

Dr. Tongue: Sign-Oriented Multi-label Detection for Remote Tongue Diagnosis

Yiliang Chen¹, Steven SC Ho¹, Cheng Xu^{1*}, Yao Jie Xie¹, Wing-Fai Yeung¹,
Shengfeng He^{2*}, Jing Qin¹

¹School of Nursing, The Hong Kong Polytechnic University, Hong Kong

²Singapore Management University, Singapore

{yiliang.chen, steven-sc.ho}@connect.polyu.hk, cschengxu@gmail.com, {grace.yj.xie, jerry-wf.yeung, harry.qin}@polyu.edu.hk, shengfenghe@smu.edu.sg

Abstract


Tongue diagnosis is a vital tool in Western and Traditional Chinese Medicine, providing key insights into a patient's health by analyzing tongue attributes. The COVID-19 pandemic has heightened the need for accurate remote medical assessments, emphasizing the importance of precise tongue attribute recognition via telehealth. To address this, we propose a Sign-Oriented multi-label Attributes Detection framework. Our approach begins with an adaptive tongue feature extraction module that standardizes tongue images and mitigates environmental factors. This is followed by a Sign-oriented Network (SignNet) that identifies specific tongue attributes, emulating the diagnostic process of experienced practitioners and enabling comprehensive health evaluations. To validate our methodology, we developed an extensive tongue image dataset specifically designed for telemedicine. Unlike existing datasets, ours is tailored for remote diagnosis, with a comprehensive set of attribute labels. This dataset will be openly available, providing a valuable resource for research. Initial tests have shown improved accuracy in detecting various tongue attributes, highlighting our framework's potential as an essential tool for remote medical assessments.

Resources — <https://github.com/tonguedx/tonguedx>.

Introduction

The global COVID-19 pandemic has accelerated the shift towards telemedicine, highlighting the critical need for reliable and efficient remote diagnosis methods (Kichloo et al. 2020; Ohannessian et al. 2020; Portnoy, Waller, and Elliott 2020). Within this context, tongue diagnosis, a cornerstone of Traditional Chinese Medicine, has gained renewed interest due to its potential to provide valuable insights into a patient's health through the analysis of tongue attributes (Heo et al. 2022; Hu et al. 2015; Lee et al. 2016; Zhao et al. 2013; Jiang et al. 2012; Lo et al. 2012). However, effectively incorporating tongue diagnosis into telehealth platforms requires the development of sophisticated automated recognition technologies capable of overcoming challenges such as inconsistent image quality and environmental conditions, to ensure precise and reliable diagnostics.

*Corresponding authors: Cheng Xu and Shengfeng He.
Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.



(a) Low-resolution tongue images (b) Complex-lighting tongue images

Figure 1: Challenges in tongue diagnosis imaging for telehealth: (a) low resolution and (b) complex lighting conditions.

The integration of tongue diagnosis into telehealth frameworks faces significant challenges, particularly in interpreting remotely captured images. Low-resolution cameras and beautification features on smartphones can compromise the quality of tongue images, which is critical for accurate diagnosis (Figure 1(a)). Additionally, variable lighting conditions can alter the perceived color and texture of the tongue, leading to potential misdiagnoses (Figure 1(b)). Images taken under different lighting conditions from the same individual can show marked differences in color and texture, complicating the diagnostic process.

To address these challenges, it is crucial to develop a comprehensive tongue image dataset. Despite numerous studies (Li et al. 2018; Heo et al. 2022; Hu et al. 2019; Wang et al. 2013; Wang, Zhang, and Cai 2014; Jiang et al. 2022), there is a lack of publicly available datasets that reflect the diverse conditions of telehealth settings. Existing datasets often rely on standardized equipment like the TFDA-1 tongue diagnosis device, which does not accurately represent telemedicine scenarios (Shi et al. 2021; Jiang et al. 2021b,a). Moreover, many datasets lack detailed annotations and do not cover a wide range of tongue attributes necessary for comprehensive diagnosis (Wang, Zhang, and Cai 2014; Jiang et al. 2012).

To overcome these deficiencies, we propose to create a more comprehensive dataset that extends beyond mere tongue imagery. This dataset will include images captured in telemedicine scenarios, along with demographic details and comprehensive descriptions of tongue surface attributes, particularly tongue fur. Images will be collected using various cameras under different environmental conditions to en-

sure representative diversity. Our approach shifts the focus from identifying a single attribute to the concurrent recognition of multiple tongue surface attributes.

To address the challenges in tongue image analysis, such as environmental interference and posture variations, we propose a novel methodology comprising two main components: The Adaptive Tongue Feature Extraction module and the Sign-Oriented Attribute Detection network (SignNet). The Adaptive Tongue Feature Extraction module employs GroundingDINO (Liu et al. 2023) and SAM models (Kirillov et al. 2023; Ke et al. 2024) to precisely segment the tongue from complex backgrounds and align it to a standardized orientation. This standardization reduces the impact of environmental and postural variations, similar to object orientation methods (Fu et al. 2008), simplifying subsequent analysis.

Building on the standardized input, SignNet employs a multi-branch architecture to capture inter-attribute relationships effectively. The network consists of whole tongue, tongue body, and tongue edge sections, aligned with the attribute-sign relationships outlined in Table 2. To enhance inter-attribute learning, the design incorporates several key mechanisms. The tongue body and edge sections integrate regional information, enabling the detection of location-specific features. The whole tongue branch extracts and integrates color and fur information, reorganizing color attributes into four categories and simplifying fur characteristics to presence or absence. This structured information is processed through an attention mechanism, which facilitates learning inter-relationships among various features across branches.

The attention mechanism further fuses information from different branches, enabling the network to recognize patterns such as the association between specific regions, colors, and fur characteristics. For example, it can detect that a pale tongue edge with no fur corresponds to particular diagnostic patterns. This comprehensive design allows SignNet to capture the complex interrelationships between tongue attributes, considering their spatial distributions and mutual influences. Consequently, SignNet provides reliable support for remote tongue diagnosis, closely simulating the diagnostic practices of TCM practitioners. By learning these intricate relationships, our method achieves more accurate and comprehensive tongue attribute recognition, leading to more reliable and precise health assessments.

In summary, our contributions are threefold:

- We create a diverse and comprehensive tongue image dataset tailored for telemedicine, including telehealth scenario images, demographic information, and multi-label annotations. This publicly accessible dataset addresses the common issue of restricted dataset availability in this field.
- We introduce a novel Sign-Oriented Attribute Detection framework that integrates adaptive tongue feature extraction with a Sign-oriented recognition Network (SignNet), enhancing the accuracy and robustness of automated tongue diagnosis across varied environmental conditions.

- We demonstrate the framework’s adaptability to various imaging conditions, validating its efficacy and robustness in telehealth applications through experimental results.

Related Work

Zero-shot Learning in Vision-Language Models. Zero-shot learning is a key feature in Vision-Language Models (VLMs), enabling recognition of new concepts without training (Wang et al. 2022; Xian et al. 2018). CLIP (Radford et al. 2021; Yu et al. 2024) pioneered this field through natural language supervision. Models like ALIGN (Jia et al. 2021) and Florence (Yuan et al. 2023) further enhanced performance and scalability, advancing tasks like object detection (Bansal et al. 2018; Zhu, Wang, and Saligrama 2019) and segmentation (Zheng et al. 2021; Bucher et al. 2019). Unified architectures such as BLIP (Li et al. 2022) and SimVLM (Wang et al. 2021) show promise across vision-language tasks.

Recent advancements have pushed the boundaries of zero-shot capabilities in specialized domains. GroundingDINO (Liu et al. 2023) extended open-set object detection by using natural language prompts, while Segment Anything Models (SAM) (Kirillov et al. 2023; Ke et al. 2024) advanced zero-shot segmentation with a promptable interface that generates masks for arbitrary objects. These developments underscore the growing flexibility and generalization power of VLMs in novel scenarios. In our method, we leverage GroundingDINO and SAM to process tongue images within complex backgrounds, enhancing the robustness of our diagnostic framework.

Tongue Images Dataset and Studies. The specialized domain of automated tongue diagnosis has received limited attention, largely due to its roots in traditional Eastern medicine rather than mainstream Western practices. A significant early contribution was made by Wang et al. (Wang et al. 2013), who developed the first publicly accessible dataset for tongue diagnosis. However, this dataset was limited in scope, offering only basic demographic details like age and gender alongside the tongue images. Moreover, it has since become unavailable, creating a resource gap for researchers in this under-explored field.

Wang et al. also proposed several traditional machine learning techniques (Pang et al. 2004; Zuo et al. 2004; Zhi et al. 2007; Huang et al. 2010) to explore the relationship between abnormal tongue appearances and various diseases, focusing on aspects like color and shape analysis. As deep learning evolved (Xu et al. 2022; Jiang et al. 2024), research in tongue image analysis began to address more complex attributes. Li et al. (Li et al. 2018) employed basic CNNs to recognize tooth-marked tongues, and another team (Yuan et al. 2023) used CNNs to investigate the connection between tongue images and stomach cancer, both as straightforward binary classification tasks.

More recently, Jiang et al. (Jiang et al. 2022) advanced the field by using deep learning to identify multi-label tongue annotations. Their images, collected using the TFDA-1 tongue diagnosis instrument, achieved high accuracy due to the controlled imaging conditions, unaffected by angle or

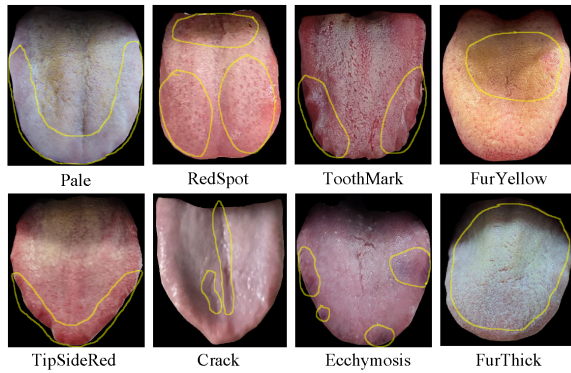


Figure 2: Cropped images with eight representative tongue attributes, with characteristics encircled in yellow.




Figure 3: Original samples of tongue images from our dataset.

lighting variations. However, while this method is effective, it is not ideal for real-world telehealth scenarios where conditions are less controlled. To address this, we implemented a Sign-Oriented Multi-label Attributes Detection framework in our experiments, aiming to replicate the real diagnostic process and improve accuracy in diverse and realistic telehealth environments.

The TongueDx Dataset

Our tongue image dataset comprises a diverse compilation of images captured in natural settings using various smartphone and laptop cameras. This results in a dataset with inherent variability in environmental conditions, image resolutions, and tongue capture angles. Some cameras may have slight beauty filters enabled by default, while others do not. Additionally, the dataset includes images with different facial inclusions; some feature the entire face, while others only show a portion of the face.

For the annotation of tongue bounding boxes, we employed the GroundingDINO (Liu et al. 2023) model, which can accurately localize tongues in complex backgrounds based on the prompt text “tongue”. While this model provides precise annotations, it requires significant computational time for inference. To address its computational limitations, we used GroundingDINO’s outputs as high-quality training data for a more lightweight network, significantly reducing detection time while maintaining accuracy.

We define classes for tongue images based on several representative tongue attributes, identified through consultations with five professional practitioners. We established eight fundamental tongue attributes: pale tongue color (Pale), reddish tip or edges (TipSideRed), red spots (RedSpot), ecchymosis (Ecchymosis), deep cracks

Tongue attributes	No _i	Yes _i	No _s	Yes _s
Pale	4485	624	4083	567
TipSideRed	2866	2243	2691	1959
RedSpot	2678	2431	2457	2193
Ecchymosis	4638	471	4219	431
Crack	835	4274	754	3896
ToothMark	2125	2984	1856	2794
FurThick	136	4973	119	4531
FurYellow	4276	833	3891	759

Table 1: Tongue Attributes Statistical Distribution of the Proposed Dataset. No_i denotes the number of images without the attribute, while Yes_i denotes the number of images with the attribute. Similarly, No_s and Yes_s indicate the number of subjects without and with the attribute, respectively.

(Crack), pronounced tooth marks (ToothMark), thick coating (FurThick), and yellow coating (FurYellow). Figure 2 illustrates these attributes with yellow circles highlighting areas of interest. Each attribute is binary-classified, where 1 indicates its presence. ‘Pale’ suggests a lighter tongue body color; ‘TipSideRed’ denotes tip redness; ‘RedSpot’ refers to small, red dots in the coating; ‘Ecchymosis’ indicates bruise-like discolorations; ‘Crack’ points out significant fissures in the central area; ‘ToothMark’ shows dental impressions along the edges; ‘FurThick’ suggests a thick coating (normally transparent); and ‘FurYellow’ denotes a yellowish coating. A tongue may exhibit multiple attributes simultaneously. These attributes were identified through in-person observations at a clinic, with corresponding images captured either at the clinic or participants’ homes. Three professional practitioners completed the final annotation of surface attributes.

Our dataset comprises a total of 5109 images collected from 4650 unique subjects. Each subject have one or two images captured from different distances or angles. This variability is considered to capture a more comprehensive representation of the subjects within the dataset. For the subjects in our study, we present the number of occurrences of the attributes in Table 1. In addition to these attributes, we have also labeled other demographic data such as age and gender for each subject, laying the groundwork for future research. We showcase several samples of our dataset in Figure 3.

Methodology

In this section, we elaborate the architecture of our model comprehensively. We start by outlining the problem formulation of the multi-label tongue attribute recognition task. Following this, we introduce the proposed Sign-Oriented framework as shown in Figure 5. This method begins with an Adaptive Tongue Feature Extraction (ATFE) module to standardize tongue images and mitigate environmental factors, followed by a sign-oriented network (SignNet) that combines various signs to identify specific tongue attributes. This approach emulates the diagnostic process of practitioners, enabling thorough patient assessment.




Figure 4: Two examples of the Tongue Image Upright Orientation Algorithm: Input image (Before Orientation) and the upright oriented tongue image (After Orientation).

Formulation

Let (X, B, M, T) denote the tongue image dataset, where X , B , M , and T represent images, bounding boxes, segmentation masks, and tongue attributes, respectively. Our approach consists of four steps: 1) Detection, where a lightweight network identifies the tongue’s position within each image $x \in X$, predicting its bounding box $b \in B$; 2) Segmentation, using the detected bounding box as a prompt for SAM to generate a precise segmentation mask $m \in M$; 3) Upright Orientation, transforming the segmented tongue image to ensure an upright orientation; and 4) Attribute identification, identifying the basic attributes of the tongue. Formally, we aim to construct a multi-stage framework $f_\theta : x_n \rightarrow (b_n, m_n, r_n, t_n)$, where θ , x_n , b_n , m_n , and r_n represent model parameters, the input image, the predicted bounding box, the segmentation mask, and the upright orientation transformation, respectively. $t_n \in \{0, 1\}^s$ is the binary label vector for attributes, with $s = |T|$ being the total number of possible attributes. Each image may exhibit multiple attributes simultaneously.

Adaptive Tongue Feature Extraction Module

Tongue Detection and Segmentation. To accurately capture the tongue body in complex image backgrounds, we employ a multi-step approach. Initially, we utilize GroundingDINO to generate precise ground truth bounding boxes for each tongue image in our dataset. By inputting the text prompt “tongue” and selecting the detection box with the highest confidence score, we obtained high-quality annotations. While GroundingDINO’s accuracy proved sufficient, its computational demands as a large model made it impractical for real-time applications. To address this efficiency issue, we opted for YOLOv5-MobileNetV3 (Howard et al. 2019; Jocher et al. 2022) as our detection network to identify tongue positions. Given the simplicity of this single-object localization task, we can easily achieve over 99% mAP performance. Subsequently, we input the predicted bounding box into MobileSAM (Zhang et al. 2023) for segmentation, adding a click prompt at the box’s center to enhance accuracy. Our observations showed minimal impact of different SAM variants on segmentation quality. Thus, considering the performance-efficiency trade-off, we choose MobileSAM for our implementation.

Upright Orientation of Tongue. To standardize the orientation of tongue images for mitigating complex environmental factors and ensuring consistent feature extraction, we

Algorithm 1: Tongue Image Upright Orientation Process

Input: Color images X_1, X_2, \dots, X_n , Grayscale conversion function $G(\cdot)$, Contour smoothing functions $S_u(\cdot)$ and $S_l(\cdot)$, Tongue contour detection function $F(\cdot)$, Image rotation function $R(\cdot, \theta)$, Image translation function $T(\cdot, \Delta x, \Delta y)$, Margin m , smoothness s , and angle threshold α

Output: Uprighted tongue images Y_1, Y_2, \dots, Y_n

```

1: for  $i = 1$  to  $n$  do
2:    $G_i = G(X_i)$ 
3:    $U, L = F(G_i, \alpha)$  // Detect edges and filter
                        // points where angle between points  $> \alpha$ 
4:    $U_s = S_u(U, G_i, s)$ 
5:    $L_s = S_l(L, G_i, s)$ 
6:    $T_{\text{tip}} = L_s[\lfloor |L_s|/2 \rfloor]$  // Find tongue tip point
7:    $T_{\text{top}} = U_s[\lfloor |U_s|/2 \rfloor]$  // Find tongue top point
8:    $\theta = \arctan 2(T_{\text{tip}.y} - T_{\text{top}.y}, T_{\text{tip}.x} - T_{\text{top}.x})$ 
                        // Calculate rotation angle
9:    $X_r = R(X_i, \theta)$  // Rotate image
10:   $\Delta x = \text{target}x - T_{\text{tip}.x}$ ,  $\Delta y = \text{target}y -$ 
                         $T_{\text{tip}.y}$ 
11:   $X_t = T(X_r, \Delta x, \Delta y)$  // Translate image
12:   $Y_i = \text{CropImage}(X_t, m)$ 
13: end for
14: return  $Y_1, Y_2, \dots, Y_n$ 

```

customize an approach to adjust the tongue to an upright position. This process ensures that the subsequent analysis is performed on uniformly oriented images, which is crucial to the accuracy and reliability of our tongue diagnosis system.

Our Tongue Image Upright Orientation Algorithm, as detailed in Algorithm 1, consists of several key steps. First, we convert the input color image to grayscale using the conversion operation $G(\cdot)$ for simplified processing. We then apply a contour detection approach $F(\cdot)$ to the grayscale image to identify the tongue’s upper and lower contours, filtering out points where the angle of points exceeds a threshold α . Note that our method assumes typical situations where the two contours must be located in the upper and lower regions of the image, and does not address extreme cases like completely upside-down or horizontal tongue images. This method efficiently traces the upper and lower contours of the tongue. To refine these contours for high-quality separation, we apply smoothing techniques using functions $S_u(\cdot)$ and $S_l(\cdot)$ for the traced contours. From these smoothed contours, we identify the tongue tip and top, which are the middle points of the lower and upper contours, respectively. The rotation angle is then calculated based on the line connecting the tongue tip and top, and the image is rotated by $R(\cdot, \theta)$. Afterwards, the image is translated using $T(\cdot, \Delta x, \Delta y)$ to align the tongue tip with a target position. Finally, the image is cropped with a specified margin m . In our experiments, we found that angle thresholds between 50° and 70° all produce similarly satisfactory results, with 60° chosen for our final implementation. This process ensures that the tongue is consistently positioned in an upright orientation, facilitating more accurate and standardized feature extraction in subsequent steps, thus helping mitigate environmental variations

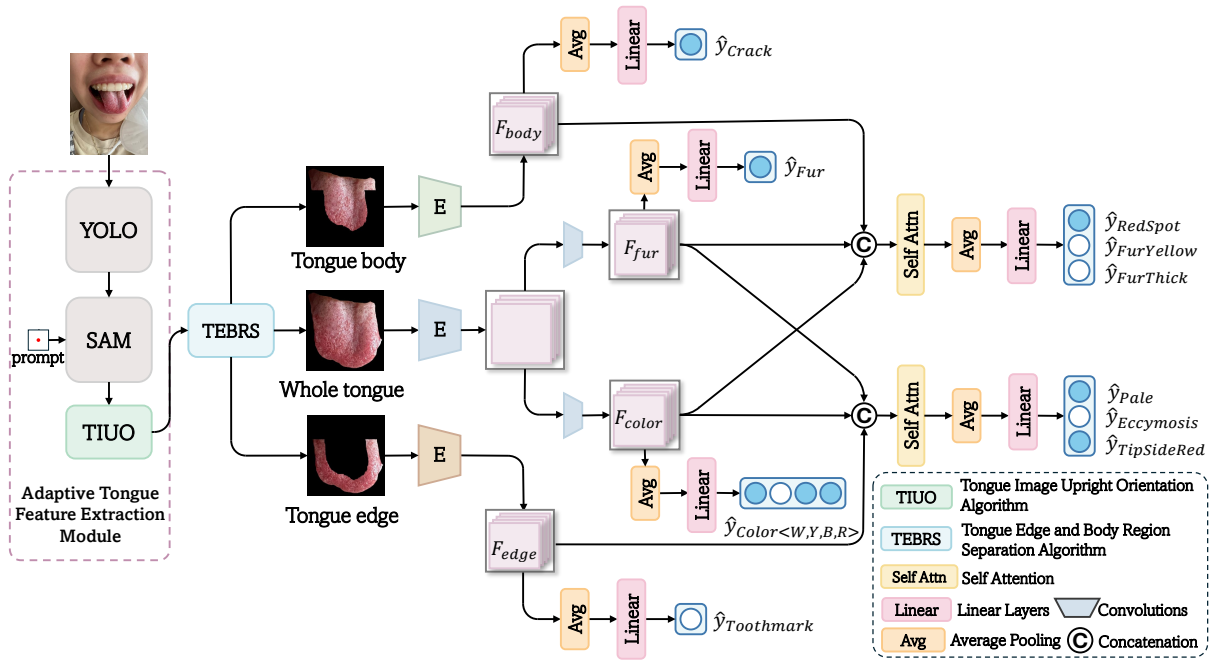


Figure 5: Overview of our Sign-Oriented Attributes Detection Framework. The framework consists of three main stages: (1) Adaptive Tongue Feature Extraction (ATFE) module, which employs a detection network to locate the tongue, followed by Mobile-SAM (Segment Anything Model) for precise tongue segmentation. The segmented tongue image then undergoes upright orientation for standardization. (2) Tongue Edge and Body Region Separation (TEBS), where the normalized tongue image is further separated into the tongue body and edge regions. (3) SignNet Pipeline, where the whole tongue, tongue body, and tongue edge images are fed into separate network branches to predict eight distinct tongue attributes.

Surface attributes	Color	Shape	Region	Fur
Pale	✓		△E	×
TipSideRed	✓		E	×
RedSpot	✓		△B	✓
Ecchymosis	✓		△E	×
Crack		✓	B	
Toothmark		✓	E	
FurYellow	✓		△B	✓
FurThick	✓		△B	✓

Table 2: Relationship between tongue attributes and their characteristic signs: color, shape, region, and fur.

such as lighting conditions and image qualities.

Attribute-Sign Relationships Graph

After standardizing tongue images, we tackle the identification of eight basic surface attributes. However, using a single ResNet50 or ViT yields poor results for some attributes, particularly pale and ecchymosis in Table 3. This may stem from the network’s lack of understanding of the specific areas to be identified. For instance, the pale attribute image in Figure 2 only refers to the tongue region without fur. Although the fur area may appear pale in color, it is not included in the scope of pale attribute identification. This con-

tradictory situation, where areas are mutually exclusive but similar in color, is likely to confuse the network.

To address this, we established a relationship table (Table 2) between the eight tongue attributes and four characteristic signs: color, shape, region, and fur. In the table, ‘✓’ indicates a definite connection, ‘×’ denotes mutual exclusivity, ‘△’ suggests a highly likely relationship and blank spaces indicate no connection.

Table 2 shows that color is associated with six attributes, excluding crack and toothmark. Shape relates only to crack and toothmark. Region-wise, some attributes are area-specific: toothmark and tipsidered occur at the edge (E), crack in the body (B), while pale and ecchymosis likely appear at the edge (△E). RedSpot, FurYellow, and FurThick likely appear in the body region (△B). Regarding fur, redspot, furyellow, and furthick describe fur characteristics, while pale, ecchymosis, and tipsidered areas are definitely fur-free.

To effectively analyze these region-specific attributes, we segment the tongue into body and edge regions for more in-depth analysis. These segmented regions and the original image serve as inputs to our SignNet architecture.

Sign-Oriented Attributes Detection Network

To realize the complex attribute relationships described in Table 2 and simulate the real-world diagnostic process of Practitioners, we propose a multi-branch network architec-

ture as shown in Figure 5. Our network consists of three main parts: the whole tongue analysis branch, the tongue body analysis branch, and the tongue edge analysis branch.

First, the whole tongue branch processes the entire tongue image, predicting fur color and the presence of fur. The color classification is achieved by combining multiple related attributes into four main categories: the White category is determined by the presence of Pale or Furthick attributes, the Yellow category corresponds to the FurYellow attribute, the Black category is associated with the Ecchymosis attribute, and the Red category is determined by the presence of TipSideRed or RedSpot attributes. This approach simplifies the learning process by transforming the difficult structural recognition problem into a simpler color recognition task while helping mitigate similar-appearing attributes under various imaging conditions. Also, fur is a fine-grained structure that is challenging to identify. By detecting the presence of fur, the network can initially capture the concept of fur and further learn based on the information in Table 2 and other complementary data. The tongue body branch focuses on the central area of the tongue, with one of its branches dedicated to identifying the crack attribute that only appears in the body region. Similarly, the tongue edge branch analyzes the edge area of the tongue, with one branch identifying tooth marks and predicting the toothmark attribute specific to the edge. Subsequently, we employ attention mechanisms to comprehensively fuse the color and fur information from the whole tongue network with the tongue edge and tongue body network information, respectively. For feature fusion, we apply:

$$F_{\text{cat},i} = [F_{\text{fur}}; F_{\text{color}}; F_i], \quad i \in \{\text{edge}, \text{body}\} \quad (1)$$

$$Q_i = W_Q F_{\text{cat},i}, \quad K_i = W_K F_{\text{cat},i}, \quad V_i = W_V F_{\text{cat},i} \quad (2)$$

$$A_i = \text{softmax} \left(\frac{Q_i K_i^T}{\sqrt{d_k}} \right) V_i, \quad (3)$$

where F_{fur} , F_{color} , F_{edge} , and F_{body} represent features from fur, color, tongue edge, and tongue body, respectively. $F_{\text{cat},i}$ is the concatenated feature vector. W_Q , W_K , and W_V are learnable weight matrices. The subscript i denotes either the edge or body branch. Following the attention mechanism, we apply a Feed-Forward Network (FFN):

$$\text{FFN}(x) = \sigma(xW_1 + b_1)W_2 + b_2, \quad (4)$$

where W_1 , W_2 , b_1 , and b_2 are learnable parameters, and $\sigma(\cdot)$ represents the activation function. The complete feature fusion process can be expressed as:

$$\begin{aligned} F_i' &= \text{LayerNorm}(F_i + A_i) \\ F_i'' &= \text{LayerNorm}(F_i' + \text{FFN}(F_i')), \end{aligned} \quad (5)$$

where i denotes either the edge or body branch. This process applies self-attention and a feed-forward network to each branch, with layer normalization after each step to enhance and stabilize the feature representations.

This network architecture, based on the attribute distribution in Table 2, simulates the TCM practitioners' diagnostic process. As our Attribute-Sign Relationships Graph

illustrates, each tongue attribute has its corresponding regional, color, and fur characteristics. Our multi-branch structure reflects these relationships: edge-related attributes (pale, ecchymosis, tipsidered) and body-related attributes (redspot, furyellow, furthick) are processed in their respective branches. By fusing whole tongue color and fur information with edge and body features through attention mechanisms, our network effectively predicts basic tongue attributes for each region. This approach mimics the practitioners' diagnostic method, analyzing tongue attributes through multiple signs.

Loss Function

We apply several loss functions to govern the training of SignNet in terms of three main tasks: classification of color, fur, and eight binary attributes. Each task is weighted to balance their contributions to the overall loss:

$$\mathcal{L}_{\text{total}} = w_{\text{color}} \mathcal{L}_{\text{color}} + w_{\text{fur}} \mathcal{L}_{\text{fur}} + \mathcal{L}_{\text{attr}}. \quad (6)$$

Here, $\mathcal{L}_{\text{color}}$ uses sigmoid cross-entropy losses (Xu et al. 2023) for the multi-label color classification, \mathcal{L}_{fur} employs sigmoid cross-entropy loss for fur classification, and $\mathcal{L}_{\text{attr}}$ is a weighted sigmoid cross-entropy loss for the eight binary features:

$$\mathcal{L}_{\text{attr}} = \sum_{j=1}^8 \alpha_j \mathcal{L}_{\text{BCE}}(y_j, \hat{y}_j), \quad (7)$$

where α_j represents frequency-based weights for each binary feature, mitigating potential class imbalances. These weights are calculated as:

$$\alpha_j = \frac{\text{median}}{F_j}. \quad (8)$$

Here, m is the median frequency of all binary features in the training set, and F_j is the frequency of binary feature j .

Experiments

Compared Methods. To evaluate our method's effectiveness, we compared several models for tongue attribute recognition. Our baseline models include ResNet50 and Vision Transformer (ViT). We also reproduced the highest-accuracy method in Yuan et al. (Yuan et al. 2023), which is based on TransFG (He et al. 2022). Other methods, such as Jiang et al. (Jiang et al. 2022), use bounding box annotations for subtle tongue features, making direct comparison impossible, and their code and datasets are also not publicly available. We further tested ResNet50+ATFE, where ATFE is our proposed tongue image standardization process (including detection, segmentation, and upright orientation), and our proposed SignNet.

Additionally, we conducted an experiment involving human evaluation. We invited three registered practitioners and five laypersons to classify eight simple attributes based on images from our test set, with their accuracy rates averaged. The laypersons received a brief training session on attribute recognition as illustrated in Figure 2. It is important to note that while our annotations were made through direct observation in a clinical setting, the human participants made their classifications solely from the images.

Model	Pale		TipSideRed		Spot		Ecchymosis		Crack	
	Acc	F1	Acc	F1	Acc	F1	Acc	F1	Acc	F1
Resnet50	88.31±0.56	32.68±2.04	72.94±1.17	66.19±0.81	78.06±1.31	76.44±1.26	90.75±0.58	24.72±2.83	84.92±0.86	91.17±0.55
Vit	89.45±0.13	1.61±2.20	66.64±0.56	43.41±1.57	60.92±1.34	50.90±5.15	90.84±0.00	0.00±0.00	83.46±0.00	90.99±0.00
Resnet50+ATFE	88.16±2.15	39.72±3.81	70.68±3.06	67.26±0.83	81.72±2.28	80.70±1.87	91.17±1.03	29.78±3.97	84.56±2.48	90.86±1.72
TransFG	89.61±0.00	0.00±0.00	64.02±1.11	31.08±7.12	62.68±3.30	43.63±8.56	90.84±0.00	0.00±0.00	83.46±0.00	90.99±0.00
Laypersons	48.94	18.54	49.39	45.62	51.06	49.42	49.72	14.77	47.71	59.59
Practitioners	86.82	43.81	76.20	71.64	82.91	81.81	85.14	38.14	87.71	92.46
Ours	90.21±1.10	46.00±5.80	74.08±1.72	68.22±1.06	82.35±0.70	81.09±0.84	92.13±0.48	38.02±3.05	86.23±0.88	91.90±0.56

Model	ToothMark		FurThick		FurYellow		Average		Jaccard
	Acc	F1	Acc	F1	Acc	F1	Acc	F1	
Resnet50	80.72±0.64	84.31±0.55	97.34±0.43	98.65±0.17	92.25±0.98	76.27±3.09	85.66±0.23	68.80±0.48	74.21±0.43
Vit	60.87±1.10	73.62±0.75	97.21±0.00	98.58±0.00	82.57±0.00	0.00±0.00	78.99±0.17	44.89±1.04	63.68±0.41
Resnet50+ATFE	80.12±2.50	83.63±1.73	97.34±0.09	98.65±0.05	92.11±0.44	70.92±0.80	85.73±1.12	70.92±0.80	74.86±0.99
TransFG	60.89±0.70	75.52±0.24	97.21±0.00	98.58±0.00	82.57±0.00	0.00±0.00	78.91±0.61	42.48±1.89	63.40±1.11
Laypersons	50.95	54.51	49.61	65.44	48.94	23.45	49.54	41.42	31.13
Practitioners	82.68	86.39	91.17	95.29	89.16	73.85	85.22	72.92	77.63
Ours	79.24±0.98	82.82±0.61	97.14±0.19	98.54±0.10	92.07±0.35	76.41±1.00	86.68±0.30	72.87±0.63	75.83±0.37

Note: Acc columns show accuracy scores, and F1 columns show F1-scores. The Jaccard index provides a measure of overall performance across all attributes.

Table 3: Quantitative comparison among different methods using five-fold cross-validation (values in mean% ± std%).

Model	Pale	TipSideRed	Spot	Ecchymosis	Crack	ToothMark	FurThick	FurYellow	Average	Jaccard
Resnet50	32.68±2.04	66.19±0.81	76.44±1.26	24.72±2.83	91.17±0.55	84.31±0.55	98.65±0.17	76.27±3.09	68.80±0.48	74.21±0.43
Resnet50+det	38.43±2.63	65.75±0.91	79.70±0.98	27.27±2.06	91.35±0.52	83.50±0.68	98.51±0.14	75.02±1.28	69.94±0.60	74.65±0.46
Resnet50+seg	34.62±4.23	67.13±1.87	80.21±0.90	29.48±3.45	91.79±0.33	81.84±0.97	98.52±0.16	74.77±1.67	69.80±0.51	74.61±0.49
Resnet50+upright	39.72±3.81	67.26±0.83	80.70±1.87	29.78±3.97	90.86±1.72	83.63±1.73	98.65±0.05	76.77±1.55	70.92±0.80	74.86±0.99
SignNet w/o color	40.65±2.04	66.10±0.94	81.05±0.37	34.94±4.47	91.10±0.38	82.23±0.86	98.55±0.22	75.99±2.53	71.33±0.38	74.62±0.60
SignNet w/o fur	42.16±4.64	67.03±1.20	79.86±0.56	34.94±3.23	91.77±0.51	81.55±0.52	98.46±0.17	77.10±2.02	71.61±0.64	74.74±0.80
SignNet	46.00±5.80	68.22±1.06	81.09±0.84	38.02±3.05	91.90±0.56	82.82±0.61	98.54±0.10	76.41±1.00	72.87±0.63	75.83±0.37

Note: All columns except the last one show F1-scores. The Jaccard index in the last column provides a measure of overall performance across all attributes.

Table 4: Ablations of different model variants using five-fold cross-validation: F1-score and Jaccard (mean% ± std%).

Quantitative Results. As shown in Tables 3, the ViT-based models significantly underperform compared to ResNet50, with F1-scores for attributes like Pale and Ecchymosis approaching zero, indicating a fundamental lack of understanding of the task. This degradation can be attributed to ViT’s limited capacity to capture high-frequency features. Similarly, the transformer-based TransFG architecture also exhibits comparable performance deficiencies.

Comparing ResNet50 with ResNet50+ATFE, we observe a significant improvement in F1-scores, with an average increase of about 2%. Given the highly imbalanced nature of our dataset as shown in Table 1, we prioritize F1-score over accuracy. Our SignNet design further improves performance, achieving 1% higher average accuracy and about 4% higher average F1-score compared to the baseline ResNet50.

Additionally, all methods show particularly low F1-scores for pale and ecchymosis attributes. This could be attributed to the influence of overly bright or dark lighting conditions in the images. Notably, our network not only surpasses laypersons’ accuracy but also approaches the performance of practitioners. This demonstrates the effectiveness of our proposed method in handling tongue attributes recognition.

Ablation Study

In this section, we present an ablation study on the efficacy of our main components in Table 4. ResNet50 serves as our baseline network. The Resnet50+det, Resnet50+seg, and

Resnet50+upright models represent steps of our proposed ATFE module, corresponding to tongue detection, segmentation, and upright orientation, respectively. These components aim to standardize tongue image inputs. From Table 4, we can observe that images processed through the ATFE module show improvements in both average and F1-score metrics, demonstrating its effectiveness in mitigating varying imaging conditions.

Furthermore, we examine the impact of removing the color branch and fur branch from our proposed SignNet. The results labeled as ‘SignNet w/o color’ and ‘SignNet w/o fur’ illustrate the performance when these branches are individually removed. Notably, we can see that only when both branches are simultaneously included do we achieve the highest accuracy, showing their complementary roles in mitigating imaging variations. Therefore, color and fur branches play crucial roles in capturing tongue attributes and contributing to performance.

Conclusion

In this study, we developed a comprehensive tongue image dataset and proposed a novel framework for remote diagnosis, combining adaptive feature extraction and a sign-oriented recognition network. Our approach achieves near-practitioner performance despite dataset limitations. Future work aims to utilize all available labels, including demographics, to further enhance accuracy.

Ethical statement

All procedures performed in this study with human participants were in accordance with ethical standards and approved by the Institutional Review Board of Hong Kong Polytechnic University. Informed consent was obtained from all individual participants. All tongue images were specially processed to protect personal privacy and prevent disclosure of identifying information. All data collected during the study was anonymized and securely stored, with access restricted to the research team.

Acknowledgements

This work is partially supported by a Shenzhen-Hong Kong-Macao Science and Technology Plan Project (Category C Project) under the Shenzhen Municipal Science and Technology Innovation Commission (project no. SGDX20230821092359002), a grant for Collaborative Research with World-leading Research Groups of The Hong Kong Polytechnic University (project no. G-SACF), the Guangdong Natural Science Funds for Distinguished Young Scholars (Grant 2023B1515020097), the AI Singapore Programme under the National Research Foundation Singapore (Grant AISG3-GV-2023-011), and the Lee Kong Chian Fellowships.

References

Bansal, A.; Sikka, K.; Sharma, G.; Chellappa, R.; and Divakaran, A. 2018. Zero-shot object detection. In *Proceedings of the European conference on computer vision (ECCV)*, 384–400.

Bucher, M.; Vu, T.-H.; Cord, M.; and Pérez, P. 2019. Zero-shot semantic segmentation. *Advances in Neural Information Processing Systems*, 32.

Fu, H.; Cohen-Or, D.; Dror, G.; and Sheffer, A. 2008. Upright orientation of man-made objects. In *ACM SIGGRAPH 2008 papers*, 1–7.

He, J.; Chen, J.-N.; Liu, S.; Kortylewski, A.; Yang, C.; Bai, Y.; and Wang, C. 2022. Transfg: A transformer architecture for fine-grained recognition. In *Proceedings of the AAAI conference on artificial intelligence*, volume 36, 852–860.

Heo, J.; Lim, J. H.; Lee, H. R.; Jang, J. Y.; Shin, Y. S.; Kim, D.; Lim, J. Y.; Park, Y. M.; Koh, Y. W.; Ahn, S.-H.; et al. 2022. Deep learning model for tongue cancer diagnosis using endoscopic images. *Scientific reports*, 12(1): 6281.

Howard, A.; Sandler, M.; Chu, G.; Chen, L.-C.; Chen, B.; Tan, M.; Wang, W.; Zhu, Y.; Pang, R.; Vasudevan, V.; et al. 2019. Searching for mobilenetv3. In *Proceedings of the IEEE/CVF international conference on computer vision*, 1314–1324.

Hu, J.; Han, S.; Chen, Y.; Ji, Z.; et al. 2015. Variations of tongue coating microbiota in patients with gastric cancer. *BioMed research international*, 2015.

Hu, Y.; Wen, G.; Liao, H.; Wang, C.; Dai, D.; and Yu, Z. 2019. Automatic construction of chinese herbal prescriptions from tongue images using CNNs and auxiliary latent therapy topics. *IEEE transactions on cybernetics*, 51(2): 708–721.

Huang, B.; Wu, J.; Zhang, D.; and Li, N. 2010. Tongue shape classification by geometric features. *Information Sciences*, 180(2): 312–324.

Jia, C.; Yang, Y.; Xia, Y.; Chen, Y.-T.; Parekh, Z.; Pham, H.; Le, Q.; Sung, Y.-H.; Li, Z.; and Duerig, T. 2021. Scaling up visual and vision-language representation learning with noisy text supervision. In *International conference on machine learning*, 4904–4916. PMLR.

Jiang, B.; Liang, X.; Chen, Y.; Ma, T.; Liu, L.; Li, J.; Jiang, R.; Chen, T.; Zhang, X.; and Li, S. 2012. Integrating next-generation sequencing and traditional tongue diagnosis to determine tongue coating microbiome. *Scientific reports*, 2(1): 936.

Jiang, T.; Guo, X.-j.; Tu, L.-p.; Lu, Z.; Cui, J.; Ma, X.-x.; Hu, X.-j.; Yao, X.-h.; Cui, L.-t.; Li, Y.-z.; et al. 2021a. Application of computer tongue image analysis technology in the diagnosis of NAFLD. *Computers in Biology and Medicine*, 135: 104622.

Jiang, T.; Hu, X.-j.; Yao, X.-h.; Tu, L.-p.; Huang, J.-b.; Ma, X.-x.; Cui, J.; Wu, Q.-f.; and Xu, J.-t. 2021b. Tongue image quality assessment based on a deep convolutional neural network. *BMC Medical Informatics and Decision Making*, 21(1): 1–14.

Jiang, T.; Lu, Z.; Hu, X.; Zeng, L.; Ma, X.; Huang, J.; Cui, J.; Tu, L.; Zhou, C.; Yao, X.; et al. 2022. Deep Learning Multi-label Tongue Image Analysis and Its Application in a Population Undergoing Routine Medical Checkup. *Evidence-Based Complementary and Alternative Medicine*, 2022.

Jiang, X.; Zheng, C.; Xu, X.; Liu, B.; Zheng, W.; Zhang, H.; and He, S. 2024. VrdONE: One-stage Video Visual Relation Detection. In *Proceedings of the 32nd ACM International Conference on Multimedia*, 1437–1446.

Jocher, G.; Chaurasia, A.; Stoken, A.; Borovec, J.; Kwon, Y.; Michael, K.; Fang, J.; Wong, C.; Yifu, Z.; Montes, D.; et al. 2022. ultralytics/yolov5: v6. 2-yolov5 classification models, apple m1, reproducibility, clearml and deci. ai integrations. *Zenodo*.

Ke, L.; Ye, M.; Danelljan, M.; Tai, Y.-W.; Tang, C.-K.; Yu, F.; et al. 2024. Segment anything in high quality. *Advances in Neural Information Processing Systems*, 36.

Kichloo, A.; Albosta, M.; Dettloff, K.; Wani, F.; El-Amir, Z.; Singh, J.; Aljadah, M.; Chakinala, R. C.; Kanugula, A. K.; Solanki, S.; et al. 2020. Telemedicine, the current COVID-19 pandemic and the future: a narrative review and perspectives moving forward in the USA. *Family medicine and community health*, 8(3).

Kirillov, A.; Mintun, E.; Ravi, N.; Mao, H.; Rolland, C.; Gustafson, L.; Xiao, T.; Whitehead, S.; Berg, A. C.; Lo, W.-Y.; et al. 2023. Segment anything. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 4015–4026.

Lee, T.-C.; Lo, L.-C.; Wu, F.-C.; et al. 2016. Traditional Chinese medicine for metabolic syndrome via TCM pattern differentiation: Tongue diagnosis for predictor. *Evidence-Based Complementary and Alternative Medicine*, 2016.

- Li, J.; Li, D.; Xiong, C.; and Hoi, S. 2022. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *International conference on machine learning*, 12888–12900. PMLR.
- Li, X.; Zhang, Y.; Cui, Q.; Yi, X.; and Zhang, Y. 2018. Tooth-marked tongue recognition using multiple instance learning and CNN features. *IEEE transactions on cybernetics*, 49(2): 380–387.
- Liu, S.; Zeng, Z.; Ren, T.; Li, F.; Zhang, H.; Yang, J.; Li, C.; Yang, J.; Su, H.; Zhu, J.; et al. 2023. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. *arXiv preprint arXiv:2303.05499*.
- Lo, L.-c.; Chen, Y.-F.; Chen, W.-J.; Cheng, T.-L.; Chiang, J. Y.; et al. 2012. The study on the agreement between automatic tongue diagnosis system and traditional Chinese medicine practitioners. *Evidence-Based Complementary and Alternative Medicine*, 2012.
- Ohannessian, R.; Duong, T. A.; Odone, A.; et al. 2020. Global telemedicine implementation and integration within health systems to fight the COVID-19 pandemic: a call to action. *JMIR public health and surveillance*, 6(2): e18810.
- Pang, B.; Zhang, D.; Li, N.; and Wang, K. 2004. Computerized tongue diagnosis based on Bayesian networks. *IEEE Transactions on biomedical engineering*, 51(10): 1803–1810.
- Portnoy, J.; Waller, M.; and Elliott, T. 2020. Telemedicine in the era of COVID-19. *The Journal of Allergy and Clinical Immunology: In Practice*, 8(5): 1489–1491.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PMLR.
- Shi, Y.; Hu, X.; Cui, J.; Cui, L.; Huang, J.; Ma, X.; Jiang, T.; Yao, X.; Lan, F.; Li, J.; et al. 2021. Clinical data mining on network of symptom and index and correlation of tongue-pulse data in fatigue population. *BMC Medical Informatics and Decision Making*, 21(1): 1–14.
- Wang, H.; Zhang, X.; and Cai, Y. 2014. Research on teeth marks recognition in tongue image. In *2014 International Conference on Medical Biometrics*, 80–84. IEEE.
- Wang, J.; Lan, C.; Liu, C.; Ouyang, Y.; Qin, T.; Lu, W.; Chen, Y.; Zeng, W.; and Philip, S. Y. 2022. Generalizing to unseen domains: A survey on domain generalization. *IEEE transactions on knowledge and data engineering*, 35(8): 8052–8072.
- Wang, X.; Zhang, B.; Yang, Z.; Wang, H.; and Zhang, D. 2013. Statistical analysis of tongue images for feature extraction and diagnostics. *IEEE Transactions on Image Processing*, 22(12): 5336–5347.
- Wang, Z.; Yu, J.; Yu, A. W.; Dai, Z.; Tsvetkov, Y.; and Cao, Y. 2021. Simvlm: Simple visual language model pretraining with weak supervision. *arXiv preprint arXiv:2108.10904*.
- Xian, Y.; Lampert, C. H.; Schiele, B.; and Akata, Z. 2018. Zero-shot learning—a comprehensive evaluation of the good, the bad and the ugly. *IEEE transactions on pattern analysis and machine intelligence*, 41(9): 2251–2265.
- Xu, C.; Chen, Z.; Mai, J.; Xu, X.; and He, S. 2023. Pose-and attribute-consistent person image synthesis. *ACM Transactions on Multimedia Computing, Communications and Applications*, 19(2s): 1–21.
- Xu, C.; Qu, W.; Xu, X.; and Liu, X. 2022. Multi-scale flow-based occluding effect and content separation for cartoon animations. *IEEE Transactions on Visualization and Computer Graphics*, 29(9): 4001–4014.
- Yu, Y.; Liu, B.; Zheng, C.; Xu, X.; Zhang, H.; and He, S. 2024. Beyond textual constraints: Learning novel diffusion conditions with fewer examples. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7109–7118.
- Yuan, L.; Yang, L.; Zhang, S.; Xu, Z.; Qin, J.; Shi, Y.; Yu, P.; Wang, Y.; Bao, Z.; Xia, Y.; et al. 2023. Development of a tongue image-based machine learning tool for the diagnosis of gastric cancer: a prospective multicentre clinical cohort study. *EClinicalMedicine*, 57.
- Zhang, C.; Han, D.; Qiao, Y.; Kim, J. U.; Bae, S.-H.; Lee, S.; and Hong, C. S. 2023. Faster segment anything: Towards lightweight sam for mobile applications. *arXiv preprint arXiv:2306.14289*.
- Zhao, Y.; Gou, X.-j.; Dai, J.-y.; Peng, J.-h.; Feng, Q.; Sun, S.-j.; Cao, H.-j.; Zheng, N.-n.; Fang, J.-w.; Jiang, J.; et al. 2013. Differences in metabolites of different tongue coatings in patients with chronic hepatitis B. *Evidence-Based Complementary and Alternative Medicine*, 2013.
- Zheng, Y.; Wu, J.; Qin, Y.; Zhang, F.; and Cui, L. 2021. Zero-shot instance segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2593–2602.
- Zhi, L.; Zhang, D.; Yan, J.-q.; Li, Q.-L.; and Tang, Q.-l. 2007. Classification of hyperspectral medical tongue images for tongue diagnosis. *Computerized Medical Imaging and Graphics*, 31(8): 672–678.
- Zhu, P.; Wang, H.; and Saligrama, V. 2019. Zero shot detection. *IEEE Transactions on Circuits and Systems for Video Technology*, 30(4): 998–1010.
- Zuo, W.; Wang, K.; Zhang, D.; and Zhang, H. 2004. Combination of polar edge detection and active contour model for automated tongue segmentation. In *Third International Conference on Image and Graphics (ICIG'04)*, 270–273. IEEE.