquire an overall cognition of the patient’s condition. As shown in the blue box of Fig 1, the Selective Attention Network (SAN) is a shallow network with 2 convolution groups that takes the  $4 \times$  down-sampled input images  $I^{ds}$  as input. Each convolution block contains two  $3 \times 3 \times 3$  convolution layers (with group normalization [47] and ReLU), followed by a 3D max-pooling layer. Notations such as  $M_2$  denote the output of each convolution group.

When medical experts skim on the images, the selective attention mechanism would make them instinctively gaze at suspicious regions [50]. Therefore, we supervise the SAN with the collected medical expert’s gaze. As shown in Fig 1, the output of the second convolution group  $M_2$  is forwarded to a  $1 \times 1 \times 1$  convolution layer to generate the gaze estimation map  $G \in \mathbb{R}^{1,d,w,h}$ , where  $d, w, h$  are  $\frac{1}{16}$  of the original image size. By learning to estimate human




Figure 2: Architecture of the Auxiliary Attention Block. This block gets information from the SAN and feeds the backbone network information in the form of attention refined features.

gaze,  $M_2$  selectively encodes the information (e.g. location, shape) of the suspicious regions. Therefore, we regard  $M_2$  as a representation of selective attention and pass it to the backbone.

### 3.2 Auxiliary Attention Block

To better leverage the human underlying attention, we design a novel Auxiliary Attention Block and insert it between the encoder and the task heads. Specifically, our AAB fuses the backbone encoded feature  $E_4$  with the latent representation of underlying attention  $M_2$  from the SAN.

We illustrate the proposed AAB in Fig. 2. The core of the AAB is a modified multi-head attention layer [43]: using a *query* vector generated from  $M_2$ , an attention weight on each location of the *key* vector (generated from  $E_4$ ) is computed. This layer thus generates attention weights to refine the feature by only focusing on the important, relevant parts. This is inspired from human visual attention mechanism, where human searches the visible space ( $E_4$ ) using the selective attention ( $M_2$ ) in mind [44] as query. More formally speaking, the encoded feature  $E_4$  is first down-sampled by a max pooling layer to match the spatial size of the cognition feature  $M_2$ , denoted as  $E_4^{ds} \in \mathbb{R}^{c_e \times d \times w \times h}$ . Following the previous success [52], we add a fixed 3D positional encoding  $P^{D,W,H} \in \mathbb{R}^{3,D,W,H}$  where in each position,  $P_{i,j,k}^{D,W,H} = (\frac{2i-D+1}{D-1}, \frac{2j-W+1}{W-1}, \frac{2k-H+1}{H-1})$  and concatenate with  $E_4^{ds}$ . We then reshape the features by flattening all the spatial dimensions, and project  $E_4^{ds}$  as the value  $V$  and key  $K$ , and project  $M_2$  as the query vector  $Q$ . The shape of  $Q, K, V$  are all  $\mathbb{R}^{(dwh) \times c_h}$  where  $c_h = 64$  is the hidden dimension. The core component in AAB is a multi-head attention block [43], where the query-key-value vectors are computed in the following manner:

$$Att(Q, K, V) = Softmax\left(\frac{QK^T}{\sqrt{c_h}}\right)V \quad (1)$$

$$MultiHead(Q, K, V) = Concat(Att_1, Att_2, \dots, Att_k)W^O. \quad (2)$$

Here  $W^O$  is a learnable weight matrix.

As shown in Figure 2, the output of the multi-head attention block further passes through normalization layers and a small two-layer feed forward network (FFN) to get the refined feature  $M'$ . After up-sampling,  $M'$  is then added with the original input feature  $E_4$ , forming the output  $E_4'$ . Finally, the refined feature  $E_4'$  is forwarded to the segmentation head to generate the segmentation output  $S$ .




Figure 3: (a) GUI for Gaze data collection, (b) ground-truth segmentation labels, (c) gaze map of medical expert, (d) estimated gaze generated by SAN, (e) gaze map of non-expert.

### 3.3 Loss Functions and Model Training

All the components of our framework can be trained using one hybrid loss function. As for the segmentation, denoting the segmentation ground-truth as  $S^{gt}$ , we use the dice loss function [9, 37] based on the previous success of brain tumor segmentation:

$$\mathcal{L}_{seg} = \frac{1}{N} \sum_{n=1}^N \frac{2 \times S_n \cap S_n^{gt}}{S_n^2 + (S_n^{gt})^2}, \quad (3)$$

where  $N$  is the number of output channels.

The task of gaze estimation is trained using binary cross entropy loss [45]:

$$\mathcal{L}_g = -\frac{1}{P} \sum_{i=1}^P (G_i^{gt} \cdot \log(G_i) + (1 - G_i^{gt}) \cdot \log(1 - G_i)), \quad (4)$$

where  $P$  is the total volume, and  $G^{gt}$  is the measured expert's gaze positions obtained as explained in Section 4.1.

The joint loss function is a weighted sum of the loss functions:  $\mathcal{L} = w_1 \mathcal{L}_{seg} + w_2 \mathcal{L}_g$ . For the task of chest X-ray classification, we apply the typical cross-entropy loss to replace  $\mathcal{L}_{seg}$ .

## 4 Experiment

In this section, we validate our proposed framework on two datasets: (1) brain tumor segmentation on the BraTS 2020 dataset [4, 2, 41] and (2) disease classification on the MIMIC-CXR-gaze dataset [25].

### 4.1 Brain tumor segmentation task

We conduct our experiments on the training set and validation set of the BraTS 2020 dataset. The dataset contains 3D MRI data of four modalities (T1, T1ce, T2, and FLAIR), each with shape  $240 \times 240 \times 155$ . This dataset also provides the survival state (short, medium, and long) of each patient. The size of the training set and validation set is 369 and 125, respectively. The targets of segmentation are three regions namely Enhancing Tumor (ET), Whole Tumor (WT), and Tumor Core (TC). The performance of 5-fold cross-validation on the training set and the performance on the validation set assessed by the online evaluation server are both used to evaluate our method. We report the dice score as the evaluation metric.

**Gaze data collection** Since BraTS dataset does not provide human gaze data, we collect gaze data on its training set using an eyeTribe eye-tracker<sup>1</sup>. We also designed a GUI as shown in Fig. 3 to facilitate the gaze collection. This GUI is placed in a fixed position of the screen

<sup>1</sup><https://theeyetribe.com/>

while the timestamps of the user’s actions are automatically recorded. All the buttons in the GUI have corresponding shortcut keys to avoid looking at them during data collection. During data collection, we ask the user to diagnose the severeness of each patient and find the tumor regions from the MRI images in the GUI similar to the daily screening.

Among all 369 data samples in the BraTS training set, we collect a total of 180 samples with gaze data from two-board certificated radiologists. The two doctors have 4 years and 2 years of professional experience, respectively. On average, each data sample is checked by the doctor for 38 seconds, and we use the code from [15] to filter out the saccade points and keep only the fixation points. For comparison, we also collect the gaze data of these samples from two non-experts, however, they are not used in the training. The ground truth for gaze estimation  $G^{gt}$  is generated by adding a gaussian on the fixation points with kernel size 10, following the previous work [15]. An example is shown in Fig. 3, where (c) is the processed gaze map of an expert, and (e) is that of a non-expert. Note that our framework does not require human gaze data during inference.

**Model training** We empirically set  $w_1 = 1, w_2 = 0.5$  in all of the experiments. The model is trained for 210 epochs using Ranger as the optimizer [48, 49], with an initial learning rate of  $1e-4$ , and reduced by a cosine decay after 100 epochs. Including data augmentation, we keep other training and inference details the same as *pipeline A* in [13], except that we do not use test time augmentation during inference.

**Experiments** We first show the experiment of 3D tumor segmentation. As for baseline, we use the backbone encoder alone (**Backbone**) without adding our framework. Here we mainly choose to use the backbone of one of the top solutions [13] from the BraTS 2020 challenge, since we experimentally verify that this backbone works better on small data compared with the state of the art [20]. We also show that our proposed method can well cooperate with multiple backbones in the later part of this section. The baseline **ours** indicates the backbone is equipped with the proposed Auxiliary Attention Block and the Selective Attention Network. To better evaluate the effectiveness of the proposed framework, since there is no previous work that uses gaze to enhance 3D segmentation task to the best of our knowledge, we construct another baseline, in which we add a new segmentation head with 1 channel output for gaze estimation, and use the estimated gaze as soft attention map during inference [29]. We denote this baseline as **+gaze\***. To validate the effectiveness of the proposed AAB, we use a baseline **+SAN\***, indicating that the information of SAN is not processed by AAB, but is fused to the backbone network by direct concatenation. To validate the improvement brought by gaze data, we use **ours-gaze** baseline, indicating that the SAN is trained also with ground-truth segmentation mask, but not gaze data.

We also tested whether the framework can benefit from multi-task learning, by adding a classification head for survival state prediction, and denote this as **ours+surv**. To evaluate the effectiveness of our proposed framework using a smaller amount of training data, we conduct experiments using a part of training data that have collected gaze data under different proportions: 0.2, 0.3, 0.5, and 0.7. The experiment using a full set of training data, and the experiment using only the subset with gaze data, can be found in the supplementary materials.

From Table 1, the performance increase of **ours** compared with the backbone clearly demonstrates the advantage of our proposed framework when using small training data. From the comparison of **+gaze\*** and **ours** baselines, we can see that our proposed SAN and AAB can better leverage the auxiliary information provided by human gaze. The fact that **ours** outperforms **ours-gaze** also demonstrates the human expert’s gaze can provide complementary information other than segmentation masks, for example, experts’ gaze would



Mean (std)	Enhancing Tumor				Whole Tumor				Tumor Core			
	0.2	0.3	0.5	0.7	0.2	0.3	0.5	0.7	0.2	0.3	0.5	0.7
Data ratio												
Backbone	72.79 (0.022)	74.57 (0.019)	76.32 (0.027)	77.78 (0.007)	85.41 (0.018)	85.70 (0.016)	85.82 (0.022)	88.82 (0.006)	80.11 (0.017)	81.20 (0.017)	82.45 (0.012)	84.24 (0.016)
+ gaze*	72.97 (0.015)	74.24 (0.019)	75.77 (0.014)	78.01 (0.024)	83.42 (0.038)	85.22 (0.023)	86.01 (0.016)	88.71 (0.020)	79.97 (0.020)	81.14 (0.009)	82.03 (0.015)	84.65 (0.012)
+SAN*	71.84 (0.025)	72.06 (0.017)	74.37 (0.018)	75.90 (0.030)	82.21 (0.035)	84.58 (0.027)	85.10 (0.024)	86.56 (0.023)	77.74 (0.016)	79.66 (0.013)	80.73 (0.018)	82.34 (0.015)
ours-gaze	72.07 (0.020)	73.84 (0.016)	76.50 (0.017)	78.02 (0.023)	84.37 (0.021)	85.28 (0.028)	86.70 (0.022)	87.67 (0.015)	79.83 (0.017)	81.11 (0.010)	82.50 (0.018)	84.35 (0.014)
ours	74.69 (0.022)	76.46 (0.007)	79.05 (0.025)	<b>81.43</b> (0.021)	85.05 (0.019)	85.44 (0.024)	89.14 (0.013)	89.13 (0.010)	82.51 (0.009)	82.27 (0.008)	85.46 (0.012)	85.42 (0.020)
ours+surv	<b>76.69</b> (0.007)	<b>77.72</b> (0.012)	<b>79.10</b> (0.009)	80.36 (0.028)	<b>86.76</b> (0.013)	<b>87.29</b> (0.022)	<b>89.27</b> (0.009)	<b>90.29</b> (0.010)	<b>82.88</b> (0.005)	<b>84.07</b> (0.009)	<b>85.67</b> (0.010)	<b>86.83</b> (0.011)

Table 1: Training set cross validation result (Dice score %) using part of training data. The numbers in parenthesis are the std values.




Figure 4: Qualitative comparison of different methods. (a) Performance of Backbone with our framework and multi-task learning (ours+surv), when using different amount of training data. (b) Performance of the backbone equipped with SAN and AAB (ours), and further multi-task training (ours+surv), using 20% of training data.

give more focus on the boundaries. The effectiveness of the proposed AAB is validated by the comparison between **ours** and **+SAN\***. Results on the validation set are in the supplementary materials, showing similar performances.

A qualitative example of the segmentation result is shown in Fig. 4. The first row shows the performance of the **ours+surv** framework, when trained using different proportions of training data. The second row shows a comparison of different methods when using 20% training data. From both rows, we can observe performance increase from left to right, which validates the benefit of our proposed framework. Fig. 3 (d) shows the gaze estimation output from the SAN, which is highly consistent with the expert’s gaze in Fig. 3 (c).

To further test the effect of gaze, we compare the performance of the network trained using expert gaze (EG) and non-expert gaze (n-EG) and show the results in Table 2. The

Mean (std)	Enhancing Tumor		Whole Tumor		Tumor Core	
	0.2	0.5	0.2	0.5	0.2	0.5
EG	<b>74.69</b> (0.022)	<b>79.05</b> (0.025)	<b>85.05</b> (0.019)	<b>89.14</b> (0.013)	<b>82.51</b> (0.009)	<b>85.46</b> (0.012)
n-EG	72.65(0.020)	76.02(0.011)	84.77(0.013)	84.96(0.014)	78.65(0.019)	81.13(0.006)

Table 2: Cross validation results on training set using expert gaze (EG) or non-expert gaze (n-EG) during training (Dice score %). The numbers in parenthesis are the std values.



Type	ET				WT				TC			
	0.2		0.5		0.2		0.5		0.2		0.5	
Data ratio												
Metric	Dice	H-95	Dice	H-95	Dice	H-95	Dice	H-95	Dice	H-95	Dice	H-95
Unet [13]	72.79	34.05	76.32	30.02	85.41	10.98	85.82	9.52	80.11	10.51	82.45	7.78
nnUnet [20]	71.86	38.52	73.74	28.42	83.79	13.53	85.67	9.90	79.87	17.07	80.78	16.21
DMFNet [9]	70.85	36.02	72.01	33.81	82.67	16.04	85.12	10.36	76.89	16.96	80.52	9.08
Unet [13]+ours	<b>74.69</b>	<b>33.61</b>	<b>79.05</b>	<b>21.28</b>	85.05	<b>10.61</b>	<b>89.14</b>	7.54	<b>82.51</b>	<b>6.60</b>	<b>85.46</b>	8.21
nnUnet [20]+ours	72.80	35.77	77.50	26.65	<b>86.55</b>	15.11	88.10	7.12	81.39	12.47	82.85	9.72
DMFNet [9]+ours	73.20	34.03	74.99	33.42	86.12	10.89	88.67	<b>7.07</b>	80.07	10.51	83.38	<b>7.94</b>
Unet+CBAM [46]	69.96	44.19	78.12	28.83	83.87	12.53	86.91	10.82	77.34	25.44	85.03	9.24
Unet+TASN [52]	70.03	44.20	73.70	36.19	83.61	17.07	83.99	13.53	79.81	11.68	81.96	17.46

**Table 3:** Model cooperation with different backbones (top) and different attention modules (bottom). We experiment with Unet from [13], nnUnet [20] and DMFNet [9] as backbones. We also replace our proposed framework with existing attention modules CBAM [46] and TASN [52] to test the effectiveness. Dice score (Dice) and the mean of Hausdorff95 distance (H-95) are shown as evaluation metrics.

performance of the framework using expert gaze consistently outperforms the one using non-expert gaze. As shown in Fig. 3(e), non-expert gaze is more of a distraction for our framework as reading MRI data requires professional knowledge.

It is worth noting that our method can work with different backbone models and improve their performance using gaze information as demonstrated in our experiment. We test Unet from [13], nnUnet [20] and DMFNet [9]. nnUnet is the current state-of-the-art method for brain tumor segmentation, while DMFNet is an efficient backbone that could work real-time. As shown in Table 3, with the help of our proposed modules, the performance of all backbones are increased.

For further testing the efficacy of the proposed method, we replace our method with several current common spatial attention methods. We choose CBAM [46] and TASN [52] for comparison, and supervise the attention map with ground truth gaze map. For TASN we supervise both the "structure attention" and "detailed attention". The results can be found in the last two rows of Table 3. Clearly, we can see that directly apply existing spatial attention methods cannot help the backbone model under most circumstances.

## 4.2 Chest disease classification task

We also conduct experiments on the MIMIC-CXR-gaze dataset [25] for normal versus abnormal chest X-ray classification. This dataset contains 1083 CXR images with gaze collected from an American Board of Radiology certified radiologist. We use the pre-analyzed fixations to generate ground-truth of gaze map by putting an isotropic gaussian on each fixation point. Following [18], we use AUROC scores to evaluate the model performance on the classification of Atelectasis, Cardiomegaly, Consolidation, Edema, and Pleural Effusion. Other details including dataset split and model training can be found in the supplementary material. Note that the training is not conducted on the full MIMIC-CXR dataset [25].

In this experiment, following [25] we use the encoder part of EfficientNet-b0 [40] as the backbone encoder. Our experiments show that this backbone performs better than ResNet [11] and MobileNet [44] on small training data. The classification head output after sigmoid activation contains 5 channels, indicating the existence of each disease. We add two baselines for comparison. (1) Similarly to the brain tumor segmentation task, we add the **+gaze\*** baseline that directly adds a gaze prediction head on top of the backbone encoder, without

Ratio	Method	Average	Atelectasis	Cardiomegaly	Consolidation	Edema	Pleural Effusion
0.7	Backbone	55.14	61.06	57.72	32.56	62.83	61.55
	+gaze*	55.07	59.99	52.52	34.65	67.88	60.31
	+gaze loss [10]	55.30	59.25	55.22	33.67	69.14	59.21
	ours	<b>59.26</b>	65.68	61.31	33.23	74.14	61.94
0.5	Backbone	50.43	52.30	54.47	28.06	66.40	50.94
	+gaze*	54.67	52.24	55.58	34.92	67.66	62.95
	+gaze loss [10]	54.05	53.97	55.24	34.54	64.38	62.11
	ours	<b>55.52</b>	61.85	55.86	31.35	67.70	60.86
0.3	Backbone	46.60	42.02	49.77	36.32	59.87	46.01
	+gaze*	49.64	41.19	48.75	32.20	60.98	65.06
	+gaze loss [10]	51.43	48.37	48.77	33.25	63.98	62.70
	ours	<b>55.23</b>	56.92	53.63	32.40	69.10	64.10
0.2	Backbone	42.59	38.61	48.09	37.88	41.92	46.43
	+gaze*	43.96	42.22	45.80	36.08	41.75	53.93
	+gaze loss [10]	45.68	46.51	46.70	35.64	47.16	52.38
	ours	<b>48.11</b>	52.71	48.25	29.86	53.15	56.59

Table 4: Results on the MIMIC-CXR-gaze dataset. AUROC score in percentage is used as the evaluation metric. Results are the average of the best 30 epochs.

using the proposed SAN and AAB. (2) We modify the recent classification work with gaze enhancement [10] and use its gaze loss for supervising the gaze estimation head.

Quantitative results are shown in Table 4. Despite "Consolidation", where the backbone results surprisingly performs better when the amount of training data decreases, for all other four classes and the averaged performance, using our framework with gaze information, the performance can be significantly better when using limited training data. The comparison between our proposed framework and two baselines using gaze strongly proves that the performance improvement comes from the design of our framework, but not just from the gaze data. However, we can also see that in situations where the amount of data is limited, using non-invasively collected gaze as additional training data can benefit the model learning. We believe this shows the potential of using human gaze data in real-world clinical tasks.

We also test the performance of different backbones (EfficientNet [10], MobileNet-v2 [4], ResNet-18 [10]) in our framework. The table of quantitative results can be found in the supplementary material. Among the three backbones, EfficientNet and MobileNet-v2 perform better than ResNet-18, which is reasonable since the amount of training data is limited, and ResNet is more complicated and requires more data for training when compared with the other two backbones.

## 5 Conclusion

In this work, we explored using human gaze to enhance model performance, especially on limited training data. We propose a novel framework that could leverage human gaze information to enhance multiple medical image analysis tasks such as brain tumor segmentation and chest disease classification. On both the BraTS 2020 dataset and the MIMIC-CXR-gaze dataset, our method achieves superior performance using a limited number of training samples when compared to existing methods. We also collected the real human gaze of both medical experts and non-experts on the BraTS 2020 dataset. Our experiment shows the potential of gaze and eye-tracker for real-world medical applications.

## Acknowledgement

This work was supported by JST under grant Moonshot R&D Grant Number JPMJMS2011, ACT-X Grant Number JPMJAX190D, JST AIP Acceleration Research Grant Number JPMJCR20U1 and JSPS KAKENHI Grant Number JP20H04205, Japan.

## References

- [1] Spyridon Bakas, Hamed Akbari, Aristeidis Sotiras, Michel Bilello, Martin Rozycki, Justin S Kirby, John B Freymann, Keyvan Farahani, and Christos Davatzikos. Advancing the cancer genome atlas glioma mri collections with expert segmentation labels and radiomic features. *Scientific data*, 4(1): 1–13, 2017.
- [2] Spyridon Bakas, Mauricio Reyes, Andras Jakab, Stefan Bauer, Markus Rempfler, Alessandro Crimi, Russell Takeshi Shinohara, Christoph Berger, Sung Min Ha, Martin Rozycki, et al. Identifying the best machine learning algorithms for brain tumor segmentation, progression assessment, and overall survival prediction in the brats challenge. *arXiv preprint arXiv:1811.02629*, 2018.
- [3] Donald Eric Broadbent. *Perception and communication*. Elsevier, 2013.
- [4] Adria Casamitjana, Santi Puch, Asier Aduriz, and Verónica Vilaplana. 3d convolutional neural networks for brain tumor segmentation: a comparison of multi-resolution architectures. In *International Workshop on Brainlesion: Glioma, Multiple Sclerosis, Stroke and Traumatic Brain Injuries*, pages 150–161. Springer, 2016.
- [5] Chen Chen, Xiaopeng Liu, Meng Ding, Junfeng Zheng, and Jiangyun Li. 3d dilated multi-fiber network for real-time brain tumor segmentation in mri. In *International Conference on Medical Image Computing and Computer Assisted Intervention*, 2019.
- [6] Wei Chen, Boqiang Liu, Suting Peng, Jiawei Sun, and Xu Qiao. S3d-unet: separable 3d u-net for brain tumor segmentation. In *International MICCAI Brainlesion Workshop*, pages 358–368. Springer, 2018.
- [7] Özgün Çiçek, Ahmed Abdulkadir, Soeren S. Lienkamp, Thomas Brox, and Olaf Ronneberger. 3d u-net: Learning dense volumetric segmentation from sparse annotation. In Sebastien Ourselin, Leo Joskowicz, Mert R. Sabuncu, Gozde Unal, and William Wells, editors, *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2016*, pages 424–432, Cham, 2016. Springer International Publishing.
- [8] Özgün Çiçek, Ahmed Abdulkadir, Soeren S Lienkamp, Thomas Brox, and Olaf Ronneberger. 3d u-net: learning dense volumetric segmentation from sparse annotation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 424–432. Springer, 2016.
- [9] Michal Drozdal, Eugene Vorontsov, Gabriel Chartrand, Samuel Kadoury, and Chris Pal. The importance of skip connections in biomedical image segmentation. In *Deep learning and data labeling for medical applications*, pages 179–187. Springer, 2016.
- [10] Alireza Fathi, Yin Li, and James M Rehg. Learning to recognize daily actions using gaze. In *European Conference on Computer Vision*, pages 314–327. Springer, 2012.
- [11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016.

- [12] Sen He, Hamed R Tavakoli, Ali Borji, and Nicolas Pugeault. Human attention in image captioning: Dataset and analysis. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 8529–8538, 2019.
- [13] Theophraste Henry, Alexandre Carre, Marvin Lrousseau, Theo Estienne, Charlotte Robert, Nikos Paragios, and Eric Deutsch. Top 10 brats 2020 challenge solution: Brain tumor segmentation with self-ensembled, deeply-supervised 3d-unet like neural networks. *arXiv preprint arXiv:2011.01045*, 2020.
- [14] Andrew G Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*, 2017.
- [15] Yifei Huang, Minjie Cai, Zhenqiang Li, and Yoichi Sato. Predicting gaze in egocentric video by learning task-dependent attention transition. In *European Conference on Computer Vision*, pages 754–769, 2018.
- [16] Yifei Huang, Minjie Cai, Zhenqiang Li, Feng Lu, and Yoichi Sato. Mutual context network for jointly estimating egocentric gaze and action. *IEEE Transactions on Image Processing*, 29: 7795–7806, 2020.
- [17] Yifei Huang, Minjie Cai, and Yoichi Sato. An ego-vision system for discovering human joint attention. *IEEE Transactions on Human-Machine Systems*, 50(4):306–316, 2020.
- [18] Jeremy Irvin, Pranav Rajpurkar, Michael Ko, Yifan Yu, Silvana Ciurea-Ilcus, Chris Chute, Henrik Marklund, Behzad Haghgoo, Robyn Ball, Katie Shpanskaya, et al. Chexpert: A large chest radiograph dataset with uncertainty labels and expert comparison. In *AAAI Conference on Artificial Intelligence*, pages 590–597, 2019.
- [19] Fabian Isensee, Philipp Kickingereder, Wolfgang Wick, Martin Bendszus, and Klaus H Maier-Hein. Brain tumor segmentation and radiomics survival prediction: Contribution to the brats 2017 challenge. In *International MICCAI Brainlesion Workshop*, pages 287–297. Springer, 2017.
- [20] Fabian Isensee, Paul F Jaeger, Simon AA Kohl, Jens Petersen, and Klaus H Maier-Hein. nnu-net: a self-configuring method for deep learning-based biomedical image segmentation. *Nature Methods*, 18(2):203–211, 2021.
- [21] William James. *The principles of psychology*. New York: Dover, 1890.
- [22] Alistair EW Johnson, Tom J Pollard, Nathaniel R Greenbaum, Matthew P Lungren, Chihying Deng, Yifan Peng, Zhiyong Lu, Roger G Mark, Seth J Berkowitz, and Steven Horng. Mimic-cxr-jpg, a large publicly available database of labeled chest radiographs. *arXiv preprint arXiv:1901.07042*, 2019.
- [23] Konstantinos Kamnitsas, Enzo Ferrante, Sarah Parisot, Christian Ledig, Aditya V Nori, Antonio Criminisi, Daniel Rueckert, and Ben Glocker. Deepmedic for brain tumor segmentation. In *International workshop on Brainlesion: Glioma, multiple sclerosis, stroke and traumatic brain injuries*, pages 138–149. Springer, 2016.
- [24] Georgios Kapidis, Ronald Poppe, Elsbeth van Dam, Lucas Noldus, and Remco Veltkamp. Multitask learning to improve egocentric action recognition. In *International Conference on Computer Vision Workshops*, 2019.
- [25] Alexandros Karargyris, Satyananda Kashyap, Ismini Lourentzou, Joy T Wu, Arjun Sharma, Matthew Tong, Shafiq Abedin, David Beymer, Vandana Mukherjee, Elizabeth A Krupinski, et al. Creation and validation of a chest x-ray dataset with eye-tracking and report dictation for ai development. *Scientific data*, 8(1):1–18, 2021.

- [26] Baris Kayalibay, Grady Jensen, and Patrick van der Smagt. Cnn-based segmentation of medical imaging data. *arXiv preprint arXiv:1701.03056*, 2017.
- [27] Yin Li, Alireza Fathi, and James M Rehg. Learning to predict gaze in egocentric video. In *IEEE International Conference on Computer Vision*, pages 3216–3223, 2013.
- [28] Yin Li, Zhefan Ye, and James M Rehg. Delving into egocentric actions. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 287–295, 2015.
- [29] Yin Li, Miao Liu, and Jame Rehg. In the eye of the beholder: Gaze and actions in first person video. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021.
- [30] Yang Liu, Lei Zhou, Xiao Bai, Yifei Huang, Lin Gu, Jun Zhou, and Tatsuya Harada. Goal-oriented gaze estimation for zero-shot learning. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2021.
- [31] Bjoern H Menze, Andras Jakab, Stefan Bauer, Jayashree Kalpathy-Cramer, Keyvan Farahani, Justin Kirby, Yuliya Burren, Nicole Porz, Johannes Slotboom, Roland Wiest, et al. The multimodal brain tumor image segmentation benchmark (brats). *IEEE transactions on medical imaging*, 34(10):1993–2024, 2014.
- [32] Fausto Milletari, Nassir Navab, and Seyed-Ahmad Ahmadi. V-net: Fully convolutional neural networks for volumetric medical image segmentation. In *2016 fourth international conference on 3D vision (3DV)*, pages 565–571. IEEE, 2016.
- [33] Kyle Min and Jason J Corso. Integrating human gaze into attention for egocentric activity recognition. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 1069–1078, 2021.
- [34] Yuval Nirkin, Lior Wolf, and Tal Hassner. Hyperseg: Patch-wise hypernetwork for real-time semantic segmentation. *arXiv preprint arXiv:2012.11582*, 2020.
- [35] Arijit Patra, Yifan Cai, Pierre Chatelain, Harshita Sharma, Lior Drukker, Aris T. Papageorgiou, and J. Alison Noble. Efficient ultrasound image analysis models with sonographer gaze assisted distillation. In Dinggang Shen, Tianming Liu, Terry M. Peters, Lawrence H. Staib, Caroline Essert, Sean Zhou, Pew-Thian Yap, and Ali Khan, editors, *International Conference on Medical Image Computing and Computer Assisted Intervention*, pages 394–402, 2019.
- [36] Michael I Posner and Steven E Petersen. The attention system of the human brain. *Annual review of neuroscience*, 13(1):25–42, 1990.
- [37] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.
- [38] Joseph N Stember, Haydar Celik, E Krupinski, Peter D Chang, Simukayi Mutasa, Bradford J Wood, A Lignelli, G Moonis, LH Schwartz, Sachin Jambawalikar, et al. Eye tracking for deep learning segmentation using convolutional neural networks. *Journal of digital imaging*, 32(4): 597–604, 2019.
- [39] Joseph N Stember, Haydar Celik, David Gutman, Nathaniel Swinburne, Robert Young, Sarah Eskreis-Winkler, Andrei Holodny, Sachin Jambawalikar, Bradford J Wood, Peter D Chang, et al. Integrating eye tracking and speech recognition accurately annotates mr brain images for deep learning: Proof of principle. *Radiology: Artificial Intelligence*, 3(1):e200047, 2020.
- [40] Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International Conference on Machine Learning*, pages 6105–6114, 2019.

- [41] Sabine Kastner Ungerleider and Leslie G. Mechanisms of visual attention in the human cortex. *Annual review of neuroscience*, 23(1):315–341, 2000.
- [42] Arun Balajee Vasudevan, Dengxin Dai, and Luc Van Gool. Object referring in videos with language and human gaze. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4129–4138, 2018.
- [43] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in Neural Information Processing Systems*, pages 6000–6010, 2017.
- [44] Xiaosong Wang, Yifan Peng, Le Lu, Zhiyong Lu, Mohammadhadi Bagheri, and Ronald M. Summers. Chestx-ray8: Hospital-scale chest x-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
- [45] Xiaosong Wang, Yifan Peng, Le Lu, Zhiyong Lu, and Ronald M. Summers. Tienet: Text-image embedding network for common thorax disease classification and reporting in chest x-rays. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2018.
- [46] Sanghyun Woo, Jongchan Park, Joon-Young Lee, and In So Kweon. Cbam: Convolutional block attention module. In *Proceedings of the European conference on computer vision (ECCV)*, pages 3–19, 2018.
- [47] Yuxin Wu and Kaiming He. Group normalization. In *European Conference on Computer Vision*, pages 3–19, 2018.
- [48] Hongwei Yong, Jianqiang Huang, Xiansheng Hua, and Lei Zhang. Gradient centralization: A new optimization technique for deep neural networks. In *European Conference on Computer Vision*, pages 635–652, 2020.
- [49] Michael R Zhang, James Lucas, Geoffrey Hinton, and Jimmy Ba. Lookahead optimizer: k steps forward, 1 step back. *arXiv preprint arXiv:1907.08610*, 2019.
- [50] Ruohan Zhang, Akanksha Saran, Bo Liu, Yifeng Zhu, Sihang Guo, Scott Niekum, Dana Ballard, and Mary Hayhoe. Human gaze assisted artificial intelligence: A review. In *International Joint Conference on Artificial Intelligence*, pages 4951–4958, 2020.
- [51] Yixiao Zhang, Xiaosong Wang, Ziyue Xu, Qihang Yu, Alan Yuille, and Daguang Xu. When radiology report generation meets knowledge graph. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34:12910–12917, 2020.
- [52] Heliang Zheng, Jianlong Fu, Zheng-Jun Zha, and Jiebo Luo. Looking for the devil in the details: Learning trilinear attention sampling network for fine-grained image recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5012–5021, 2019.