

Image Quality-aware Diagnosis via Meta-knowledge Co-embedding

Haoxuan Che Siyu Chen Hao Chen*
 The Hong Kong University of Science and Technology

{hche, schende, jhc}@cse.ust.hk

Abstract

Medical images usually suffer from image degradation in clinical practice, leading to decreased performance of deep learning-based models. To resolve this problem, most previous works have focused on filtering out degradation-causing low-quality images while ignoring their potential value for models. Through effectively learning and leveraging the knowledge of degradations, models can better resist their adverse effects and avoid misdiagnosis. In this paper, we raise the problem of **image quality-aware diagnosis**, which aims to take advantage of low-quality images and image quality labels to achieve a more accurate and robust diagnosis. However, the diversity of degradations and superficially unrelated targets between image quality assessment and disease diagnosis makes it still quite challenging to effectively leverage quality labels to assist diagnosis. Thus, to tackle these issues, we propose a novel **meta-knowledge co-embedding network**, consisting of two subnets: *Task Net* and *Meta Learner*. *Task Net* constructs an explicit quality information utilization mechanism to enhance diagnosis via knowledge co-embedding features, while *Meta Learner* ensures the effectiveness and constrains the semantics of these features via meta-learning and joint-encoding masking. Superior performance on five datasets with four widely-used medical imaging modalities demonstrates the effectiveness and generalizability of our method.

1. Introduction

Medical imaging is one of the most valuable sources of diagnostic information about anatomical structures and pathological characteristics [1]. Advanced deep learning-based methods applied to high-quality (HQ) medical images have shown significant potential in disease analysis and diagnosis [2, 3], achieving favorable results compared with human healthcare professionals [4]. However, in clinical practice, obtaining HQ images is not always feasible. Medical images often exhibit significant variations in imaging quality due to factors such as patient movements or en-




Figure 1. Illustration of impact of image degradations on diagnosis semantics for fundus (top) and OCTA (bottom) images. **Top:** Degradations obscure parts of the vessel structure and present lesion-like spots. **Bottom:** Degradations result in a fake enlargement of the foveal avascular zone (central circular area).

vironmental conditions [5, 6]. For instance, a medical image quality assessment study of 28,780 fundus images revealed that approximately 41.6% of them contained image artifacts and corruption and were considered low-quality (LQ) [7]. Such degradations can increase the uncertainty in pathological observation, leading to misdiagnosis [8, 9].

Medical image degradations can significantly affect diagnostic semantics, as illustrated in Figure 1. For instance, shadow degradation can obscure anatomical structures crucial for diagnosis, while spot artifacts can obfuscate pathological signs that typically manifest as circular shapes [10]. Furthermore, image degradations can also affect diagnostic measurements, such as the vessel area density in optical coherence tomography angiography (OCTA) images, rendering them unreliable [6]. These close relationships raise challenges in distinguishing degradations from actual abnormalities [10], leading to false knowledge of lesions and undesired misdiagnosis during training and deployment [9]. Aware of the profound influence of image quality on diagnosis, many previous works have focused on utilizing image quality assessment to select relatively HQ images for training and testing, thereby avoiding the influence of LQ im-

*Corresponding author: Hao Chen, email: jhc@cse.ust.hk.

ages [7, 11–14]. However, discarding LQ images containing diagnostically valuable information results in a waste of precious clinical data [7]. Including LQ images and corresponding quality information in training can assist models in recognizing potential false abnormalities, thus achieving more robust and accurate diagnosis [15–17].

In this paper, we reconsider the value of LQ images and corresponding image quality labels, and introduce the problem of *image quality-aware diagnosis* (IQAD). The goal of IQAD is to enable models to leverage LQ images while simultaneously learning image quality labels to achieve an accurate and robust diagnosis. However, effectively leveraging quality labels for diagnosis is non-trivial with a multi-task learning framework. Specifically, image quality assessment can be considered as a task “unrelated” to disease diagnosis [18], since it focuses on capturing image degradations, while diagnosis emphasizes identifying lesions. This distinction makes it challenging for models to effectively utilize image quality labels. Further, commonly-used coarse annotations of quality may not sufficiently reflect the diversity of image degradation, making it difficult to provide information that could be useful for precise diagnosis.

To achieve IQAD, we propose a novel *meta-knowledge co-embedding network* (MKCNet) consisting of two subnets, Task Net and Meta Learner. To enable leveraging potential benefits of quality information, Task Net conducts diagnosis predictions by explicitly leveraging knowledge co-embedding features with desired knowledge of image quality and disease diagnosis. These features are constructed by learning auxiliary label embeddings from Meta Learner. Further, we employ meta-learning and joint-encoding masking to ensure the effectiveness and semantics of auxiliary label embedding and circumvent the barrier of obtaining fine-grained image quality labels. Specifically, joint-encoding masking selects a part of the Meta Learner output as auxiliary label embedding through combinations of quality and diagnosis labels. Moreover, Meta Learner learns to provide auxiliary label embedding in a meta-learning manner to assist Task Net optimization, encouraging it to learn effective knowledge co-embedding features. Our main contributions are highlighted as follows:

(1) We tackle a novel problem named IQAD. To the best of our knowledge, this is the first work to discuss and analyze this critical and practical problem.

(2) We propose a novel method, MKCNet, to effectively handle the challenges posed by annotation granularity and task focus discrepancy via leveraging quality information explicitly and introducing a meta-learning paradigm.

(3) We conduct extensive experiments on five datasets spanning four different yet widely-used medical imaging modalities. Our in-depth analytical study demonstrates the effectiveness and generalizability of MKCNet.

2. Related Work

Disease diagnosis. Many deep-learning methods have been developed to diagnose diseases [19–25]. For example, He et al. [19] proposed CABNet for learning discriminative features associated with different severities of diabetic retinopathy (DR), and Liu et al. [20] developed a convolutional graph networks-based method to explore potential relationships among grades of DR. However, these methods may not perform well when dealing with image degradations as they do not consider image quality issues [9]. By contrast, MKCNet leverages LQ images and quality labels to achieve a more robust and accurate diagnosis.

Quality assessment and image enhancement. Aiming at avoiding the effect of LQ images, several methods have been proposed to assess image quality to select relatively HQ images [7, 11, 14]. For example, Fu et al. [7] selected usable samples and rejected nearly 20% of images as unsuitable for model learning. However, these images would still be diagnosable for physicians. Rejecting diagnostically valuable LQ images is wasteful, since LQ images are not only useful in model generalization ability improvement in training [15, 26], but also effective in evaluating the robustness of models in testing [27, 28]. Another solution is to leverage generative models to enhance the quality of LQ images [29–32]. For example, Shen et al. [10] improve the quality of LQ fundus images by generating degradations on images to simulate pseudo-paired samples. However, it is costly to train such generative models, which requires a large number of images for desirable performance [32]. Additionally, the simulation may only represent a partial distribution of realistic degradation [33], and it is rarely helpful in improving recognition performance [34]. Meanwhile, capturing real pairs with different image qualities proves to be extremely challenging [35]. Instead of requiring a high cost, MKCNet leverages the multi-task learning framework to tackle the IQAD problem in a less costly and more effective manner.

Multi-task learning. It has been shown that multi-task learning is an effective method for training a generalized model that can simultaneously handle multiple tasks [18]. In the medical domain, many studies have used it to improve the performance of models by exploring internal relationships among diseases [36–40], as well as utilizing auxiliary tasks to assist primary tasks [41, 42]. However, most of these works ignore the potential benefits of quality information on diagnosis, which is critical for the IQAD problem. Among them, Zhou et al. [16] treated quality assessment as an auxiliary task. However, they did not obtain a significant performance boost due to the absence of an explicit mechanism for utilizing quality information. By contrast, MKCNet explicitly models and explores the potential assistance of quality information in diagnosis.




Figure 2. The overview of MKCNet with two subnets (Task Net \mathcal{M}_θ and Meta Learner \mathcal{M}_ϕ). In the first stage, \mathcal{M}_θ learns to construct f_θ^ω by leveraging y_ω . Simultaneously, it adopts global attention to learn an informative and generalizable F_θ , while it explicitly utilizes f_θ^ω and f_θ^d to make diagnoses. In the second stage, \mathcal{M}_ϕ learns to provide y_ω with desired knowledge of image quality and disease diagnosis. \mathcal{M}_ϕ ensures the effectiveness of f_θ^ω while constraining its semantics by utilizing the joint-encoding masking and meta-optimization.

3. Methodology

This section first provides an introduction to IQAD and presents a preliminary experiment, as well as highlights the challenges involved. We then introduce our proposed solution, MKCNet, which explicitly addresses these challenges. Figure 2 shows an overview of MKCNet.

3.1. Image Quality-aware Diagnosis

Preliminaries. Given an image $x \in \mathcal{X}$, $y_d \in \mathcal{Y}_D$ and $y_q \in \mathcal{Y}_Q$ denote the corresponding disease diagnosis and image quality labels, respectively. The target of IQAD is to train a model $\mathcal{F} : \mathcal{X} \rightarrow \mathcal{Y}_D$ to achieve robust and accurate diagnosis by leveraging both HQ and LQ images together with corresponding diagnosis and quality labels.

Intuitively, treating image quality assessment (IQA) as an auxiliary branch in a multi-task learning framework may seem like a straightforward solution. Thus, to explore this approach, we conduct a preliminary experiment on a VGG16 model (denoted as Vanilla) to assess the impact of learning LQ images and leveraging quality labels, as shown in Figure 3. As expected, leveraging LQ images is beneficial for model learning. Unexpectedly, however, we observed that incorporating quality labels only marginally enhances or even hinders the diagnostic performance, despite the fact that these labels may imply false abnormalities and provide extra useful diagnostic information. It suggests that effectively leveraging quality labels is non-trivial.

Challenges. In the context of multi-task learning, we argue that the challenges of IQAD are two-fold. **The first challenge** arises from the non-direct relationship between

image quality and diagnosis, which makes it difficult for models to associate them without an explicit mechanism. Generally, models trained for disease diagnosis focus on lesion areas and anatomical structures, while image quality assessment requires models to capture image degradations. This discrepancy in the task focus necessitates a specialized design that enables models to effectively leverage quality information. To address this challenge, we design Task Net, which includes an explicit utilization mechanism to leverage knowledge co-embedding features for diagnosis. **The second challenge** arises due to the limited granularity of binary or multi-class quality annotations, which fail to capture the diverse types and degrees of image degradations [9, 10]. It results in a lack of detailed information that can guide the model and makes it challenging to establish fixed patterns of how image quality affects disease diagnosis [43]. Additionally, determining image-level label criteria to associate quality and diagnosis is complex, and annotating images with pixel-level information of degradations is costly [13]. Due to label limitations, even if an explicit utilization mechanism exists, the model may not effectively leverage quality information. To tackle this issue, we adopt a meta-learning paradigm, where the Meta Learner is trained to provide auxiliary label embeddings that present adaptive correlations between image quality and disease diagnosis labels, allowing the model to effectively utilize quality information.

3.2. Meta-knowledge Co-embedding Network

As shown in Figure 2, MKCNet is composed of two subnets, Task Net denoted as \mathcal{M}_θ with parameter θ , and Meta Learner denoted as \mathcal{M}_ϕ with parameter ϕ . Given an image




Figure 3. Influence of LQ images and an IQA module. It is useful for Vanilla to leverage LQ images, but the IQA module improves only marginally or harms performance. MKCNet can leverage LQ images and quality labels more effectively than Vanilla.

x and its corresponding image quality and disease diagnosis labels y_q and y_d , we obtain an auxiliary label embedding y_ω from the output vector $\mathcal{M}_\phi(x)$ of Meta Learner. The corresponding predictions of Task Net for y_q , y_d , and y_ω are denoted as $\mathcal{M}_\theta^q(x)$, $\mathcal{M}_\theta^d(x)$, and $\mathcal{M}_\theta^\omega(x)$, respectively.

Inspired by the recent success of meta-learning [44, 45], we adopt a two-stage learning paradigm to optimize MKCNet. In the first stage, Task Net is trained with y_q , y_d , and y_ω , and it explicitly leverages knowledge co-embedding features to improve diagnosis performance. In the second stage, Meta Learner learns to provide y_ω via meta-learning, to optimize knowledge co-embedding features in Task Net. The two stages are iterative during each training epoch, leading to end-to-end interaction between the two subnets.

3.2.1 Task Net for Quality-aware Diagnosing

To address the first challenge discussed earlier, Task Net is designed to construct knowledge co-embedding features while explicitly exploring their potential usefulness in improving diagnosis. To achieve this, we design the global attention block (GAB) in Task Net to extract task-specific features and the meta-knowledge assistance block (MAB) to explicitly explore their potential benefits in diagnosis. In Task Net, we denote the feature map from the backbone as F_θ , the features from three GABs as f_θ^q , f_θ^d , and f_θ^ω , and the feature from the MAB for final diagnosis as f_θ^{d*} .

Knowledge co-embedding feature. Task Net constructs the knowledge co-embedding feature f_θ^ω by learning y_ω , which contains the desired information beneficial for both image quality and disease diagnosis. To achieve this, it first needs to learn informative and generalizable F_θ , encompassing comprehensive semantics, and then to construct the f_θ^ω with applicable information for quality-aware diagnosis from F_θ . Therefore, We designed three learning branches with GAB in Task Net, which receive supervision signals

y_q , y_d , and y_ω , respectively. Intuitively, learning y_q and y_d helps Task Net to learn a joint feature space of image quality and disease diagnosis, leading to more informative F_θ with multiple semantics [46]. Additionally, we adopt GAB to conduct both channel- and spatial-wise attention to highlight the potentially correlated features from F_θ to each supervision signal [47]. GABs extract features specific to y_q and y_d to encourage the construction of generalizable F_θ [36], while filtering irrelevant channels and spatiality to capture desired f_θ^ω correlated with y_ω . Finally, Task Net obtains f_θ^d and f_θ^ω containing diagnosis-related information and desired knowledge of image quality and disease diagnosis, respectively, for further explicit utilization.

Meta-knowledge assistance block. In addition, we designed the MAB to explicitly explore the potential assistance of f_θ^ω for disease diagnosis, where f_θ^ω learns joint semantics of disease diagnosis and image quality. MAB first filters channels of f_θ^ω uncorrelated with disease diagnosis via a channel-wise attention block to ensure the desired functionality of f_θ^ω . It then utilizes the concatenated feature f_θ^{d*} of f_θ^d and a filtered f_θ^ω for the diagnosis. This “filter and utilize” pattern is critical for quality-aware diagnosis since f_θ^ω may contain undesired or even misleading information. Furthermore, the explicit utilization encourages the models to be aware of the potential benefits and assistance of image quality in the diagnosis. Finally, the objective function \mathcal{L}_θ of Task Net \mathcal{M}_θ is denoted as

$$\mathcal{L}_\theta = \mathcal{L}_D(\mathcal{M}_\theta^d(x), y_d) + \mathcal{L}_Q(\mathcal{M}_\theta^q(x), y_q) + \mathcal{L}_\Omega(\mathcal{M}_\theta^\omega(x), y_\omega), \quad (1)$$

where \mathcal{L}_D , \mathcal{L}_Q and \mathcal{L}_Ω represent the loss functions for disease diagnosis, image quality prediction, and knowledge co-embedding learning, respectively. We would like to clarify that our focus is on the design philosophy rather than specific implementation details, and our proposed design is compatible with any attention modules such as [48, 49]. We argue that a practical and appropriate design is the key to solving the IQAD problem. In this paper, we adopt a classic attention module CBAM [47] rather than other modern attention modules to demonstrate the effectiveness of our design. Additionally, we conduct experiments to explore the significance of the design of Task Net, i.e., the existence of GAB and MAB, in the experiment section.

3.2.2 Meta Learner to Co-embed Knowledge

As discussed in Section 3.1, the current quality labels do not capture the diversity of degradations and their corresponding impact on diagnostic semantics. We overcome this limitation by Meta Learner, which constrains the semantics and ensures the effectiveness of f_θ^ω via *joint-encoding masking* and *meta-auxiliary learning*.

Joint-encoding masking. We design the joint-encoding masking approach to ensure the effectiveness and semantic

constraints of f_θ^ω by constraining Meta Learner to provide a discriminative y_ω for both quality and diagnosis labels. To achieve this, we select y_ω from $\mathcal{M}_\phi(x)$ based on a joint encoding $y_{d,q}$ of y_d and y_q , and create a binary mask $\mathcal{B}_{y_{d,q}}$ that assigns a value of 1 to positions related to $y_{d,q}$ and 0 to others. Consider the case that both y_q and $y_d \in \{0, 1\}$, and the length of $\mathcal{M}_\phi(x)$ is 4. In this case, the joint label encoding will be $y_{d,q} \in \{00, 01, 10, 11\}$, and $\mathcal{B}_{y_{d,q}}$ will be $[1, 0, 0, 0]$, $[0, 1, 0, 0]$, $[0, 0, 1, 0]$ and $[0, 0, 0, 1]$, respectively. Further, y_ω is derived from $\mathcal{M}_\phi(x)$ as follows:

$$y_\omega = \mathcal{B}_{y_{d,q}}(\mathcal{M}_\phi(x)). \quad (2)$$

At each optimization step, the Meta Learner optimizes $\mathcal{M}_\phi(x)$ using a specific slice y_ω , encouraging the construction of a discriminative auxiliary embedding among different label combinations. This mechanism ties f_θ^ω to the joint semantic information of image quality and disease diagnosis. By working in conjunction with meta-auxiliary learning, it contributes to adaptively and effectively capturing the interplay between image quality and disease diagnosis, ultimately leading to improved model performance.

Meta-auxiliary learning. Meta Learner is designed to optimize Task Net to construct f_θ^ω effectively for diagnostic purposes. This is essentially a ‘‘learning-to-learn’’ problem, for which we adopt meta-auxiliary learning to optimize the effectiveness of f_θ^ω by providing a beneficial and desired y_ω [45]. The idea is that training Task Net on y_ω should result in improved performance. Therefore, Meta Learner \mathcal{M}_ϕ should optimize ϕ in the direction that minimizes the objective function \mathcal{L}_θ of Task Net \mathcal{M}_θ , i.e., $\arg \min_\phi \mathcal{L}_\theta$. To accomplish this, we first perform a one-step pseudo update with learning rate α to simulate the effect of y_ω on θ as

$$\tilde{\theta} := \theta - \alpha \nabla_\theta [\mathcal{L}_D(\mathcal{M}_\theta^d(x), y_d) + \mathcal{L}_Q(\mathcal{M}_\theta^q(x), y_q) + \mathcal{L}_\Omega(\mathcal{M}_\theta^\omega(x), y_\omega)]. \quad (3)$$

Then, we use the second derivative trick to update ϕ , similar to [50, 51]:

$$\phi := \phi - \beta \nabla_\phi [\mathcal{L}_D(\mathcal{M}_\theta^d(x), y_d) + \mathcal{L}_Q(\mathcal{M}_\theta^q(x), y_q) + \mathcal{R}(\mathcal{B}_{y_{d,q}}(\mathcal{M}_\phi(x)))], \quad (4)$$

where β is the learning rate of Meta Learner, and $\mathcal{R}(\cdot)$ is a regularization term to avoid auxiliary label embedding collapse by increasing the entropy of $\mathcal{M}_\phi(x)$ [45, 52].

Meta Learner leverages gradient information to adaptively generate $\mathcal{M}_\phi(x)$ to assist Task Net in constructing f_θ^ω . Minimizing \mathcal{L}_θ allows Meta Learner to provide useful $\mathcal{M}_\phi(x)$ for solving diagnosis and quality assessment tasks [51], which in turn ensures that f_θ^ω contains the desired semantic information derived from learning y_ω . Furthermore, this meta-optimization procedure takes into account the MAB where f_θ^ω is utilized for diagnosis, which further enhances the effectiveness of f_θ^ω .

4. Experiment

4.1. Dataset and Implementation Details

Dataset. We employ five widely-used public datasets from four imaging modalities for three diagnosis tasks. In the data preprocessing step, all images were resized to 256×256 and then individually normalized to zero mean and unit variance in intensity values. Table 1 presents the statistics of the datasets, and their settings are briefly described as follows: **DRAC** [59]: OCTA images are graded into three levels, standing for no DR, non-proliferative DR (NPDR), and proliferative DR (PDR), respectively. Each image is annotated with HQ or LQ. **DeepDR** [5]: Same with DRAC, fundus images are graded into three levels of DR and two levels of image quality. **EyeQ** [7]: The diagnosis annotation of fundus images is the same as DeepDR, while the quality of images is labeled as good (HQ), usable (LQ-U), and poor (LQ-P). **CT-IQAD**: We combine computed tomography images from COVID-X [60] and SARS [61] to make up this dataset. Since COVID-X images are collected in the wild, e.g., downloaded from papers, their quality cannot be guaranteed, and thus we label these images LQ and others are HQ. **CXR-IQAD**: This dataset is sampled from child chest X-ray images of CXR-P [62] and adult CXR images of CXR14 [63]. The diagnosis label is Normal and Pneumonia. We downsample images with a bicubic kernel to simulate low-dose CXR images [64]. Thus, images have three parts, i.e., child images (HQ), low-dose child images (LQ-C), and low-dose adult images (LQ-A). During training, we use the original data splits for DeepDR and DRAC. We split EyeQ into three parts: 60% for training, 10% for validation, and 30% for testing. As for the rest of the datasets, we divide them into 70% for training, 10% for validation, and 20% for testing. To evaluate the performance fairly, we used three metrics: the area under the ROC curve (AUC), accuracy (ACC), and macro F1-score (F1).

Implementation. We adopt VGG16 [65] as backbones in the experiment, and the focal loss [66] and entropy loss for training. The learning rate and weight decay are kept the same for all datasets at 0.01 and 0.0005, respectively, and we adopt SGD as the optimizer. The number of training epochs is set as 100 for EyeQ and 200 for the rest of the others. We followed the train-validation-test paradigm to report results for all datasets except DRAC, i.e., we selected the best models in validation and evaluated them on the test set. For DRAC, we reported the last epoch result because its official data split does not include a validation set. Further details about this sub-section can be found in the appendix.

4.2. Comparison with Other Methods

Experiment setting. We compare MKCNet with recent state-of-the-art methods from three groups: ophthalmic disease diagnosis (ODS) methods [19, 20, 53, 54], multi-task

Table 1. Statistics for lung disease diagnosis datasets (left group) and ophthalmic disease diagnosis datasets (right group).

Dataset	CT-IQAD			CXR-IQAD				Label	DRAC			DeepDR			EyeQ			
	ALL	HQ	LQ	ALL	HQ	LQ-C	LQ-A		ALL	HQ	LQ	ALL	HQ	LQ	ALL	HQ	LQ-U	LQ-P
Normal	1,149	800	349	2,582	792	791	999	NoDR	545	444	101	914	532	382	20,680	12,308	4,553	3,747
Covid-19	1,197	800	397	-	-	-	-	NPDR	344	303	41	974	506	468	7,533	4,405	1,690	1,438
Pneumonia	-	-	-	5,704	2,137	2,136	1,431	PDR	108	93	15	112	50	62	558	104	191	353
ALL	2,346	1,600	746	8,286	2,929	2,927	2,430	ALL	997	840	157	2,000	1,088	912	28,789	16,817	6,434	5,538

Table 2. Comparison with state-of-the-art approaches in disease diagnosis.

Dataset	CT-IQAD			CXR-IQAD			DRAC			DeepDR			EyeQ			Avg.		
	AUC	ACC	F1	AUC	ACC	F1	AUC	ACC	F1	AUC	ACC	F1	AUC	ACC	F1	AUC	ACC	F1
MMCNN [53]	86.57	86.54	86.51	89.13	89.07	91.82	82.41	77.20	66.18	76.39	79.50	54.40	69.36	73.22	35.48	80.77	81.11	66.88
BIRA-Net [54]	88.46	88.46	88.21	89.88	90.76	93.23	84.00	77.46	67.00	80.67	78.50	64.10	80.13	72.50	56.42	84.63	81.54	73.79
GREEN [20]	92.54	92.52	92.44	90.06	91.30	93.67	84.37	73.83	63.03	81.14	76.50	64.88	81.23	77.21	59.08	85.87	82.27	74.64
CABNet [19]	91.01	91.03	90.79	88.38	89.73	92.51	81.94	72.28	59.07	82.39	77.75	55.97	79.21	76.95	53.48	84.59	81.55	70.36
Mixstyle [55]	91.06	91.03	90.99	88.85	90.52	93.14	85.03	73.32	60.91	81.89	73.50	60.34	82.10	74.92	<u>61.25</u>	85.79	80.66	73.33
AugMix [56]	89.17	89.10	89.26	88.21	89.79	92.58	84.94	70.47	57.66	<u>84.79</u>	<u>80.00</u>	54.74	77.38	77.00	44.26	84.90	81.27	67.70
DDAIG [57]	92.72	92.74	92.54	88.19	90.28	93.00	84.95	77.72	62.91	79.40	75.75	51.66	70.71	74.64	37.84	83.19	82.23	67.59
Mixup [58]	92.12	92.10	92.04	88.98	90.70	93.28	82.41	72.80	62.84	81.69	76.00	63.57	82.35	<u>77.74</u>	63.07	85.51	81.87	<u>74.89</u>
QGNet [16]	92.71	92.74	92.50	89.51	90.70	93.22	84.38	63.73	47.80	84.21	73.50	<u>64.66</u>	82.40	74.73	60.36	86.64	79.08	71.71
MAXL [45]	93.35	93.38	93.16	87.99	89.86	92.66	<u>87.24</u>	<u>77.72</u>	<u>69.20</u>	81.59	71.25	57.58	80.64	73.10	58.81	86.16	81.06	74.28
CANet [36]	93.36	93.38	93.19	<u>90.51</u>	91.49	93.78	86.70	75.13	64.38	81.83	77.50	62.74	81.89	77.20	60.26	<u>86.86</u>	<u>82.94</u>	74.87
MT-Net [39]	92.56	92.52	92.51	90.07	90.58	93.05	84.63	73.06	61.89	82.40	76.75	64.51	<u>82.66</u>	75.97	59.92	86.46	81.78	74.38
MTMR-Net [38]	<u>94.46</u>	<u>94.44</u>	<u>94.37</u>	89.69	<u>91.61</u>	<u>93.96</u>	85.79	76.42	65.12	82.30	63.75	53.61	79.71	76.57	60.50	86.39	80.56	73.51
DETACH [37]	92.53	92.52	92.41	89.52	90.64	93.16	84.81	76.94	65.24	83.61	71.75	62.14	81.68	74.80	59.64	86.43	81.33	74.52
MKCNet (Ours)	95.73	95.73	95.65	90.71	91.91	94.12	88.68	82.38	73.87	86.58	80.00	54.73	83.83	78.01	59.68	89.11	85.61	75.61

and auxiliary learning (MTAL) methods [16,36–39,45], and other adoptable methods [55–58]. The chosen methods can be adopted in IQAD without or with only simple modifications, and their brief descriptions are in the appendix.

Comparison of results. Columns 3-5 of Table 2 show the quantitative results of ophthalmic disease diagnosis. MKCNet outperforms other methods in overall performance and improves at least two of the AUC, ACC, and F1 metrics on all datasets. Typically, ODS methods do not consider using image quality information and only design modules for achieving good performance on HQ images. Thus, most of them perform non-ideally on the IQAD problem, where numerous LQ images exist. For example, GREEN [20] considers relationships among categories to achieve a robust diagnosis, yet it suffers in IQAD. It is because image degradations significantly disturb the model to construct the foundations of relationship exploration among categories, i.e., correlations between images and labels. We modified methods from the MTAL set to learn image quality information as an auxiliary task. As observed in the last column of Table 2, a rough sense is that the average performance of MTAL methods outperforms OSD methods. Such advance illustrates the importance of image quality information in the IQAD problem. Nevertheless, MKCNet outperforms these MTAL methods due to its explicit image quality-assist modules and meta-learning techniques that bypass the problem of coarse quality labels and provide sufficient information to diagnose problems. It is worth mentioning that CANet explicitly explores the internal relationships between two diseases, and we modified it to explore the assistance of quality assessment in diagnosis. Al-

though the coarse quality label only contributes scant information, CANet still surpasses other methods, demonstrating the importance of explicit utilization mechanisms. We also tested MKCNet on two different lung disease diagnosis tasks on two different imaging modalities to verify its generalization ability. Columns 1-2 of Table 2 present the quantitative results of lung disease diagnosis. As expected, MKCNet significantly outperforms other methods even when facing diagnosis tasks on another organ with different imaging modalities, owing to its unique modeling for the IQAD problem. MKCNet outperforms the methods ranked second with a significant margin, achieving an average AUC, ACC, and F1 improvement of **2.25%**, **2.67%**, and **0.72%**, respectively, highlighting its effectiveness.

4.3. Ablation Study

Contribution of each component. We first validate the effects of the two critical components in our method by removing Meta Learner (denoted as w/o Meta, which will leverage f_{θ}^q to assist diagnosis) and removing the explicit quality information utilization mechanism (denoted as w/o MAB, where Meta Learner will still optimize Task Net). We also include a baseline model with the same basic experimental setting as ours, denoted as Vanilla. As shown in Figure 4, removing either component results in decreased performance of MKCNet on all three metrics, but the performance is still better than Vanilla. This result is reasonable and suggests that these two components play complementary roles in addressing the IQAD problem.

Influence of LQ images. We evaluated the effectiveness of MKCNet in handling different numbers of LQ images




Figure 4. Ablation study for the effect of proposed components.

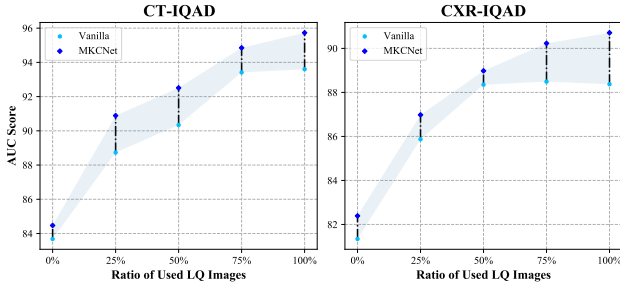


Figure 5. Performance of MKCNet and Vanilla on different ratios of LQ images used in training on CT-IQAD and CXR-IQAD.

during training to demonstrate its generalization ability, as illustrated in Figure 5. To conduct this experiment, we divided LQ images into four equal parts and trained Vanilla and MKCNet with 0%, 25%, 50%, 75%, and 100% of LQ images while retaining the same number of HQ images. As expected, the performance of both Vanilla and MKCNet increased with the addition of LQ images, but the performance gap between the two methods became more pronounced. Furthermore, we compared the performance of MKCNet with Vanilla on LQ images. As shown at the top of Figure 6, MKCNet achieved a significant performance improvement over Vanilla. These experiments highlight the importance of learning from LQ images and demonstrate that MKCNet can more effectively leverage LQ images by utilizing quality labels compared to Vanilla.

Effect of designs in Task Net. We conducted an in-depth analysis to demonstrate the effectiveness of designs in Task Net. First, we discussed modifications to investigate the need for the GAB and the way of explicit utilization mechanism. We replaced the GAB with fully connected layers and kept other components unchanged, denoted as MKCNet-FC. Additionally, we designed MKCNet-Con, which diagnosed based on the concatenation of f_{θ}^d and the entire meta co-embedding feature f_{θ}^{ω} , without filtering diagnosis-irrelevant channels. As shown in the middle of Figure 6, the performance drop of these variants indicates the effectiveness and necessity of these designs. The results suggest two critical issues: a) the GAB contributes to desired feature representation construction, and b) designing a practical utilization mechanism is crucial for IQAD.

Analysis of Meta Learner. We first conducted experiments to verify the effectiveness of joint-encoding masking, which is used to select the auxiliary label embedding




Figure 6. Histogram of AUC score analysis on four datasets. **Top:** Comparison on LQ images. **Middle:** The influence of Task Net design. **Bottom:** The impact of different quality information.

y_{ω} . We compared the results of using different masking strategies and found that only using the image quality label for masking decreased performance significantly, as shown at the bottom of Figure 6. This illustrates the importance of joint-encoding masking, which uses joint label encoding to ensure that the semantics of y_{ω} correlate with both image quality and disease diagnosis. We further analyzed the cosine similarity among the auxiliary label embedding loss gradient (AuxGrad), diagnosis loss gradient (D-Grad), and quality assessment loss gradient (IQ-Grad) on the shared backbone in Task Net. Figure 7 shows that y_{ω} provides information related to both diagnosis and image quality, while there is little similarity among the loss gradients of the disease diagnosis and image quality assessment branch. We also observed that the label type used in the embedding masking affects the auxiliary loss gradient and its correlation with the corresponding task. If we use the diagnosis label to obtain y_{ω} , the gradient of that learning branch will be highly related to disease diagnosis but uncorrelated with image quality assessment, and vice versa. These analysis support the effectiveness of knowledge co-embedding features, which contain the desired knowledge correlated with both quality assessment and disease diagnosis.

Importance of auxiliary label embedding y_{ω} . We conducted t-SNE analysis on the output vector $\mathcal{M}_{\phi}(x)$ of Meta Learner and the final diagnosis feature f_{θ}^{d*} to visually understand the effect of auxiliary label embedding. As shown in Figure 8, Task Net successfully learned discriminative features for disease diagnosis, and these features tended to be clustered based on image quality and disease diagnosis categories. However, Vanilla only focused on diagnosis labels and obscured features of images with different image qualities. This indicates the effectiveness of MKCNet in leveraging diagnosis and quality labels, i.e., it comprehen-




Figure 7. Gradient cosine similarity analysis on Task Net.

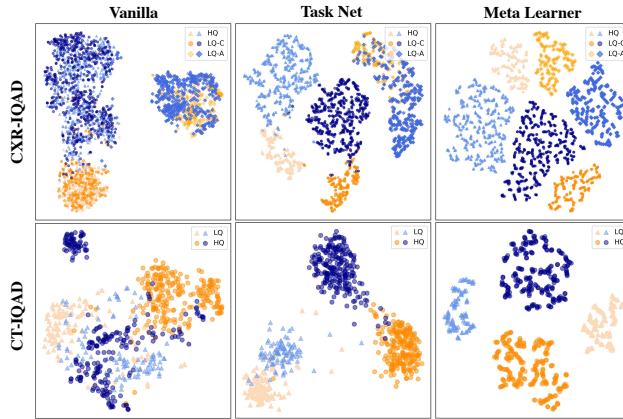


Figure 8. The t-SNE analysis. Similar colors and distinct shapes indicate the same diagnosis and different quality labels.

sively considers both diagnosis and quality information to diagnose. From the visualization of $\mathcal{M}_\phi(x)$, we can observe two facts: a) these output vectors are clustered by the joint encoding of image quality and diagnosis, and b) feature points are dispersed inside each cluster. The first fact illustrates the justification of Meta Learner, which learns fused and discriminative information regarding image quality and disease diagnosis. The latter fact illustrates that the auxiliary label embedding is diverse with images and may provide adaptively helpful information to f_θ^ω .

Qualitative analysis. Finally, we performed class activation mappings (CAMs) using [67] to qualitatively investigate how MKCNet performs on DARC with varying image qualities. The results, shown in Figure 9, demonstrate that both Vanilla and MKCNet effectively attend to the HQ images. However, Vanilla fails to deal with degradations in the images, resulting in misleading signs. In the second row of the figure, for instance, Vanilla focuses on the fake vessel disappearing in the low-signal area caused by threshold degradation. In the third row, it focuses on the false abnormality of the optic disc due to segment degradation and low-signal area. In contrast, MKCNet is more robust in dealing with such degradations and ignores mis-




Figure 9. Qualitative analysis via CAM visualization.

leading clues while capturing the foveal avascular zone or vessel density as desired. Moreover, the CAM visualization of features in MKCNet demonstrates how quality information can assist in diagnosis. In the second and third rows, f_θ^d initially perceives the area of interest as similar to the area of interest of f_θ^q ; i.e., focuses on degradations, which indicates that it has been misled. However, after leveraging information from f_θ^ω , it refocuses on the vessel area density or foveal avascular zone. This illustrates the importance of quality information and the effectiveness of our model in providing assistance or calibration to the diagnosis.

5. Conclusion

In this paper, we address the significant yet often overlooked problem of IQAD by reconsidering the value of LQ images and quality labels. To achieve this, we propose a novel MKCNet consisting of two subnets. Task Net explicitly explores the potential benefits of quality labels in diagnosis by leveraging knowledge co-embedding features, while Meta Learner learns to optimize these features to ensure their desired effectiveness and semantics. The experiments show the superior performance of our method on five datasets spanning four widely used medical imaging modalities. Our method effectively overcomes the challenges of leveraging quality labels, offering a more practical and cost-effective solution than conventional approaches. By assessing the value of LQ images and quality labels, we highlight their importance in diagnosis tasks. Furthermore, it is expected to pave the way for future research in leveraging quality labels and LQ images, stimulating the development of robust, accurate, and practical diagnostic models that can effectively handle real-world challenges in clinical settings.

Acknowledgement. This work was supported by Shenzhen Science and Technology Innovation Committee Fund (Project No. SGDX20210823103201011 and JCYJ20180507182410327), Hong Kong Innovation and Technology Fund (Project No. ITS/028/21FP) and HKUST-WeBank Joint Lab (Project No. WEB19EG01-Q). We would like to express our sincere gratitude to Huajun Zhou, Haibo Jin, Hao Jiang, Zhengrui Guo, Dawei Yang, Jingran Su, Kang Zhou, Lai Tian, and for their valuable discussions.

References

- [1] Xiaoxuan Liu, Livia Faes, Aditya U Kale, Siegfried K Wagner, Dun Jack Fu, Alice Bruynseels, Thushika Mahendiran, Gabriella Moraes, Mohith Shamdas, Christoph Kern, et al. A comparison of deep learning performance against health-care professionals in detecting diseases from medical imaging: a systematic review and meta-analysis. *The Lancet Digital Health*, 1(6):e271–e297, 2019. 1
- [2] Tao Li, Wang Bo, Chunyu Hu, Hong Kang, Hanruo Liu, Kai Wang, and Huazhu Fu. Applications of deep learning in fundus images: A review. *Medical Image Analysis*, 69:101971, 2021. 1
- [3] Hao Jiang, Yanning Zhou, Yi Lin, Ronald CK Chan, Jiang Liu, and Hao Chen. Deep learning for computational cytology: A survey. *Medical Image Analysis*, page 102691, 2022. 1
- [4] Myura Nagendran, Yang Chen, Christopher A Lovejoy, Anthony C Gordon, Matthieu Komorowski, Hugh Harvey, Eric J Topol, John PA Ioannidis, Gary S Collins, and Mahiben Maruthappu. Artificial intelligence versus clinicians: systematic review of design, reporting standards, and claims of deep learning studies. *British Medical Journal*, 368, 2020. 1
- [5] Ruhan Liu, Xiangning Wang, Qiang Wu, Ling Dai, Xi Fang, Tao Yan, Jaemin Son, Shiqi Tang, Jiang Li, Zijian Gao, et al. Deepdrid: Diabetic retinopathy—grading and image quality estimation challenge. *Patterns*, page 100512, 2022. 1, 5
- [6] Ian C Holmen, Sri Meghana Konda, Jeong W Pak, Kyle W McDaniel, Barbara Blodi, Kimberly E Stepien, and Amitha Domalpally. Prevalence and severity of artifacts in optical coherence tomographic angiograms. *JAMA Ophthalmology*, 138(2):119–126, 2020. 1
- [7] Huazhu Fu, Boyang Wang, Jianbing Shen, Shanshan Cui, Yanwu Xu, Jiang Liu, and Ling Shao. Evaluation of retinal image quality assessment networks in different color-spaces. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 48–56. Springer, 2019. 1, 2, 5
- [8] Emma Beede, Elizabeth Baylor, Fred Hersch, Anna Iurchenko, Lauren Wilcox, Paisan Ruamviboonsuk, and Laura M Vardoulakis. A human-centered evaluation of a deep learning system deployed in clinics for the detection of diabetic retinopathy. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, pages 1–12, 2020. 1
- [9] Haofeng Liu, Haojin Li, Xiaoxuan Wang, Heng Li, Mingyang Ou, Luoying Hao, Yan Hu, and Jiang Liu. Understanding how fundus image quality degradation affects cnn-based diagnosis. In *2022 44th Annual International Conference of the IEEE Engineering in Medicine & Biology Society*, pages 438–442. IEEE, 2022. 1, 2, 3
- [10] Ziyi Shen, Huazhu Fu, Jianbing Shen, and Ling Shao. Modeling and enhancing low-quality retinal fundus images. *IEEE Transactions on Medical Imaging*, 40(3):996–1006, 2020. 1, 2, 3
- [11] TANG Ziqi, Xi Wang, Anran Ran, Fangyao Tang, Yu Cai, Haoxuan Che, Dawei Gabriel Yang, Luyang Luo, Quande Liu, Yiu Lun Wong, et al. Using deep learning for assessing image-quality of 3d macular scans from spectral-domain optical coherence tomography. *Investigative Ophthalmology & Visual Science*, 63(7):204–F0051, 2022. 2
- [12] Yaxin Shen, Bin Sheng, Ruogu Fang, Huating Li, Ling Dai, Skylar Stolte, Jing Qin, Weiping Jia, and Dinggang Shen. Domain-invariant interpretable fundus image quality assessment. *Medical Image Analysis*, 61:101654, 2020. 2
- [13] Junhao Hu, Chenyang Zhang, Kang Zhou, and Shenghua Gao. Chest x-ray diagnostic quality assessment: How much is pixel-wise supervision needed? *IEEE Transactions on Medical Imaging*, 41(7):1711–1723, 2022. 2, 3
- [14] Xinqiang Zhou, Yicheng Wu, and Yong Xia. Retinal image quality assessment via specific structures segmentation. In *International Workshop on Ophthalmic Medical Image Analysis*, pages 53–61. Springer, 2020. 2
- [15] Fabian SL Yui, Raman Dutt, Tom MacGillivray, Baljean Dhillon, Miguel Bernabeu, and Niall Strang. Rethinking retinal image quality: Treating quality threshold as a tunable hyperparameter. In *International Workshop on Ophthalmic Medical Image Analysis*, pages 73–83. Springer, 2022. 2
- [16] Kang Zhou, Zaiwang Gu, Annan Li, Jun Cheng, Shenghua Gao, and Jiang Liu. Fundus image quality-guided diabetic retinopathy grading. In *Computational Pathology and Ophthalmic Medical Image Analysis*, pages 245–252. Springer, 2018. 2, 6
- [17] Luyang Luo, Lequan Yu, Hao Chen, Quande Liu, Xi Wang, Jiaqi Xu, and Pheng-Ann Heng. Deep mining external imperfect data for chest x-ray disease screening. *IEEE transactions on medical imaging*, 39(11):3583–3594, 2020. 2
- [18] Sebastian Ruder. An overview of multi-task learning in deep neural networks. *arXiv preprint arXiv:1706.05098*, 2017. 2
- [19] Along He, Tao Li, Ning Li, Kai Wang, and Huazhu Fu. Cabnet: category attention block for imbalanced diabetic retinopathy grading. *IEEE Transactions on Medical Imaging*, 40(1):143–153, 2020. 2, 5, 6
- [20] Shaoteng Liu, Lijun Gong, Kai Ma, and Yefeng Zheng. Green: a graph residual re-ranking network for grading diabetic retinopathy. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 585–594. Springer, 2020. 2, 5, 6
- [21] Xi Wang, Fangyao Tang, Hao Chen, Carol Y Cheung, and Pheng-Ann Heng. Deep semi-supervised multiple instance learning with self-correction for dme classification from oct images. *Medical Image Analysis*, 83:102673, 2023. 2
- [22] Xi Wang, Hao Chen, Caixia Gan, Huangjing Lin, Qi Dou, Efstratios Tsougenis, Qitao Huang, Muyan Cai, and Pheng-Ann Heng. Weakly supervised deep learning for whole slide lung cancer image analysis. *IEEE transactions on cybernetics*, 50(9):3950–3962, 2019. 2
- [23] Xi Wang, Hao Chen, Huiling Xiang, Huangjing Lin, Xi Lin, and Pheng-Ann Heng. Deep virtual adversarial self-training

- with consistency regularization for semi-supervised medical image classification. *Medical image analysis*, 70:102010, 2021. 2
- [24] Long Bai, Liangyu Wang, Tong Chen, Yuanhao Zhao, and Hongliang Ren. Transformer-based disease identification for small-scale imbalanced capsule endoscopy dataset. *Electronics*, 11(17):2747, 2022. 2
- [25] Luyang Luo, Dunyuan Xu, Hao Chen, Tien-Tsin Wong, and Pheng-Ann Heng. Pseudo bias-balanced learning for debiased chest x-ray classification. In *Medical Image Computing and Computer Assisted Intervention—MICCAI 2022: 25th International Conference, Singapore, September 18–22, 2022, Proceedings, Part VIII*, pages 621–631. Springer, 2022. 2
- [26] Robert Geirhos, Patricia Rubisch, Claudio Michaelis, Matthias Bethge, Felix A Wichmann, and Wieland Brendel. Imagenet-trained cnns are biased towards texture; increasing shape bias improves accuracy and robustness. In *International Conference on Learning Representations*, 2018. 2
- [27] Yaniv Ovadia, Emily Fertig, Jie Ren, Zachary Nado, David Sculley, Sebastian Nowozin, Joshua Dillon, Balaji Lakshminarayanan, and Jasper Snoek. Can you trust your model’s uncertainty? evaluating predictive uncertainty under dataset shift. *Advances in Neural Information Processing Systems*, 32, 2019. 2
- [28] Robert Geirhos, Jörn-Henrik Jacobsen, Claudio Michaelis, Richard Zemel, Wieland Brendel, Matthias Bethge, and Felix A Wichmann. Shortcut learning in deep neural networks. *Nature Machine Intelligence*, 2(11):665–673, 2020. 2
- [29] Haofeng Liu, Heng Li, Huazhu Fu, Ruoxiu Xiao, Yunshu Gao, Yan Hu, and Jiang Liu. Degradation-invariant enhancement of fundus images via pyramid constraint network. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 507–516. Springer, 2022. 2
- [30] Bingyu Yang, He Zhao, Lvchen Cao, Hanruo Liu, Ningli Wang, and Huiqi Li. Retinal image enhancement with artifact reduction and structure retention. *Pattern Recognition*, 133:108968, 2023. 2
- [31] Zhuo Deng, Yuanhao Cai, Lu Chen, Zheng Gong, Qiqi Bao, Xue Yao, Dong Fang, Wenming Yang, Shaochong Zhang, and Lan Ma. Rformer: Transformer-based generative adversarial network for real fundus image restoration on a new clinical benchmark. *IEEE Journal of Biomedical and Health Informatics*, 26(9):4645–4655, 2022. 2
- [32] Weiwen Zhang, Dawei Yang, Carol Y Cheung, and Hao Chen. Frequency-aware inverse-consistent deep learning for oct-angiogram super-resolution. In *Medical Image Computing and Computer Assisted Intervention—MICCAI 2022: 25th International Conference, Singapore, September 18–22, 2022, Proceedings, Part II*, pages 645–655. Springer, 2022. 2
- [33] Andrés Romero, Luc Van Gool, and Radu Timofte. Unpaired real-world super-resolution with pseudo controllable restoration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 798–807, 2022. 2
- [34] Zhuang Liu, Hungju Wang, Tinghui Zhou, Zhiqiang Shen, Bingyi Kang, Evan Shelhamer, and Trevor Darrell. Exploring simple and transferable recognition-aware image processing. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022. 2
- [35] Xin Yang, Yi Lin, Zhiwei Wang, Xin Li, and Kwang-Ting Cheng. Bi-modality medical image synthesis using semi-supervised sequential generative adversarial networks. *IEEE Journal of Biomedical and Health Informatics*, 24(3):855–865, 2019. 2
- [36] Xiaomeng Li, Xiaowei Hu, Lequan Yu, Lei Zhu, Chi-Wing Fu, and Pheng-Ann Heng. Canet: cross-disease attention network for joint diabetic retinopathy and diabetic macular edema grading. *IEEE Transactions on Medical Imaging*, 39(5):1483–1493, 2019. 2, 4, 6
- [37] Haoxuan Che, Haibo Jin, and Hao Chen. Learning robust representation for joint grading of ophthalmic diseases via adaptive curriculum and feature disentanglement. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 523–533. Springer, 2022. 2, 6
- [38] Lihao Liu, Qi Dou, Hao Chen, Jing Qin, and Pheng-Ann Heng. Multi-task deep model with margin ranking loss for lung nodule analysis. *IEEE Transactions on Medical Imaging*, 39(3):718–728, 2019. 2, 6
- [39] Qingyu Chen, Yifan Peng, Tiarnan Keenan, Shazia Dharssi, Elvira Agro, Wai T Wong, Emily Y Chew, Zhiyong Lu, et al. A multi-task deep learning model for the classification of age-related macular degeneration. *AMIA Summits on Translational Science Proceedings*, 2019:505, 2019. 2, 6
- [40] Long Bai, Sihang Chen, Mingyang Gao, Leila Abdelrahman, Manal Al Ghamdi, and Mohamed Abdel-Mottaleb. The influence of age and gender information on the diagnosis of diabetic retinopathy: based on neural networks. In *2021 43rd Annual International Conference of the IEEE Engineering in Medicine & Biology Society*, pages 3514–3517. IEEE, 2021. 2
- [41] Yi Lin, Luyan Liu, Kai Ma, and Yefeng Zheng. Seg4reg+: Consistency learning between spine segmentation and cobb angle regression. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 490–499. Springer, 2021. 2
- [42] Ming-En Hsieh and Vincent Tseng. Boosting multi-task learning through combination of task labels-with applications in ecg phenotyping. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 7771–7779, 2021. 2
- [43] Luyang Luo, Hao Chen, Yongjie Xiao, Yanning Zhou, Xi Wang, Varut Vardhanabhuti, Mingxiang Wu, Chu Han, Zaiyi Liu, Xin Hao Benjamin Fang, et al. Rethinking annotation granularity for overcoming shortcuts in deep learning-based radiograph diagnosis: A multicenter study. *Radiology: Artificial Intelligence*, 4(5):e210299, 2022. 3
- [44] Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy Hospedales. Learning to generalize: Meta-learning for do-

- main generalization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018. 4
- [45] Shikun Liu, Andrew Davison, and Edward Johns. Self-supervised generalisation with meta auxiliary learning. *Advances in Neural Information Processing Systems*, 32, 2019. 4, 5, 6
- [46] Bernardino Romera Paredes, Andreas Argyriou, Nadia Berthouze, and Massimiliano Pontil. Exploiting unrelated tasks in multi-task learning. In *Artificial intelligence and statistics*, pages 951–959. PMLR, 2012. 4
- [47] Sanghyun Woo, Jongchan Park, Joon-Young Lee, and In So Kweon. Cbam: Convolutional block attention module. In *European Conference on Computer Vision*, pages 3–19, 2018. 4
- [48] Zequn Qin, Pengyi Zhang, Fei Wu, and Xi Li. Fcanet: Frequency channel attention networks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 783–792, 2021. 4
- [49] Abhijit Guha Roy, Nassir Navab, and Christian Wachinger. Recalibrating fully convolutional networks with spatial and channel “squeeze and excitation” blocks. *IEEE Transactions on Medical Imaging*, 38(2):540–549, 2018. 4
- [50] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *International Conference on Machine Learning*, pages 1126–1135. PMLR, 2017. 5
- [51] Yabin Zhang, Hui Tang, and Kui Jia. Fine-grained visual categorization using meta-learning optimization with sample selection of auxiliary data. In *European Conference on Computer Vision*, pages 233–248, 2018. 5
- [52] Shikun Liu, Stephen James, Andrew J Davison, and Edward Johns. Auto-lambda: Disentangling dynamic task relationships. *arXiv preprint arXiv:2202.03091*, 2022. 5
- [53] Kang Zhou, Zaiwang Gu, Wen Liu, Weixin Luo, Jun Cheng, Shenghua Gao, and Jiang Liu. Multi-cell multi-task convolutional neural networks for diabetic retinopathy grading. In *2018 40th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, pages 2724–2727. IEEE, 2018. 5, 6
- [54] Ziyuan Zhao, Kerui Zhang, Xuejie Hao, Jing Tian, Matthew Chin Heng Chua, Li Chen, and Xin Xu. Bira-net: Bilinear attention net for diabetic retinopathy grading. In *2019 IEEE International Conference on Image Processing*, pages 1385–1389. IEEE, 2019. 5, 6
- [55] Kaiyang Zhou, Yongxin Yang, Yu Qiao, and Tao Xiang. Domain generalization with mixstyle. In *International Conference on Learning Representations*, 2020. 6
- [56] Dan Hendrycks, Norman Mu, Ekin Dogus Cubuk, Barret Zoph, Justin Gilmer, and Balaji Lakshminarayanan. Augmix: A simple data processing method to improve robustness and uncertainty. In *International Conference on Learning Representations*, 2019. 6
- [57] Kaiyang Zhou, Yongxin Yang, Timothy Hospedales, and Tao Xiang. Deep domain-adversarial image generation for domain generalisation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 13025–13032, 2020. 6
- [58] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. *arXiv preprint arXiv:1710.09412*, 2017. 6
- [59] DRAC. The diabetic retinopathy analysis challenge website. <https://drac22.grand-challenge.org>. Accessed October 20, 2022. 5
- [60] Jinyu Zhao, Yichen Zhang, Xuehai He, and Pengtao Xie. Covid-ct-dataset: a ct scan dataset about covid-19. *arXiv preprint arXiv:2003.13865*, 490, 2020. 5
- [61] Eduardo Soares, Plamen Angelov, Sarah Biaso, Michele Higa Froes, and Daniel Kanda Abe. Sars-cov-2 ct-scan dataset: A large dataset of real patients ct scans for sars-cov-2 identification. *MedRxiv*, 2020. 5
- [62] Daniel S Kermany, Michael Goldbaum, Wenjia Cai, Carolina CS Valentim, Huiying Liang, Sally L Baxter, Alex McKeown, Ge Yang, Xiaokang Wu, Fangbing Yan, et al. Identifying medical diagnoses and treatable diseases by image-based deep learning. *Cell*, 172(5):1122–1131, 2018. 5
- [63] 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*, pages 2097–2106, 2017. 5
- [64] Liming Xu, Xianhua Zeng, Zhiwei Huang, Weisheng Li, and He Zhang. Low-dose chest x-ray image super-resolution using generative adversarial nets with spectral normalization. *Biomedical Signal Processing and Control*, 55:101600, 2020. 5
- [65] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 5
- [66] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2980–2988, 2017. 5
- [67] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. Learning deep features for discriminative localization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2921–2929, 2016. 8