Distilling Knowledge from Object Classification to Aesthetics Assessment

Jingwen Hou, Henghui Ding, Weisi Lin, Fellow, IEEE, Weide Liu, Yuming Fang, Senior Member, IEEE

Abstract—In this work, we point out that the major dilemma of image aesthetics assessment (IAA) comes from the abstract nature of aesthetic labels. That is, a vast variety of distinct contents can correspond to the same aesthetic label. On the one hand, during inference, the IAA model is required to relate various distinct contents to the same aesthetic label. On the other hand, when training, it would be hard for the IAA model to learn to distinguish different contents merely with the supervision from aesthetic labels, since aesthetic labels are not directly related to any specific content. To deal with this dilemma, we propose to distill knowledge on semantic patterns for a vast variety of image contents from multiple pre-trained object classification (POC) models to an IAA model. Expecting the combination of multiple POC models can provide sufficient knowledge on various image contents, the IAA model can easier learn to relate various distinct contents to a limited number of aesthetic labels. By supervising an end-to-end single-backbone IAA model with the distilled knowledge, the performance of the IAA model is significantly improved by 4.8% in SRCC compared to the version trained only with ground-truth aesthetic labels. On specific categories of images, the SRCC improvement brought by the proposed method can achieve up to 7.2%. Peer comparison also shows that our method outperforms 10 previous IAA methods.

Index Terms-deep learning, image aesthetics assessment

I. INTRODUCTION

MAGE aesthetics is significant in a variety of scenarios, including image recommendation [1], image editing [2], image retrieval [3], and photo management [4]. As a result, image aesthetics assessment (IAA) approaches are sought for evaluating visual aesthetic experiences automatically. State-ofthe-art (SOTA) methods [5]-[9] are mainly based on deep learning, which relies on neural networks for learning to extract aesthetic features (i.e., features for distinguishing different aesthetic levels) in a data-driven manner. As some works [10]–[12] have suggested, semantic information can help to improve the effectiveness of a deep IAA model. Kao et al. [11] intuitively explained that semantic information is useful for IAA since humans need to understand the content of an image before assessing it. Besides such an intuitive explanation, we believe that the reason why semantic information is useful in IAA is that semantic information can make up for the shortcoming of the abstractness of aesthetic labels.

Essentially, IAA can be regarded as a process that maps different image contents into different aesthetic levels (as shown in Fig. 1). And image contents relevant to aesthetics are described by aesthetic features, which represent different



Fig. 1: Motivation for distilling knowledge from POC models. IAA by a deep model can be depicted as a process that maps contents of a test image into contents corresponding to different aesthetic levels known to the IAA model. In the given example, image pairs with boundaries of the same color represent relevant mappings. Since POC models can capture semantic patterns for a large variety of image contents, we distill knowledge from POC models to teach an IAA model on semantic patterns for relating a larger variety of contents to a limited number of aesthetic labels.

combinations of semantic patterns relevant to IAA. For semantic patterns, we refer to a collection of pixels that are organized in a certain way so that such a pattern can be clearly identified across different images, and the combination of such patterns is connected to a certain semantic meaning (e.g., an object). And constructing discriminative aesthetic features requires sufficiently diverse semantic patterns to deal with a vast variety of contents. In the example of Fig. 1, if the IAA model does not know the semantic patterns for representing a macro-photo of a flower, the IAA model cannot confidently map the test image of a flower to the class of high-aesthetics. However, since similar aesthetic labels can refer to images with various contents, it is hard for deep models to learn semantic patterns from aesthetic labels. To make up for the handicap of aesthetic labels in providing semantic guidance (*i.e.*, guiding the IAA model to learn about semantic patterns), one instant way is to define predicting semantic information as an auxiliary task [10], [11], or use semantic information as an auxiliary input [12]. However, these approaches require extra human labels describing the semantics contained in images.

Thus, we aim to provide semantic guidance to improve an IAA model without using labels from humans. Specifically, for an IAA model that produces less discriminative aesthetic features, we wish to provide extra supervision without human labels to guide the IAA model to capture more relevant

Jingwen Hou and Henghui Ding are co-first authors. Part of the work is done by Jingwen Hou during his internship at ByteDance.

Corresponding author: Weisi Lin.



(a) Accurate matching by GSF

(b) Inaccurate matching by GSF

Fig. 2: Comparing generic semantic features (GSFs) to aesthetic features (AFs) constructed from GSFs. The score for the test image is estimated by matching it to the most similar training image with GSFs or AFs. GT scores are given in the topright corner. GSFs are likely to match visually-similar images with distinct scores, implying that some semantic patterns in GSFs may not be relevant to IAA. Thus, we take knowledge on IAA-relevant semantic patterns from GSFs instead of directly using GSFs for semantic guidance. Quantitative evaluations of GSFs and AFs on IAA by matching are given in Table IV.

semantic patterns for constructing more discriminative aesthetic features that can deal with a large variety of contents. Typically, we consider a baseline IAA model constructed with a single pre-trained backbone trained merely with aesthetic labels. Therefore, the aesthetic features of the IAA model are mostly constructed from the semantic patterns known to its selected pre-trained backbone, if its backbone fail to learn extra semantic patterns from aesthetic labels (and this is the case when training an IAA model merely with aesthetic labels, see Table V and Sec. IV-C for details). In this case, we may introduce extra pre-trained object classification (POC) models besides the selected pre-trained backbone to provide semantic guidance so that the IAA model can learn extra semantic patterns from them. Thus, one possible solution that provides semantic guidance to an IAA model with extra POC models is to encourage the IAA model to produce extra features as the extra POC models, which will force the IAA model to capture extra semantic patterns to construct extra features as the extra POC models. However, this cannot guarantee that the extra semantic patterns learned by the IAA model are relevant to the downstream IAA task, since POC models are trained on object classification instead of IAA.

For simplicity, we call features from the backbone of a POC model as generic semantic features (GSFs). We give an example in Fig. 2 to show whether GSFs can match images with similar aesthetic levels. In some cases, images with similar GSFs can have similar aesthetic scores (as shown in Fig. 2(a)). While in some other cases, though the GSFs can match images with similar visual appearances, the GSFs cannot fully distinguish images with different aesthetic levels (as shown in Fig. 2(b)). This implies that some semantic patterns represented by GSFs may not be helpful for distinguishing aesthetic levels, but can even lead to confusion. Thus, we propose to use a knowledge distiller to select IAA-relevant semantic patterns from GSFs for constructing aesthetic features. Specifically, we train an IAA model as the knowledge distiller that can directly predict image aesthetics from the GSFs, and the inputs

to the output layer of the knowledge distiller are regarded as aesthetic features. As shown in Fig. 2 (and Table IV), the aesthetic features constructed by the knowledge distiller can better distinguish different aesthetic levels than GSFs. Finally, the knowledge distiller can serve as the teacher model to provide a student IAA model with semantic guidance: the aesthetic features and predictions of the knowledge distiller are regarded as knowledge on semantic patterns distilled for IAA (i.e., teacher knowledge). The distilled knowledge is then used for imposing supervision to the student aesthetic features (i.e., input features to student's output layer) and student predictions. To ensure that the distilled knowledge can allow the student IAA model to acquire extra knowledge on semantic patterns for more discriminative aesthetic features, we choose POC models deeper or trained with more data than the student's backbone for constructing the teacher model, and verify that the teacher model has a higher IAA performance than the student. Our contributions are summarized as follows:

- We point out the dilemma of IAA caused by the abstract nature of aesthetic labels. To deal with this dilemma, we propose a KD method to allow an end-to-end singlebackbone IAA model can learn about semantic patterns relevant to IAA from multiple POC models via a knowledge distiller.
- The knowledge distiller with the combined POC models as feature extractors is an effective IAA model (*i.e.*, teacher model) which outperforms 10 previous IAA methods. Compared to the best-performed method among the contenders, the model achieves 5% higher SRCC performance.
- With the proposed KD scheme, the teacher model provides knowledge on semantic patterns to the training of a single-backbone end-to-end IAA model (*i.e.*, student model). The performance of an end-to-end IAA model can be significantly improved by 4.8% in SRCC compared to the version trained only with GT aesthetic labels. On specific categories of images, the improvement brought by the proposed KD scheme can achieve up to 7.2%. Compared to the teacher model, the student model has a 99% lower computational cost, with only 3% lower SRCC performance. Peer comparison shows that the student model also outperforms 10 previous IAA methods. Compared to the best-performed end-toend IAA model among the contenders, the student model achieves 7.1% higher SRCC performance.

II. RELATED WORKS

Semantic patterns in image aesthetics assessment (IAA). To build a robust IAA model, major efforts have been made to construct image features that distinguish different aesthetic levels (*i.e.*, aesthetic features). Assuming aesthetic levels can be distinguished by judging whether a photo follows known photography rules, early approaches attempt to predict image aesthetics from hand-crafted features following photography rules [13]–[15]. However, the number of well-defined photography rules is too limited to explain images on a large scale. Therefore, SOTA methods [5]–[9] are mainly based on

deep learning, which allows the model to learn to construct aesthetic features in a data-driven manner. However, there is still not a consensus about what a deep IAA model has learned to distinguish images with different aesthetics. In this work, we hypothesize that IAA is a process that maps different combinations of semantic patterns (represented by aesthetic features) into different aesthetic levels. Thus, the IAA model is required to recognize more diverse relevant semantic patterns for aesthetic features when more diverse image contents are needed to be dealt with. This hypothesis can explain some findings in previous works: 1) image aesthetics can be easier distinguished among images with similar GSFs [16], since it requires less semantic patterns to construct sufficiently discriminative aesthetic features; 2) image aesthetics can be accurately predicted from GSFs [7], [9], since GSFs contain semantic patterns useful for constructing aesthetic features; 3) introducing extra semantic information can help an IAA model to achieve higher performance [10]-[12], since it allows the IAA model to learn more semantic patterns for constructing aesthetic features. According to this hypothesis, we further propose methods to allow aesthetic features to be constructed with more diverse semantic patterns, so that the aesthetic features can deal with a larger variety of image contents: 1) we go beyond previous methods [7], [9] that construct aesthetic features with GSFs from a single POC model, we prepare more diverse semantic patterns with stacked GSFs from multiple POC models, which allows SOTA models to be achieved; 2) we also go beyond previous methods [10]-[12] that introduce extra semantic labels, we distill knowledge on semantic patterns from multiple POC models to provide extra supervision to a weaker student model (e.g., a singlebackbone end-to-end model), which allows us to significantly improve the performance of the student model without extra human efforts on labeling.

Knowledge distillation (KD) beyond object classification. KD has firstly been proposed by Hinton et al. [17] for object classification, which transfers knowledge in the form of softened output logits from a deep teacher model to a rather shallower student model. Later on, various methods have been proposed for KD on object classification [18]-[22]. However, KD methods designed for object classification may not be fully applicable to other tasks. Therefore, various methods have been specifically designed for different tasks. In object detection, Chen et al. [23] proposed to distill the knowledge on imbalanced classification and bounding box regression. In semantic segmentation, He et al. [24] proposed to distill the knowledge on capturing long-term dependencies. In road marking detection, Hou et al. [25] proposed to distill the knowledge on the structural relationship of road scenes. In our case, we expect to improve an IAA model from knowledge on general object classification, where IAA and general object classification are two different tasks. As discussed, aesthetic features are essential for IAA performance, while more discriminative aesthetic features need to be built upon sufficiently diverse semantic patterns. To this end, we conduct KD to distill IAA-relevant knowledge on semantic patterns from object classification models. As a relevant topic, Zhang et al. [26] has proposed to adopt a POC model for

constructing an image quality assessment (IQA) model with continual learning that can deal with various IQA scenarios, which supports our hypothesis that GSFs can be used to guide the downstream IAA model to produce more discriminative features for distinguishing image aesthetics.

III. OUR APPROACH

A. Problem Statement

We consider a typical IAA model [27] that estimates the aesthetic rating distribution directly from an image. Particularly, given the *i*-th image I_i from an IAA dataset, the IAA model $\mathcal{M}_{\theta}(\cdot)$ predicts the aesthetic rating distribution $\hat{\mathcal{D}}_i$:

$$\hat{D}_i = \mathcal{M}_{\theta}(\mathbf{I_i}).$$
 (1)

A direct way for obtaining the model $\mathcal{M}_{\theta}(\cdot)$ parameterized by θ is to directly optimize its parameter θ towards the ground-truth (GT) aesthetic rating distribution \mathcal{D}_i :

$$\theta = \arg\min_{\theta} \sum_{i=1}^{N} \mathcal{L}(\mathcal{D}_i, \hat{\mathcal{D}}_i), \qquad (2)$$

where the most commonly-used loss function for $\mathcal{L}(\cdot)$ is earth mover distance (EMD) loss [27]:

$$EMD(y,\hat{y}) = \sqrt{\frac{1}{n} \sum_{k=1}^{n} |\text{CDF}_y(k) - \text{CDF}_{\hat{y}}(k)|^2}, \quad (3)$$

where CDF_y and $\text{CDF}_{\hat{y}}$ are cumulative density function for GT distribution y and predicted distribution \hat{y} of length n, respectively.

However, this approach overlooks the abstract nature of the aesthetic labels \mathcal{D}_i . As we previously discussed, IAA can be regarded as a process that maps different combinations of semantic patterns into different aesthetic levels. On the one hand, during inference, the IAA model is required to relate various combinations of semantic patterns to the same aesthetic label. On the other hand, when training, it would be hard for the IAA model to learn to distinguish different combinations of semantic patterns merely with the supervision from aesthetic labels since aesthetic labels are not directly related to any specific contents. To make up for the abstractness of aesthetic labels, one instant way is to assign an extra semantic label to each of the training samples for IAA. Because each semantic label is a direct description of the image contents (e.g., themes), the IAA model can learn about semantic patterns related to each semantic label along with the IAA objective, and aesthetic features covering more image contents can be constructed from the semantic patterns learned from semantic labels.

Nevertheless, there are several problems with semantic labels: 1) it requires extra human efforts to assign semantic labels to each of IAA training samples; 2) it is hard to define what semantics in the image will be relevant to the downstream IAA task, and therefore, it is hard to find a standard for introducing semantic labels. Thus, our goal is to find a better representation that provides semantic guidance, so that the IAA model can learn about semantic patterns relevant to IAA



Fig. 3: Details of training a single-backbone end-to-end IAA model (*i.e.*, student model) under the supervision of the teacher knowledge distilled from POC models via a knowledge distiller. The teacher aesthetic features and teacher predictions together are deemed as teacher knowledge distilled from GSFs. The teacher aesthetic features and teacher predictions are stored for imposing supervision on the aesthetic features and predictions of the student model when training the student model. Note that the preparation of teacher knowledge is conducted separately from the training of the student model in practice (see Table XIII for the analysis on extra computational costs).



Fig. 4: LayerCAM [28] visualization for explaining how the proposed KD works in training the student model. For the same image, different GSFs tend to capture different semantic patterns. The knowledge distiller takes useful parts from GSFs to produce the teacher aesthetic feature. The teacher aesthetic features impose supervision to the student aesthetic features. Through back-propagation, the supervision from teacher aesthetic features further guides the student model to capture more IAA-relevant semantic features for producing more discriminative aesthetic features as the teacher. Details of training the knowledge distiller are given in Fig. 5.

for constructing more discriminative aesthetic features for a large variety of image contents:

$$\theta = \arg\min_{\theta} \sum_{i=1}^{N} (\mathcal{L}_a(\mathcal{D}_i, \hat{\mathcal{D}}_i) + \mathcal{L}_s(\mathcal{S}_i, \hat{\mathcal{S}}_i)), \qquad (4)$$

where S_i and \hat{S}_i are the GT and predicted representation for semantic guidance, D_i and \hat{D}_i are the GT and predicted aesthetic rating distributions, and $\mathcal{L}_s(\cdot)$ and $\mathcal{L}_a(\cdot)$ are loss functions for semantic and aesthetic guidance, respectively. Here, semantic guidance is provided by imposing an extra supervision to the training objective.

B. Knowledge Distillation for Semantic Guidance

As discussed in Sec. III-A, we aim for an extra supervision for semantic guidance besides extra semantic labels. Since POC models can recognize a vast variety of image contents, our idea is to distill knowledge from POC models on semantic patterns relevant to IAA. To cover a large variety of contents, we can take multiple POC models trained on different datasets to provide sufficiently diverse semantic patterns. And the semantic patterns can be taken from the GSFs of the selected POC models, as we earlier described in Sec. I. By combining GSFs from different POC models, we can obtain representations that can describe a large variety of contents:

$$f_C = [\bar{f}_1 | \bar{f}_2 | \dots | \bar{f}_K], \tag{5}$$

where | denotes the concatentation operation that combines pooled GSFs \bar{f}_k from K different POC models, resulting in the combined feature f_C . For extracting features from different POC models, we adopt multi-layer spatial pooling (MLSP) [7] as the pooling strategy to cover both low-level and high-level semantic information in the resulting feature, which is denoted as $\bar{f}_k = MLSP(f_k)$, where f_k is the raw GSF from the k-th POC model.

However, directly supervising the IAA model by the combined GSFs is ineffective, since not all semantic patterns represented by the combined GSFs are relevant to IAA. We further describe the relationship between the combined GSFs and aesthetic features by Fig. 4. As shown, red, green, and blue circles are sets of patterns that can be captured by different POC models. Since these models are trained with different datasets, different models may be sensitive to a different set of patterns. However, considering different POC models trained on different datasets may share similar categories of semantic patterns, these sets also have overlapped portions. For example, for classifying the same bird, some models may tend to classify by the beak, while some other models may tend to



Fig. 5: Training of the knowledge distiller. Table III presents performance of the teacher IAA model constructed with the knowledge distiller and different POC models.

TABLE I: Comparison of GSFs for IAA from different models trained with ImageNet on the train-test split from [29]. Full-resolution images are used.

Ν		Results			
ResNet18	ResNet50	ResNet101	SRCC	PLCC	Acc
1			0.669	0.678	78.7%
	1		0.713	0.718	80.2%
		1	0.720	0.724	80.3%
·			0.718	0.723	80.3%
	1	1	0.730	0.734	80.6%
1		1	0.724	0.727	80.5%
·		,	0.730	0.734	80.7%

classify by wings. Nevertheless, it is possible that all models tend to classify an object as bird when they see the object with feather is in the sky. However, since these POC models are not trained for IAA, the patterns contained in these sets may not all be relevant to IAA. As we show in Fig. 2, when GSFs contain semantic patterns not relevant to IAA, GSFs may match two visually-similar images but with distinct aesthetic scores. In Fig. 4, we present the patterns relevant to IAA as yellow.

Therefore, we propose a KD method based on a knowledge distiller trained separately. To be specific, we train an IAA model as the knowledge distiller with the combined GSF f_C . As shown in Fig. 5, the knowledge distiller is constructed with a batch normalization layer followed by three linear layers with ReLU activations. To train the knowledge distiller, we use EMD loss (Eq. 3). The unified whole of selected POC models and the trained knowledge distiller can also be viewed as a teacher model (upper-stream of Fig. 3), and the single-backbone end-to-end model (lower-stream of Fig. 3) is then viewed as the student model that learns to imitate the teacher. Thus, the knowledge distiller $C(\cdot)$ predicts a feature f_t and an aesthetic rating distribution $\hat{\mathcal{D}}_t$ from GSFs given by different POC models:

$$\{f_t, \hat{\mathcal{D}}_t\} = \mathcal{C}(f_C),\tag{6}$$

where $\{f_t, \hat{\mathcal{D}}_t\}$ are deemed as teacher knowledge, including a teacher aesthetic feature f_t and a teacher prediction $\hat{\mathcal{D}}_t$.

TABLE II: Comparison of GSFs for IAA from ResNeXt101 trained with different data on the train-test split from [29]. Full-resolution images are used.

Train	ing		Results		
ImageNet1	IG^2	SWSL ³	SRCC	PLCC	Acc
1			0.723	0.726	80.4%
	1		0.756	0.760	81.8%
		1	0.755	0.758	81.5%
			0.764	0.766	81.9%
	1	1	0.770	0.773	82.1%
1		1	0.763	0.765	81.9%
·			0.773	0.775	82.0%

Thus, the teacher knowledge is directly used for supervising the student model. Accordingly, we formulate the KD loss for training the student model from Eq. 4:

$$\mathcal{L}_{\mathcal{KD}}(\hat{\mathcal{D}}_t, \hat{\mathcal{D}}_s, f_t, f_s) = EMD(\hat{\mathcal{D}}_t, \hat{\mathcal{D}}_s) + MSE(f_t, f_s), \quad (7)$$

where \mathcal{D}_s , f_s denote the student prediction and the student aesthetic feature, $EMD(\cdot)$ and $MSE(\cdot)$ refer to EMD loss and mean squared error (MSE) loss respectively. Accordingly, f_t here is expected to provide semantic guidance in the context of IAA, and the student model is also expected to summarize how the teacher predicts $\hat{\mathcal{D}}_t$ from f_t . The architecture of the student model follows a succinct design (lower-stream of Fig. 3), which is constructed with a single CNN backbone followed by a fully-connected (FC) network for aligning the student aesthetic feature to the same size as the teacher aesthetic feature. The aesthetic labels are finally predicted from the student aesthetic feature by an FC softmax layer.

Note that in Eq. 7, the teacher aesthetic feature f_t is constructed from GSFs, and the student aesthetic feature f_s is constructed from semantic patterns relevant to IAA captured by the student's own backbone. And we hypothesize that the teacher aesthetic feature f_t are more discriminative than the student aesthetic feature f_s for effective semantic guidance. Generally, to prepare more discriminative teacher aesthetic features, the POC models for constructing the teacher model should be selected according to the student model. For a student model with a known backbone, we could expect the constituent POC models of the teacher model should: 1) deeper than the student's backbone; or 2) trained with more data than the student's backbone. By combining multiple POC models deeper or trained with more data, more diverse semantic patterns are expected to be produced than the student's pretrained backbone and more discriminative aesthetic features are expected to be created. The points above are guidelines for selecting POC models that are likely to produce more diverse semantic patterns than the student's backbone. Note that since the effectiveness of a POC model also relies on the training strategy and the quality of the training data, we are not able to appropriately select POC models merely according to the parameter size and training data size. Thus, the criterion to verify whether the teacher aesthetic features are more discriminative than the student aesthetic features is to directly compare the performance of the teacher model (i.e., knowledge distiller) to the student model without KD. When the student model without KD performs poorer, it

¹ResNeXt101_32x8d: https://paperswithcode.com/lib/timm/resnext

²ig_ResNeXt101_32x8d: https://paperswithcode.com/lib/timm/ig-resnext

³swsl_ResNeXt101_32x8d: https://paperswithcode.com/model/swsl-resnext

TABLE III: Comparison of different POC models as feature extractors of the teacher model. Experiments are conducted on the same train-test split from [29], and results show that the knowledge distiller can effectively distill information needed for IAA from GSFs.

Setting		Model backbone		Full resol	ution (~64	0×640)	Resiz	zed (300×3	300)
Setting	ResNet-v2 (BiTm) ⁴	ResNeXt101 (SWSL) ⁵	ResNeXt101 (IG) ⁶	SRCC↑	PLCC↑	Acc↑	SRCC↑	PLCC↑	Acc↑
1	✓ <i>✓</i>			0.779	0.781	82.7%	0.762	0.764	82.2%
2		\checkmark		0.756	0.758	81.6%	0.739	0.742	80.8%
3			\checkmark	0.762	0.764	81.9%	0.753	0.754	81.6%
4	·			0.788	0.789	83.0%	0.771	0.773	82.3%
5		\checkmark	\checkmark	0.775	0.777	82.2%	0.762	0.764	81.8%
6	1		\checkmark	0.792	0.792	83.0%	0.779	0.779	82.4%
7	/	<i>-</i>	<i>,</i>	0.794	0.795	83.1%	0.780	0.781	82.7%

TABLE IV: Comparison of generic semantic features (GSFs) and aesthetic features (AFs) on IAA by matching (as shown in Fig. 2) on the split from [29]. Results show that AFs can better matches images of the same score.

Feature	SRCC	PLCC	Acc
GSF	0.414	0.417	71.7%
AF	0.657	0.657	78.3%

means that its aesthetic features are less discriminative than those of the teacher. As long as the teacher model performs better than the student model, the student model can learn to construct better aesthetic features from the teacher, and the selection of POC models for the teacher model is appropriate. The points above will be experimentally discussed in Sec. IV-B and Sec. IV-C.

C. Adaptation to Smaller Input Sizes

As shown in some previous works [7], [30], resizing input images can harm the effectiveness of the IAA model due to loss in high-resolution details. However, using high-resolution images can introduce large computational costs. To cope with the trade-off between input resolution and computational costs, the proposed KD scheme is also designed to allow the student model to adapt to smaller input sizes. In training the teacher model, we adopt full-resolution inputs for feature extraction, which allows semantic patterns on high-resolution details to be preserved in the teacher knowledge. While for training the student model, we adopt resized inputs and encourage the student model to excavate high-resolution details from the resized image with the teacher knowledge. Experimental evidences are given in Table VII.

IV. EXPERIMENTS

In this section, we would like to answer the following questions by experiments:

- How to construct more discriminative aesthetic features from GSFs? (Q1)
- Does the supervision from the teacher model better than ground-truth aesthetic labels? (Q2)

- How much improvement has been made in terms of efficiency, and how much effectiveness is compromised, by comparing the student to the teacher model? (Q3)
- How much improvement has been made by the teacher and the student model compared to previous works? (Q4)

A. Experimental Settings

1) Dataset: Following previous works [6], [7], [27], [31]– [36], our experiments are performed on the AVA dataset [29]. The AVA dataset includes \sim 250,000 images, which have been scored from 1 \sim 10 by 78 \sim 594 workers. We follow previous works [6], [7], [9], [31], [32], [34], [37] to use the same traintest split⁷ [29] for our experiments. The split adopts \sim 230,000 images for training and \sim 20,000 images for testing.

2) Implementation details: Our model is implemented with PyTorch. As to the implementation of feature extraction, we refer to previous works [7], [38], [39]. For the teacher model, we train the model with batch size 512 for 12 epochs with Adam optimizer with an initial learning rate $3e \times 10^{-5}$ and divided by 10 every 3 epochs. For the student model, we train the model with batch size 16 for 12 epochs with Adam optimizer with an initial learning rate $3e \times 10^{-5}$ and divided by 10 every 3 epochs. For the student model, we train the model with batch size 16 for 12 epochs with Adam optimizer with an initial learning rate $3e \times 10^{-5}$ and divided by 10 every 3 epochs.

Following previous works [7], [9], [27], [32], [34], [36], we adopt Spearman Correlation Coefficients (SRCC), Pearson Correlation Coefficients (PLCC), and Accuracy (Acc) for evaluation. For evaluating SRCC and PLCC, we convert the aesthetic rating distributions into aesthetic scores by weighted average. As previous works [6], [7], [31]–[34], we take five as the cut-off threshold for converting aesthetic scores into binary aesthetic categories for evaluating Acc. We also use Floating Point Operations (FLOPs)⁸ to evaluate computational costs.

B. Effectiveness of Generic Semantic Features (Q1)

1) Investigating aesthetic features constructed from GSFs: To answer the question: "how to construct more discriminative aesthetic features from GSFs", we have picked two sets of GSFs to investigate how GSFs from different POC models and their combinations will impact the IAA performance:

⁴resnetv2_152x4_bitm: https://paperswithcode.com/lib/timm/big-transfer ⁵swsl_ResNeXt101_32x8d: https://paperswithcode.com/model/swsl-resnext ⁶ig_ResNeXt101_32x48d: https://paperswithcode.com/lib/timm/ig-resnext

⁷Note that the original source for the AVA dataset is no more available. The official split is originated from: https://github.com/mtobeiyf/ava_downloader/blob/master/AVA_dataset/aesthetics_image_lists/generic_test.jpgl

⁸https://github.com/Swall0w/torchstat

TABLE V: Ablation study on the use of the proposed KD scheme on the same train-test split from [29]. The results imply that the proposed KD scheme enables the student's backbone to capture more IAA-relevant semantic patterns.

Setting		PI CC ⁺	Acc↑	
Trainable backbone KD scheme			Thee	nee
×	×	0.736	0.739	81.1%
✓	×	0.735	0.737	80.8%
×		0.747	0.748	81.3%
<i>✓</i>	1	0.770	0.770	82.1%

TABLE VI: Ablation study on the two terms of the proposed KD loss (Eq. 7) on the same train-test split from [29]. Results show that both terms of the KD loss contribute significantly to the overall improvement.

Supervision		SRCC↑	PLCC ↑	Acc↑
Feature	Output			
×	×	0.735	0.737	80.8%
1	×	0.766	0.766	82.0%
×	1	0.758	0.758	81.6%
1	1	0.770	0.770	82.1%

- POC models of different architecture but trained on the same data: ResNet18, ResNet50 and ResNet101 trained on ImageNet are picked (see Table I).
- POC models of the same architecture but trained on different data: ResNeXt101 trained on different data are picked. Note that both IG version and SWSL version have included the ImageNet training data (see Table II, sources for pre-trained models are given in the footnotes).

The approaches including using POC models with larger sizes or using POC models with more training data enable the GSFs to provide more diverse semantic patterns. Combining GSFs from different POC models further allows diverse semantic patterns from different POC models to be considered for aesthetic features. Thus, the experimental results imply that the performance of different settings in Table I and Table II mainly depend on whether the semantic patterns are sufficiently diverse for constructing aesthetic features that can deal with diverse contents in the AVA dataset. Considering both the teacher model and the student model training merely with aesthetic labels construct aesthetic features from semantic patterns known to their pre-trained backbones (see Sec.IV-C for details), the teacher model consists of extra POC backbones with larger size and trained with more data, which is expected to capture more diverse semantic patterns for constructing aesthetic features covering more diverse contents.

2) Selecting POC models for KD: The targeted backbone for the student model we mainly consider is ResNeXt101 (SWSL) [40], whose baseline SRCC is 0.735 (Table V). To construct the teacher model for later experiments in KD, we have selected three different POC models trained with different data, including ResNet-v2 (BiTm) [41], ResNeXt101 (IG) [42], and the same ResNeXt101 (SWSL) as the student

TABLE VII: Ablation study with smaller-sized inputs on the same train-test split from [29]. The performance drop brought by smaller input sizes is given in red subscripts. The results show that model with KD has lower performance drop, implying that the proposed KD scheme enables the student model to adapt to smaller-sized inputs.

Input size	KD	SRCC↑	PLCC↑	Acc↑	FLOPs (G)↓
640×640	X	0.756	0.758	81.6%	134.7
300×300	X	0.736_2.6%	0.739_2.5%	81.1%	30.6
640×640		0.775	0.777	82.4%	134.7
300×300	1	0.770_0.6%	0.770 _{-0.9%}	82.1%	30.6

(sources for the selected POC models are given in the footnotes of Table III). Since ResNet-v2 (BiTm) and ResNeXt101 (IG) are deeper or trained with more data than ResNeXt101 (SWSL), the combined GSFs are expected to contain semantic patterns captured by ResNeXt101 (SWSL) and extra semantic patterns captured by ResNet-v2 (BiTm) and ResNeXt101 (IG). The experimental results on teacher models using different combinations among the three POC models are given in Table III. The results for the best setting that combines all three POC models above have achieved 0.794 in SRCC. Since the performance is much higher than the student's, the selected POC models are valid for constructing the teacher model for the designated student based on ResNeXt101 (SWSL) according to the guidelines and criterion in Sec. III-B. In later sections, "teacher model" will refer to the teacher model trained with GSFs combined from ResNet-v2 (BiTm), ResNeXt101 (IG) , and ResNeXt101 (SWSL) with SRCC 0.794. Note that in the peer comparison of Sec. IV-E, we will also use the same teacher model for the student model based on ResNet18 or ResNet50 backbone pretrained on ImageNet. The selected POC models for the teacher model are all trained on ImageNet along with extra data, and are much deeper than ResNet18 and ResNet50. The baseline SRCC performance of ResNet18 and ResNet50 are 0.721 and 0.735 (Fig. 6), which are poorer than the teacher's performance. Thus, using the same teacher model also follows the guidelines and criterion in Sec. III-B and is expected to provide effective semantic guidance.

3) Comparing GSFs to aesthetic features by matching: Based on the teacher model, we also compare the combined GSFs with the resulting aesthetic features following the matching-based approach as described in Fig. 2. In this way, we are able to see whether the aesthetic features are more discriminative in aesthetics than the source GSFs. As the results shown in Table IV, the resulting aesthetic features significantly outperforms the source GSFs, which confirms the effectiveness of the proposed knowledge distiller.

C. Effectiveness of Knowledge Distillation (Q2)

To answer the question "does the supervision from the teacher model better than ground-truth aesthetic labels", we conduct ablation studies to investigate different settings of the proposed KD scheme. Specifically, we firstly investigate whether a single backbone is capable of learning teacher knowledge and adapting to smaller input sizes under the proposed KD scheme. To this end, we set up baseline models

Model	Input Resolution	Params (M)	\mid FLOPs (G) \downarrow	$ $ SRCC \uparrow	PLCC \uparrow	Acc \uparrow
Teacher Composite*	~640×640	1853.2	2940.5	0.794	0.795	83.1%
Teacher Composite*	300×300	1853.2	653.8	0.780	0.781	82.7%
Student ResNeXt101 (SWSL)	640×640	88.8	134.7	0.775	0.777	82.4%
Student ResNeXt101 (SWSL)	300×300	88.8	30.6	0.770	0.770	82.1%

TABLE VIII: Comparison between teacher and student models on the same train-test split from [29]. The results show that the student model sacrifices marginal effectiveness for a greater improvement in efficiency compared to the teachers.

*: combines features from pre-trained ResNeXt101 (SWSL), ResNeXt101 (IG) and ResNetv2 (BiTm)

TABLE IX: Comparison of results on different sub-categories of the end-to-end IAA model based on ResNeXt101(SWSL) with (\checkmark) or without (\varkappa) KD. The results are obtained from the split from [29], and we choose sub-categories with top 8 portions in the testing set for the evaluation. The percentage of improvements brought by the proposed KD scheme is labeled in red subscripts. The results show that the effectiveness of the proposed KD scheme is more obvious on some sub-categories.

Categories	Portion	KD	SRCC↑	PLCC↑	Acc↑
Still life	17.60	X	0.695	0.706	77.7%
Sun me	17.0%	1	0.745+7.2%	$0.751_{\pm 6.3\%}$	79.0%
A robitocturo	15 50%	×	0.735	0.735	81.6%
Alcintecture	15.5%	1	0.761+3.6%	$0.759_{\pm 3.2\%}$	82.7%
I andecene	15 20%	×	0.777	0.776	83.8%
Lanuscape	15.5%	1	0.805+3.6%	$0.801_{\pm 3.3\%}$	85.6%
Animala	12 70%	x	0.734	0.734	80.4%
Allinais	15.7%	1	0.763+4.0%	$0.767_{\pm 4.6\%}$	81.7%
	12 00%	x	0.689	0.693	82.8%
Fortraiture	13.0%	1	0.737+6.9%	0.736 _{+6.2%}	83.6%
	12.6%	×	0.742	0.738	79.0%
Piorai	12.070	1	0.779 _{+5.0%}	$0.774_{\pm 4.9\%}$	80.4%
Cityscope	12 5%	×	0.749	0.749	81.6%
Chyscape	12.370	1	0.775+3.6%	$0.774_{\pm 3.4\%}$	81.9%
Food	12 5%	×	0.718	0.726	80.0%
1.000	12.570	1	$0.764_{+6.4\%}$	$0.771_{+6.1\%}$	82.0%
0 11	100.00	X	0.735	0.737	80.8%
Overall	100.0%	1	0.770 _{+4.8%}	$0.770_{+4.6\%}$	82.1%

trained directly with GT aesthetic rating distributions in the situation when the ResNeXt101 (SWSL) backbone is trainable or non-trainable. Then we introduce the proposed KD scheme as comparisons. The results are given in Table V.

As the results suggest, without the proposed KD scheme, similar results are obtained for both trainable and untrainable backbone settings. This suggests that even we give the backbone the freedom of learning on new semantic patterns for more discriminative aesthetic features, the IAA model has not been improved. This supports our motivation that aesthetic labels are too abstract to guide the neural network to learn about semantic patterns. As a result, the IAA model can only learn to create aesthetic features with semantic patterns already-known to its pre-trained backbone, and merely learn new semantic patterns for improving the discriminative power of the resulting aesthetic features. Since our teacher model also constructs aesthetic features from semantic patterns alreadyknown to the backbones of the selected POC models, if we introduce extra POC models deeper or trained with more data than the pre-trained backbone of the student model, it is expected that the teacher model will construct more discriminative aesthetic features than the student model. By encouraging the student aesthetic features to be close to the teacher aesthetic features in training, the student model is guided to capture more relevant semantic patterns for more discriminative aesthetic features as its teacher (Fig. 4).

By introducing the proposed KD scheme, both trainable and untrainable backbone settings present substantial improvements in performance. The results have the following implications: 1) when the backbone is set untrainable, the student model with KD learns to construct more discriminative aesthetic features with semantic patterns already-known to its pre-trained backbone; 2) when the backbone is trainable, the backbone of the student model with KD learns relevant semantic patterns for more discriminative aesthetic features. This confirms the effectiveness of the proposed KD scheme.

The KD loss (Eq. 7) can be further divided into feature supervision (term 2 of Eq. 7) and output supervision (term 1 of Eq. 7). We detail our experiments to show the contribution of each individual term in the proposed KD loss. We set the whole network trainable in this experiment, and the model is directly supervised by GTs when the output supervision is off. The results are presented in Table VI. As the results show, both terms substantially contribute to the performance of the student model, and the student model achieves the best performance when both terms are adopted.

Note that the student model takes 300×300 resized inputs while it is supervised by teacher knowledge distilled from fullresolution inputs. As discussed in Sec. III-C, this means we expect the student model to be adaptive to the smaller-sized inputs (learn full-resolution knowledge on semantic patterns from resized inputs). To confirm, we also train a student model with almost full-resolution (640×640) inputs as a comparison. Specifically, we first pad all input images to 800×800 (where 800×800 is the largest size in the dataset), and center-crop them into 640×640 to maintain the aspect ratio. The results are presented in Table VII. The results show that the proposed KD method also enables the student to be adaptive to smallersized inputs, which allows the performance drop brought by resizing ($640 \times 640 \rightarrow 300 \times 300$) to be decreased from 2.6%

Method	Backbone	Input size	SRCC↑	PLCC↑	Acc↑
DMA-Net ₁₅ [6]	AlexNet	224×224	-	-	75.4%
MNA-CNN ₁₆ [31]	VGG16	224×224	-	-	77.1%
NIMA [†] ₁₇ [27]	ResNet50	224×224	0.654	0.662	78.6%
APM ₁₇ [37]	ResNet101	500×500	0.709	-	80.3%
Hosu <i>et al.</i> [‡] ₁₉ [7]	InceptionResNet-v2	640×640	0.756	0.757	81.7%
Zeng <i>et al.</i> ₁₉ [34]	ResNet101	384×384	0.719	0.720	80.8%
GPF-CNN ₁₉ [32]	ResNet18	224×224	0.671	0.682	80.3%
Hou <i>et al.</i> [‡] ₂₀ [9]	InceptionResNet-v2	640×640 & 320×320	0.751	0.753	81.7%
Ours (Teacher model [‡])	Composite*	$\sim 640 \times 640$	0.794	0.795	83.1%
Ours (Student model)	ResNeXt101 (SWSL)	300×300	0.770	0.770	82.1%
Ours (Student model)	ResNet50	300×300	0.745	0.745	81.4%
Ours (Student model)	ResNet18	300×300	0.719	0.722	80.5%

TABLE X: Comparison with reported results of SOTAs on the same train-test split from [29]. Best results are presented in blue, and the second best results are presented in red. Publication years are given in subscripts.

*: combines features from pre-trained ResNeXt101 (SWSL), ResNeXt101 (IG) and ResNetv2 (BiTm)

[†]: we re-implement the model with ResNet50 and evaluate the model on the same train-test split as other presented methods. [‡]: not an end-to-end model

TABLE XI: Comparison with reported results of SOTAs on the same train-test split from [43]. Best results are presented in blue, and the second best results are presented in red. Publication years are given in subscripts.

Method	Backbone	Input size	SRCC↑	PLCC↑	Acc↑
PA_IAA ₂₀ [33] HLA-GCN ₂₁ [36]	DenseNet121 ResNet101	299×299 300×300	0.666 0.665	- 0.687	82.9% 84.6%
Ours (Teacher model) Ours (Student model) Ours (Student model) Ours (Student model) Ours (Student model)	Composite* ResNeXt101 (SWSL) ResNet50 ResNet18	$ \begin{array}{ } \sim 640 \times 640 \\ \hline 300 \times 300 \\ 300 \times 300 \\ 300 \times 300 \\ \hline 300 \times 300 \\ \end{array} $	$ \begin{array}{r} 0.732\\ \overline{0.701}\\ 0.677\\ 0.652 \end{array} $	0.751 0.722 0.698 0.677	- <u>85.3%</u> <u>84.9%</u> 84.1% 83.5%

*: combines features from pre-trained ResNeXt101 (SWSL), ResNeXt101 (IG) and ResNetv2 (BiTm)

to 0.6% in SRCC (0.756 \rightarrow 0.736 *vs.* 0.775 \rightarrow 0.770), which leads to FLOPs to be saved by 77.2% (134.7 \rightarrow 30.6). This suggests that the trained student model is adaptive to lower image resolution and change of aspect ratio after learning with teacher knowledge.

To show the impact of the proposed KD scheme on the performance of the student model on different categories of images, we further present the results on specific categories of the baseline model without KD and the student model with KD. The results are presented in Table IX. As the results show, compared to the overall improvement 4.8% on SRCC, the improvement brought by the proposed KD scheme can reach 7.2% for specific categories. This further confirms the effectiveness of the proposed KD scheme.

D. Comparing Student with Teacher Model (Q3)

To answer the question "how much improvement has been made in terms of efficiency, and how much effectiveness is compromised, by comparing the student to the teacher model", we compare the efficiency and effectiveness between teacher and student models. The results are presented in Table VIII. Compared to the teacher model, the student model with 640×640 input size saves FLOPs by 95% (2940.5 \rightarrow 134.7) with 2.4% drop in SRCC (0.794 \rightarrow 0.775). By shrinking the input size to 300×300 , the FLOPs is further saved by 77.2% (134.5 \rightarrow 30.6) with only 0.6% loss in SRCC (0.775 \rightarrow 0.770).

E. Comparison with State-of-the-arts (Q4)

To answer the question "how much improvement has been made by the teacher and the student model compared to previous methods", we compare the teacher model and the student model with previous relevant methods [6], [7], [9], [27], [31]-[34], [36], [37]. We take their reported results for the comparison, and the results are presented in Table X. Among these contenders, the reported results of DMA-Net [6], MNA-CNN [31], Zeng et al.'s method [34], APM [37], Hou et al.'s method [9], GPF-CNN [32] and Hosu et al.'s method [7] are evaluated on the train-test split from [29], while PA_IAA [33] and HLA-GCN [36] are evaluated on the train-test split from [43]⁹. We experimentally observed that the results on the split from [43] typically have a higher accuracy while a lower SRCC compared to the results on the split from [29] (see more detailized discussions in Sec. V-A). To fairly compare our method with the contenders, we separately compare the performance results of our method on the same train-test split as the contenders. Note that the reported results of NIMA have evaluated neither on the split from [29] nor on the split from [43]. Therefore, we re-implement NIMA and evaluate on the split from [29]. For the comparison on the split from [29], the results are given in Table X; for the comparison on the split from [43], the results are given in Table XI.

⁹https://github.com/BestiVictory/ILGnet/tree/local/data/AVA1



(c) PLCC results on fixed split evaluation and random split cross-validation with ResNet18 backbone.

Fig. 6: Demonstration of generalizability of the proposed KD scheme across different train-test splits. Fixed A and Fixed B denote the train-test split from [29] and [43], respectively. $CV1 \sim CV12$ denote twelve cross-validation splits with non-overlapping testing sets.

For our implementation, the ResNeXt101-based setting adopts the SWSL version [40], and the ResNet50 and the ResNet18 based settings utilize ImageNet pre-trained versions¹⁰ as the student model backbones. As shown in the results, the teacher model based on combined GSFs significantly outperforms the selected contenders in terms of SRCC, PLCC and accuracy. Specifically, the relevant method by Hosu et al. [7] based on InceptionResNet-v2 MLSP features, which performs best in terms of SRCC and PLCC among all contenders, is substantially surpassed by the teacher model by 5%. Additionally, the proposed KD scheme helps to significantly improve the model efficiency while preserving the model effectiveness. As shown in Table X and Table XI, though a performance drop can be observed comparing the teacher model to the student model, the student model based on ResNeXt101 still outperforms the contenders in terms of SRCC, PLCC and accuracy, including the most recent methods [9], [33], [36]. Comparing the student model to other end-toend models among the contenders, our student model achieves 0.770 in SRCC, which is 7.1% higher than the best performed end-to-end models among the contenders (Zeng et al.'s method [34] in Table X).

V. FURTHER DISCUSSION

A. Performance Variations

As mentioned in Sec. IV-E, performance variations can be observed when the same model is evaluated on different traintest splits. As a very relevant topic, image quality assessment (IQA) methods [44]-[48] typically conduct K-fold crossvalidation or take the average or median results evaluated on a number of different random train-test splits (a.k.a. sessions) to reduce performance variations as much as possible. However, for the problem of IAA, the adopted datasets by IAA typically have a much greater scale than IQA datasets. For example, the commonly-adopted IQA dataset LIVE-IQA [49] contains only 779 images. On the contrary, the commonly-adopted IAA dataset AVA benchmark contains \sim 255,000 images. This essential difference on dataset scales makes previous works on IAA mostly evaluate the methods only on a fixed traintest split [6], [7], [9], [27], [31], [32], [34], [37] rather than several random splits. However, evaluation on a fixed split potentially causes worries about generalizability (e.g., the method only works on a certain split). Additionally, several versions of train-test splits have been adopted in previous works (as mentioned in Sec. IV-E), which potentially cause problems of unfairness when comparing the results on different

¹⁰ https://github.com/Cadene/pretrained-models.pytorch

Model	Input	Output	Loss	SRCC↑	PLCC↑	Acc↑
Baseline ResNet18	300×300 image	Rating distribution	EMD loss	0.702	0.706	80.0%
Multi-modal ResNet18	300×300 image Two-hot semantic vector	Rating distribution	EMD loss	0.709	0.712	80.2%
Multi-task ResNet18	300×300 image	Rating distribution Two-hot semantic vector	Multi-task loss	0.714	0.716	80.3%
KD ResNet18 (proposed)	300×300 image	Rating distribution	KD loss	0.719	0.722	80.7%

TABLE XII: Comparison between the proposed KD scheme and other approaches for semantic guidance. Methods are evaluated on the split from [29]. Results show that the proposed KD scheme is more effective than using extra semantic labels.



Fig. 7: Demonstration of performance variations brought by randomness in training. The teacher model has been run for 10 times on the split from [29], and the SRCC, PLCC and accuracy results are plotted in separate box plots. The range for SRCC, PLCC and accuracy are [0.792, 0.795], [0.792, 0.795] and [0.830, 0.833], respectively. Results show the performance variations brought by randomness in training are marginal.

train-test splits. Therefore, in this section, we discuss whether the effectiveness of the proposed KD scheme is an outcome of random variations, and whether the effectiveness of the proposed KD scheme only works for a certain train-test split.

To begin with, we want to find out the actual range of performance variations brought by different splits. Such variations can be estimated by directly train and test the same model on different splits. However, due to the randomness in the training of deep learning models, even with the same split, if we separately train a model with the same setting, the results are possibly different. Therefore, we need to separately consider the performance variations brought by train-test split difference alone, and the performance variations brought by randomness in training. That is:

$$\delta_{exp} = \delta_{split} + \delta_{training},\tag{8}$$

where δ_{exp} is the performance variations observed from experiments that evaluate the same model on different train-test splits, δ_{split} is the actual performance variation brought by using different train-test split, and $\delta_{training}$ is the performance variation brought by randomness in training. To evaluate the scale of δ_{split} , we first estimate the scale of $\delta_{training}$ by running the proposed teacher model 10 times on the same split. In Fig. 7, the results show that the variations (difference between the highest and the lowest results) caused by randomness in training are around 0.003 for SRCC, PLCC and accuracy.

Then we estimate the scale of performance variations brought by using different train-test splits. To this end, we adopt both fixed split evaluation and cross-validation with random splits. For fixed split evaluation, we apply the split

from [29] and the split from [43], which both left 20,000 images for testing. For cross-validation with random splits, we also remain 20,000 images for testing. Since there are \sim 255,000 images in total, twelve cross-validation splits with non-overlapping testing sets are prepared. Then we evaluate the results when a single-backbone end-to-end IAA model is cooperated with or without the proposed KD scheme. The results are presented in Fig. 6. As the results demonstrated, the proposed KD scheme introduces performance gains for all selected student backbones across all train-test splits. For the ResNet18-based model, the ranges of improvements are [0.011,0.024], [0.010,0.023], [0.5%,1.7%] for SRCC, PLCC, and Acc, respectively. For the ResNet50-based model, the ranges of improvements are [0.018,0.026], [0.015,0.041], [0.4%,1.3%] for SRCC, PLCC, and Acc, respectively. For the ResNext101-based model, the ranges of improvements are [0.019, 0.038], [0.017,0.036], [0.8%,1.7%] for SRCC, PLCC, and Acc, respectively. All improvements are larger than the random variation 0.003, and thus, the improvements on all experimented models on all splits are regarded as significant.

B. Comparing Different Strategies for Semantic Guidance

The proposed KD scheme is designed to provide semantic guidance to the training of an IAA model. Two alternative approaches for introducing semantic guidance are to adopt semantic labels as auxiliary inputs, or train the model to predict semantic labels as an auxiliary task. To compare the effectiveness of the proposed KD scheme with the alternative approaches for introducing semantic guidance, we take the semantic labels provided by the AVA dataset [29] (two labels per image and 66 distinct semantic labels in total) for building models that are compared with the model trained with the proposed KD scheme. Therefore, the following four models are built for the comparison:

1) Baseline ResNet18: A ResNet18-based end-to-end IAA model is trained with EMD loss subject to aesthetic label prediction supervised by GT aesthetic labels, which follows the same setting as the baselines in our previous experiments in Sec. IV-C.

2) Multi-modal ResNet18: The two semantic labels of each image in the AVA dataset is represented with a two-hot vector. Based on the baseline model, a separate branch constructed with two stacking fully-connected layers is introduced to encode the two-hot semantic vector into semantic features. Then the encoded semantic feature is merged with the visual feature encoded by the ResNet18 for predicting the aesthetic



Fig. 8: Comparison of the capability of different models on extracting the main subjects. We include the selected POC models and IAA models trained with or without (w/o) KD in this comparison. Every two rows form a group of examples, where the top row in each group presents source images and layerCAMs [28] obtained from the selected models; the bottom row of each group presents GT main subject masks for the source images and masks derived from the corresponding layerCAMs (with 70 percentile of the layerCAMs as the segmentation threshold). The results show that the IAA model with KD better locates main subjects than POC models and the IAA model without KD, implying teacher knowledge from POC models allows the student model to capture semantic patterns more relevant to subject areas. Further quantitative results are given in Fig. 10.



Fig. 9: LayerCAM examples of cases that do not have a clear main subject. For these cases, the IAA model trained with KD is still able to locate the most appealing parts.

label. Since only aesthetic labels are predicted, the model is trained directly with EMD loss.

3) Multi-task ResNet18: Instead of taking the two-hot semantic vector as an auxiliary input, the multi-task ResNet18 setting is designed to predict the two-hot semantic vector. Specifically, based on the baseline model, the multi-task ResNet18 model not only predicts the aesthetic labels according to the visual features from the ResNet18 backbone, but also predict the two-hot semantic vector with the visual feature. Aesthetic label prediction is learned subject to the EMD loss as previous settings; for the semantic prediction, the model is trained subject to binary cross-entropy (BCE) loss. The multitask loss is given as:

$$\mathcal{L}_{\mathcal{MT}}(\hat{\mathcal{D}}_a, \mathcal{D}_a, \hat{v}_s, v_s) = EMD(\hat{\mathcal{D}}_a, \mathcal{D}_a) + BCE(\hat{v}_s, v_s),$$
(9)

where \mathcal{D}_a denotes a GT aesthetic label (in the form of distribution), v_s represents a GT semantic label (in the form of two-hot vector), $\hat{\mathcal{D}}_a$ denotes a predicted aesthetic label, and \hat{v}_s denotes a predicted semantic label.

4) KD ResNet18: This setting denotes the ResNet18-based IAA model trained with the proposed KD scheme as introduced in Sec. III. All above-mentioned models are trained with the same setting as Sec. IV-A. The results are presented in Table XII. As the table shows, with extra semantic labels, performance gains can be observed from the results of the Multi-modal ResNet18 model or the Multi-task ResNet18. This confirms that semantic guidance can help IAA models to learn how to measure image aesthetics better. However, even with extra semantic labels, the Multi-modal ResNet18 model or the Multi-task ResNet18 are inferior to the setting with KD, this indicates that the proposed KD scheme can more effectively provide semantic guidance than using extra semantic labels, which supports our earlier discussion on weaknesses of using human semantic labels in Sec. III-A.

C. Capability of Extracting the Main Subjects

Since it is a widely accepted opinion that an aestheticallypleasant image should be able to lead the attention of the observer to the main subject [50], semantic patterns for aesthetic features should be relevant to the subject area. Therefore, a more explainable way for measuring the performance of a deep IAA model is to evaluate the capability of a deep IAA model on extracting main subjects from aesthetically-pleasant images. To this end, we collect a number of aestheticallypleasant images and manually label their main subjects. Then for different IAA models, we can derive the predicted main subjects from their class activation maps and compare the accuracy of extracting main subjects across different models.

Specifically, we collect those aesthetically-pleasant images from the AVA testing set under the split from [29]. Images with average scores greater than 6.5 are regarded as aestheticallypleasant images. Then main subjects are firstly detected with a salient object detection model [51]. We find that some types of images, such as landscapes, may not have a clear main subject. For these cases, it may be hard to conclude a definite salient

Setting	Remark	Model	Input size	Cost per input (in FLOPs(G))	# Epochs	Cost per input × # Epochs (in FLOPs(G))
1		ResNeXt101 (SWSL)	(3, 640, 640)	134.7	1	134.7
2	Feature extraction with POC models	ResNeXt101 (IG)	(3, 640, 640)	1280	1	1280.0
3		ResNet-v2	(3, 640, 640)	1525.8	1	1525.8
4	Train the knowledge distiller	in Fig.5	(1, 23424)	0.3	12	3.6
5	Distill teacher knowledge	in Fig.5	(1, 23424)	0.1	1	0.1
	2944.2					
6	Train the end-to-end IAA model with KD	ResNeXt101 (SWSL)	(3, 300, 300)	91.8	12	1101.2
7	Test the end-to-end IAA model with KD	ResNeXt101 (SWSL)	(3, 300, 300)	30.6	1	30.6
8	Train the end-to-end IAA model w/o KD	ResNeXt101 (SWSL)	(3, 300, 300)	91.8	12	1101.2
9	Test the end-to-end IAA model w/o KD	ResNeXt101 (SWSL)	(3, 300, 300)	30.6	1	30.6

TABLE XIII: Comparison between extra computational costs brought by the proposed KD scheme and the computational costs for training or testing an end-to-end IAA model.

object detection result. Therefore, we firstly select images whose detected main subjects occupy $10\%\sim30\%$ of its area. Then we manually remove images with poor salient object detection results. Finally, 504 images with salient object masks are prepared, where the salient object masks are regarded as GT labels for the main subjects. Note that for those removed cases, the IAA model trained with the proposed KD scheme is still able to locate their most appealing part as demonstrated in Fig. 9. While these cases do not have a commonly-accepted guideline to outline their main subjects, and therefore, these cases are not included in the evaluation of this section.

The model trained with the proposed KD scheme is then compared to the version trained without KD and the selected POC models in terms of locating the main subjects. Some selected cases are given in Fig. 8. We firstly extract layerCAMs [28] from the selected models. Then masks for the main subjects are derived from the layerCAMs by setting the 70 percentile of the layerCAM as the threshold for segmenting the regions of the main subjects. We can observe from Fig. 8 that: 1) different POC models tend to have different focuses, implying different POC models have sensitivities to different semantic patterns; 2) the IAA model without KD captures semantic patterns less relevant to the subject area, which supports our claims that aesthetic labels are too abstract to guide the deep model to learn about semantic patterns for constructing aesthetic features; 3) the IAA model with KD supervised by knowledge on semantic patterns tend to have a better focus on subject areas than the version without KD, implying that KD allows the student model to capture semantic patterns more relevant to the subject area; 4) the focused subject areas of the IAA model with KD are further improved from those of the POC models, which implies that the knowledge distiller can build aesthetic features upon GSFs.

Besides qualitatively analyzing the impact of teacher knowledge on the student model, to confirm the judgments in qualitative analysis, we further quantitatively evaluate the performances of the selected models on extracting main subjects with the collected dataset. To this end, we measure mean intersection-over-union (mIoU) between the masks derived from layerCAMs and the GT masks. Since the mIoU results are varied subject to different segmentation thresholds, we plot the change of mIoU subject to segmentation thresholds from 5 to 95 percentile of the layerCAMs. According to Fig. 10, from best to worst, the performance on main subject detection



Fig. 10: Curves for depicting mIoU varying by segmentation thresholds from 5 to 95 percentile. Different curves present the performance of different models on extracting the main subjects on aesthetically-pleasant images in terms of mean intersection-over-union (mIoU). We include the selected POC models and IAA models trained with or without (w/o) KD in this comparison. Higher mIoU means better performances on extracting main subjects on our collected dataset.

can be ranked as: IAA ResNeXt101 with KD > selected POC models > IAA ResNeXt101 without KD. This further confirms our judgments concluded in the qualitative analysis.

D. Cost and Benefit Analysis

The extra cost of the proposed KD scheme is brought by: 1) extracting GSFs from multiple POC models; 2) training the knowledge distiller; 3) distilling teacher knowledge from GSFs; 4) computing losses between teacher aesthetic features and student aesthetic features when training the student model. To compare the scale of these extra computational costs to the computational cost for training a baseline IAA model, the costs are estimated and listed in Table XIII. Specifically, the scale of the training cost is measured as three times of the testing costs for the same model with the same input [52]. Since all the above-mentioned four parts that bring extra computational costs can be conducted separately, the overall extra computational costs are computed by linear combination. In Table XIII, to show the total extra cost brought by the proposed KD scheme, we present the costs for feature extraction with POC models (Setting 1, 2, 3), training the knowledge distiller

TABLE XIV: Comparison of training with GT (with mixed loss or mixed label) or without GT (with teacher prediction only) along with teacher predictions for supervision (evaluated on train-test split from [29]).

Setting	SRCC	PLCC	Acc
with teacher prediction	0.770	0.770	$\frac{82.1\%}{82.2\%}$
with mixed label	0.775	0.773	82.2% 82.0%

(Setting 4), and distill teacher knowledge with the knowledge distiller (Setting 5). Compared to the forward-propagation or backward-propagation via the models, the computational costs for loss computations are negligible, and therefore, costs for loss computations are not listed in Table XIII. Therefore, the total extra cost per input is the sum of the costs brought by feature extraction with POC models, training the knowledge distiller, and distilling teacher knowledge with the knowledge distiller. As to the cost for training the end-to-end IAA model with or without KD (Setting 6, 8), their computing the KD loss is the only difference that brings the extra cost for the setting with KD. For the cost for testing the end-to-end IAA models with or without KD (Setting 7, 9), the costs are the same since their settings are the same in the testing phase.

As it can be seen from Table XIII, the total extra cost is 2944.2 FLOPs(G) while the training cost for an end-to-end IAA model is 1101.2 FLOPs(G) (Setting 6, 8). This means that the total extra cost is ~ 1.7 times higher than training an end-to-end IAA model. Do note that the total training cost (Setting 6, 8) is a multiple of the number of epochs, while our setting adopts a rather fewer number of epochs in training (*i.e.*, 12 epochs), which makes the extra cost seemingly higher. While it is common to use a few tens of epochs in training an end-to-end IAA model. For example, Li et al.'s method [33] adopts 50 epochs in training. If our model is also trained for a few tens of epochs, the extra cost will be less than the training costs for an end-to-end model. On the contrary, among all three parts within the extra cost (Setting $1 \sim 5$), only the cost brought by training the knowledge distiller is a multiple of the number of epochs (Setting 4), and other parts of the extra cost will stay constant. Since the cost brought by training the knowledge distiller is on a scale of 1, increasing the number of epochs for training the knowledge distiller does not significantly increase the overall extra cost. With the extra cost, the model performance can be effectively improved by 4.8% (in Table VIII), and such an improvement is generalizable (in Fig. 6). Therefore, we believe that the proposed KD scheme is a rather cost-effective way for improving the performance of an end-to-end IAA model.

E. Discussion on Cause of Improvements brought by KD

A better IAA model requires more discriminative aesthetic features, and more discriminative aesthetic features require more diverse semantic patterns.

For the teacher model, using POC model with larger sizes, using POC model with more training data, or combining GSFs from different POC models can significantly improve the teacher models' performance (Table I and Table II). The main

Dataset	Task	Setting	SRCC	PLCC	Acc
CLULIV DO	Classification	Teacher			96.9%
CUHK-PQ		w/o KD with KD	-	-	96.2% 96.9%
AVA	Classification	Teacher w/o KD with KD			82.8% 80.0% 81.7%
AVA	Regression	Teacher w/o KD with KD	$ \begin{array}{c} 0.784 \\ \hline 0.\overline{6}87 \\ 0.707 \end{array} $	$ \begin{array}{r} 0.785 \\ \overline{0.689} \\ 0.693 \end{array} $	- <u>82.3%</u> 78.0% 78.9%

reason is that all above-mentioned approaches can provide more diverse semantic patterns that potentially provides more relevant patterns for constructing aesthetic features covering more contents. As a result, more discriminative aesthetic features further lead to performance gains.

For the student model, the performance improvement brought by KD can be explained by the semantic patterns learned from the teacher. As we experimentally show in Sec. IV-B and Sec. IV-C, both the teacher model and the student model trained merely with aesthetic labels construct aesthetic features from semantic patterns captured by their own pretrained backbones. Since the selected POC models for the teacher model can capture more semantic patterns than the student's backbone, the teacher aesthetic features are expected to cover more contents than the student aesthetic features. Thus, teacher aesthetic features can be used for extra supervision that guides the student model to capture more relevant semantic patterns for constructing aesthetic features covering more contents, which leads to higher IAA performance. Besides ablation studies on IAA in Table V, results on detecting the subject areas (Fig. 8 and Fig. 10) also imply that the IAA model learns to capture more IAA-relevant semantic patterns (related to the subject areas).

F. Variants of KD Loss

In this section, we further investigate some variants of the proposed KD loss along with the proposed method.

1) Combine GT with teacher predictions: We first investigate whether combining GT with teacher predictions will further improve the student model's performance. To this end, we have set up two settings to supervise the student model with the mixture of GT and teacher predictions:

• Mixed loss: GT and teacher predictions are separately used for loss computation as follows:

$$\mathcal{L}_{1}(\hat{\mathcal{D}}_{t}, \hat{\mathcal{D}}_{s}, \mathcal{D}_{GT}, f_{t}, f_{s}) = \frac{1}{2} EMD(\hat{\mathcal{D}}_{t}, \hat{\mathcal{D}}_{s}) + \frac{1}{2} EMD(\mathcal{D}_{GT}, \hat{\mathcal{D}}_{s}) + MSE(f_{t}, f_{s}),$$
(10)

where \hat{D}_t , f_t denote the teacher prediction and the teacher aesthetic feature, \hat{D}_s , f_s denote the student prediction and the student aesthetic feature, \mathcal{D}_{GT} denotes GT, $EMD(\cdot)$ and $MSE(\cdot)$ refer to EMD loss and mean squared error (MSE) loss respectively. • **Mixed loss**: GT and teacher predictions are linearly combined before loss computation:

$$\mathcal{L}_{2}(\hat{\mathcal{D}}_{t}, \hat{\mathcal{D}}_{s}, \mathcal{D}_{GT}, f_{t}, f_{s}) = EMD(\frac{1}{2}\mathcal{D}_{GT} + \frac{1}{2}\hat{\mathcal{D}}_{t}, \hat{\mathcal{D}}_{s}) + MSE(f_{t}, f_{s}).$$
(11)

The above-mentioned two settings are compared with the results of supervision with only teacher predictions. All models are based on ResNeXt101 (SWSL). The experimental results are given in Table XIV. As the results show, both settings with a mixture of GT and teacher predictions as supervision do not significantly outperform the setting merely with teacher predictions as supervision, considering a 0.003 performance variations as discussed in Sec. V-A. We interpret the results as follows. Compared to directly supervising the student model with GT, supervision with teacher predictions provide easier solutions to the student to map from aesthetic features to aesthetic predictions. Even trained with easier supervisions, there is still a performance gap between the teacher model and the student model. This means even we mix the GT with the teacher predictions for supervision, it is less likely that the student model will able to learn more information from GT.

2) Generalized to other IAA tasks: Since the distribution labels are not available on some IAA datasets, it would be better to verify the effectiveness of the proposed method when binary labels or MOSs are used as the GT. To this end, we adjust the KD loss to the cases when only binary labels or MOSs are available.

• For the case of binary classification, the first term can be replaced with a binary cross-entropy (BCE) loss:

$$\mathcal{L}_{\mathcal{KD}}(\hat{y}_t, \hat{y}_s, f_t, f_s) = BCE(\hat{y}_t, \hat{y}_s) + MSE(f_t, f_s),$$
(12)

where $\hat{y}_t \in [0, 1], f_t$ denote the teacher prediction and the teacher aesthetic feature, $\hat{y}_s \in [0, 1], f_s$ denote the student prediction and the student aesthetic feature, $BCE(\cdot)$ and $MSE(\cdot)$ refer to BCE loss and mean squared error (MSE) loss respectively.

• For the case of score regression, the first term can be replaced with a MSE loss:

$$\mathcal{L}_{\mathcal{KD}}(\hat{y}_t, \hat{y}_s, f_t, f_s) = MSE(\hat{y}_t, \hat{y}_s) + MSE(f_t, f_s),$$
(13)

where \hat{y}_t, f_t denote the teacher prediction and the teacher aesthetic feature, \hat{y}_s, f_s denote the student prediction and the student aesthetic feature, $MSE(\cdot)$ and $MSE(\cdot)$ refer to MSE loss and mean squared error (MSE) loss respectively.

We have conducted experiments on both CUHK-PQ dataset [14] (binary classification) and AVA dataset (binary classification and score regression) with above-mentioned losses. The results are given in Table XV. The results show that the proposed KD scheme is still effective in the two other cases of IAA.

VI. CONCLUSION

In this paper, we have focused on the problem of abstractness of aesthetic labels. On the one hand, during inference, the IAA model is required to relate various distinct semantic patterns to the same aesthetic label. On the other hand, when training, it would be hard for the IAA model to learn to distinguish different semantic patterns merely with the supervision from aesthetic labels. When the supervision was merely provided by aesthetic labels, experimental results (Table V) have implied that an IAA model mostly learned to construct aesthetic features from semantic patterns alreadyknown to its pre-trained backbone, instead of learning new semantic patterns for more discriminative aesthetic features. Therefore, an IAA model can be improved by providing semantic guidance in training, so that the IAA model can learn extra semantic patterns for more discriminative aesthetic features.

The proposed method is inspired by the observation that different POC models tend to capture different sets of semantic patterns (Fig. 8). And therefore, an IAA model that captures less diverse semantic patterns can learn from these POC models to capture more diverse semantic patterns for more discriminative aesthetic features covering more contents. However, since knowledge from POC models may not directly applicable to IAA (Fig. 2, Fig. 4, and Table IV), we have proposed to train a separate knowledge distiller to take relevant parts from the knowledge of POC models. Thus, a singlebackbone end-to-end model can be trained for IAA with semantic guidance from multiple POC models (Fig. 3). Since the selected POC models are deeper or trained with more data, extra semantic patterns are learned by the student for more discriminative aesthetic features (Table I, Table II).

Extensive experiments showed that: 1) the proposed KD scheme enables the student IAA model to capture more IAA-relevant semantic patterns for building more discriminative aesthetic features (Table V, VI, IX, Fig. 8, 10); 2) the proposed KD scheme also allows the student IAA model to adapt to smaller-sized inputs (Table VIII, VII); 3) the proposed KD scheme is generalizable across different train-test splits (Fig. 7, 6); 4) the proposed KD scheme is a rather cost-effective way for improving the performance of an IAA model (Table XIII); 5) the proposed IAA method significantly outperforms 10 previous relevant methods (Table X, XI); 6) the proposed approach can be adapted to binary classification and regression scenarios of IAA while keeping its effectiveness (Table XV).

REFERENCES

- S. Ma, Y. Fan, and C. W. Chen, "Pose maker: A pose recommendation system for person in the landscape photographing," in *Proceedings of* the 22nd ACM International Conference on Multimedia, 2014.
- [2] W. Wang, J. Shen, and H. Ling, "A deep network solution for attention and aesthetics aware photo cropping," *IEEE transactions on pattern analysis and machine intelligence*, vol. 41, no. 7, pp. 1531–1544, 2018.
- [3] P. Obrador, X. Anguera, R. de Oliveira, and N. Oliver, "The role of tags and image aesthetics in social image search," in *Proceedings of the first* SIGMM workshop on Social media, 2009, pp. 65–72.
- [4] C. Li, A. C. Loui, and T. Chen, "Towards aesthetics: A photo quality assessment and photo selection system," in *Proceedings of the 18th ACM international conference on Multimedia*, 2010, pp. 827–830.
- [5] X. Lu, Z. Lin, H. Jin, J. Yang, and J. Z. Wang, "Rapid: Rating pictorial aesthetics using deep learning," in *Proceedings of the 22nd* ACM international conference on Multimedia, 2014, pp. 457–466.
- [6] X. Lu, Z. Lin, X. Shen, R. Mech, and J. Z. Wang, "Deep multi-patch aggregation network for image style, aesthetics, and quality estimation," in *Proceedings of the IEEE International Conference on Computer Vision*, 2015, pp. 990–998.

- [7] V. Hosu, B. Goldlucke, and D. Saupe, "Effective aesthetics prediction with multi-level spatially pooled features," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 9375–9383.
- [8] H. Talebi and P. Milanfar, "Learning to resize images for computer vision tasks," arXiv preprint arXiv:2103.09950, 2021.
- [9] J. Hou, S. Yang, and W. Lin, "Object-level attention for aesthetic rating distribution prediction," in *Proceedings of the 28th ACM International Conference on Multimedia*, 2020.
- [10] X. Jin, L. Wu, G. Zhao, X. Li, X. Zhang, S. Ge, D. Zou, B. Zhou, and X. Zhou, "Aesthetic attributes assessment of images," in *Proceedings* of the 27th ACM International Conference on Multimedia, 2019, pp. 311–319.
- [11] Y. Kao, R. He, and K. Huang, "Deep aesthetic quality assessment with semantic information," *IEEE Transactions on Image Processing*, vol. 26, no. 3, pp. 1482–1495, 2017.
- [12] X. Zhang, X. Gao, W. Lu, L. He, and J. Li, "Beyond vision: A multimodal recurrent attention convolutional neural network for unified image aesthetic prediction tasks," *IEEE Transactions on Multimedia*, vol. 23, pp. 611–623, 2020.
- [13] X. Sun, H. Yao, R. Ji, and S. Liu, "Photo assessment based on computational visual attention model," in *Proceedings of the 17th ACM international conference on Multimedia*, 2009, pp. 541–544.
- [14] W. Luo, X. Wang, and X. Tang, "Content-based photo quality assessment," in 2011 International Conference on Computer Vision. IEEE, 2011, pp. 2206–2213.
- [15] L. Zhang, Y. Gao, R. Zimmermann, Q. Tian, and X. Li, "Fusion of multichannel local and global structural cues for photo aesthetics evaluation," *IEEE Transactions on Image Processing*, vol. 23, no. 3, pp. 1419–1429, 2014.
- [16] X. Tian, Z. Dong, K. Yang, and T. Mei, "Query-dependent aesthetic model with deep learning for photo quality assessment," *IEEE Transactions on Multimedia*, vol. 17, no. 11, pp. 2035–2048, 2015.
- [17] G. Hinton, O. Vinyals, and J. Dean, "Distilling the knowledge in a neural network," arXiv preprint arXiv:1503.02531, 2015.
- [18] Y. Zhang, T. Xiang, T. M. Hospedales, and H. Lu, "Deep mutual learning," in *Proceedings of the IEEE Conference on Computer Vision* and Pattern Recognition, 2018, pp. 4320–4328.
- [19] J. H. Cho and B. Hariharan, "On the efficacy of knowledge distillation," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 4794–4802.
- [20] A. Romero, N. Ballas, S. E. Kahou, A. Chassang, C. Gatta, and Y. Bengio, "Fitnets: Hints for thin deep nets," *arXiv preprint arXiv:1412.6550*, 2014.
- [21] F. Tung and G. Mori, "Similarity-preserving knowledge distillation," in Proceedings of the IEEE/CVF International Conference on Computer Vision, 2019, pp. 1365–1374.
- [22] J. Zhu, S. Tang, D. Chen, S. Yu, Y. Liu, M. Rong, A. Yang, and X. Wang, "Complementary relation contrastive distillation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 9260–9269.
- [23] G. Chen, W. Choi, X. Yu, T. Han, and M. Chandraker, "Learning efficient object detection models with knowledge distillation," *Advances in neural information processing systems*, vol. 30, 2017.
- [24] T. He, C. Shen, Z. Tian, D. Gong, C. Sun, and Y. Yan, "Knowledge adaptation for efficient semantic segmentation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 578–587.
- [25] Y. Hou, Z. Ma, C. Liu, T.-W. Hui, and C. C. Loy, "Inter-region affinity distillation for road marking segmentation," in *Proceedings of* the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2020, pp. 12486–12495.
- [26] W. Zhang, D. Li, C. Ma, G. Zhai, X. Yang, and K. Ma, "Continual learning for blind image quality assessment," arXiv preprint arXiv:2102.09717, 2021.
- [27] H. Talebi and P. Milanfar, "Nima: Neural image assessment," *IEEE Transactions on Image Processing*, vol. 27, no. 8, pp. 3998–4011, 2018.
- [28] P.-T. Jiang, C.-B. Zhang, Q. Hou, M.-M. Cheng, and Y. Wei, "Layercam: Exploring hierarchical class activation maps," *IEEE Transactions on Image Processing*, 2021.
- [29] N. Murray, L. Marchesotti, and F. Perronnin, "Ava: A large-scale database for aesthetic visual analysis," in 2012 IEEE Conference on Computer Vision and Pattern Recognition. IEEE, 2012, pp. 2408–2415.
- [30] L. Wang, X. Wang, T. Yamasaki, and K. Aizawa, "Aspect-ratiopreserving multi-patch image aesthetics score prediction," in *Proceed*ings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops, 2019, pp. 0–0.

- [31] L. Mai, H. Jin, and F. Liu, "Composition-preserving deep photo aesthetics assessment," in *Proceedings of the IEEE conference on computer* vision and pattern recognition, 2016, pp. 497–506.
- [32] X. Zhang, X. Gao, W. Lu, and L. He, "A gated peripheral-foveal convolutional neural network for unified image aesthetic prediction," *IEEE Transactions on Multimedia*, vol. 21, no. 11, pp. 2815–2826, 2019.
- [33] L. Li, H. Zhu, S. Zhao, G. Ding, and W. Lin, "Personality-assisted multitask learning for generic and personalized image aesthetics assessment," *IEEE Transactions on Image Processing*, vol. 29, pp. 3898–3910, 2020.
- [34] H. Zeng, Z. Cao, L. Zhang, and A. C. Bovik, "A unified probabilistic formulation of image aesthetic assessment," *IEEE Transactions on Image Processing*, vol. 29, pp. 1548–1561, 2019.
- [35] Q. Chen, W. Zhang, N. Zhou, P. Lei, Y. Xu, Y. Zheng, and J. Fan, "Adaptive fractional dilated convolution network for image aesthetics assessment," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 14114–14123.
- [36] D. She, Y.-K. Lai, G. Yi, and K. Xu, "Hierarchical layout-aware graph convolutional network for unified aesthetics assessment," in *Proceedings* of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2021, pp. 8475–8484.
- [37] N. Murray and A. Gordo, "A deep architecture for unified aesthetic prediction," arXiv preprint arXiv:1708.04890, 2017.
- [38] J. Yu, G. Xie, M. Li, H. Xie, and L. Yu, "Beauty product retrieval based on regional maximum activation of convolutions with generalized attention," in *Proceedings of the 27th ACM International Conference on Multimedia*, 2019, pp. 2553–2557.
- [39] J. Hou, S. Ji, and A. Wang, "Attention-driven unsupervised image retrieval for beauty products with visual and textual clues," in *Proceedings* of the 28th ACM International Conference on Multimedia, 2020.
- [40] I. Z. Yalniz, H. Jégou, K. Chen, M. Paluri, and D. Mahajan, "Billion-scale semi-supervised learning for image classification," *CoRR*, vol. abs/1905.00546, 2019. [Online]. Available: http://arxiv.org/abs/ 1905.00546
- [41] A. Kolesnikov, L. Beyer, X. Zhai, J. Puigcerver, J. Yung, S. Gelly, and N. Houlsby, "Big transfer (bit): General visual representation learning," 2020.
- [42] D. Mahajan, R. Girshick, V. Ramanathan, K. He, M. Paluri, Y. Li, A. Bharambe, and L. van der Maaten, "Exploring the limits of weakly supervised pretraining," 2018.
- [43] X. Jin, L. Wu, X. Li, X. Zhang, J. Chi, S. Peng, S. Ge, G. Zhao, and S. Li, "Ilgnet: inception modules with connected local and global features for efficient image aesthetic quality classification using domain adaptation," *IET Computer Vision*, vol. 13, no. 2, pp. 206–212, 2019.
- [44] Y. Zhou, Y. Sun, L. Li, K. Gu, and Y. Fang, "Omnidirectional image quality assessment by distortion discrimination assisted multi-stream network," *IEEE Transactions on Circuits and Systems for Video Technology*, 2021.
- [45] W. Zhang, K. Ma, J. Yan, D. Deng, and Z. Wang, "Blind image quality assessment using a deep bilinear convolutional neural network," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 30, no. 1, pp. 36–47, 2018.
- [46] Q. Jiang, F. Shao, W. Lin, and G. Jiang, "Blique-tmi: Blind quality evaluator for tone-mapped images based on local and global feature analyses," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 29, no. 2, pp. 323–335, 2017.
- [47] Y. Fang, R. Du, Y. Zuo, W. Wen, and L. Li, "Perceptual quality assessment for screen content images by spatial continuity," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 30, no. 11, pp. 4050–4063, 2019.
- [48] X. Zhang, W. Lin, and Q. Huang, "Fine-grained image quality assessment: A revisit and further thinking," *IEEE Transactions on Circuits* and Systems for Video Technology, 2021.
- [49] H. R. Sheikh, M. F. Sabir, and A. C. Bovik, "A statistical evaluation of recent full reference image quality assessment algorithms," *IEEE Transactions on image processing*, vol. 15, no. 11, pp. 3440–3451, 2006.
- [50] Y. Luo and X. Tang, "Photo and video quality evaluation: Focusing on the subject," in *European Conference on Computer Vision*. Springer, 2008, pp. 386–399.
- [51] X. Qin, Z. Zhang, C. Huang, C. Gao, M. Dehghan, and M. Jagersand, "Basnet: Boundary-aware salient object detection," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 7479–7489.
- [52] M. Hobbhahn and J. Sevilla. (2021) What's the backward-forward flop ratio for neural networks? [Online]. Available: https://www.lesswrong.com/posts/fnjKpBoWJXcSDwhZk/ what-s-the-backward-forward-flop-ratio-for-neural-networks