

# Find Beauty in the Rare: Contrastive Composition Feature Clustering for Nontrivial Cropping Box Regression

Zhiyu Pan<sup>1\*</sup>, Yinpeng Chen<sup>1\*</sup>, Jiale Zhang<sup>1</sup>, Hao Lu<sup>1</sup>, Zhiguo Cao<sup>1†</sup>, Weicai Zhong<sup>2</sup>

<sup>1</sup> Key Laboratory of Image Processing and Intelligent Control, Ministry of Education School of Artificial Intelligence and Automation, Huazhong University of Science and Technology

<sup>2</sup> Huawei CBG Consumer Cloud Service Search Product & Big Data Platform Department  
{zhiyupan, zgcao}@hust.edu.cn

## Abstract

Automatic image cropping algorithms aim to recompose images like human-being photographers by generating the cropping boxes with improved composition quality. Cropping box regression approaches learn the beauty of composition from annotated cropping boxes. However, the bias of annotations leads to quasi-trivial recomposing results, which has an obvious tendency to the average location of training samples. The crux of this predicament is that the task is naively treated as a box regression problem, where rare samples might be dominated by normal samples, and the composition patterns of rare samples are not well exploited. Observing that similar composition patterns tend to be shared by the cropping boundaries annotated nearly, we argue to find the beauty of composition from the rare samples by clustering the samples with similar cropping boundary annotations, *i.e.*, similar composition patterns. We propose a novel Contrastive Composition Clustering (C2C) to regularize the composition features by contrasting dynamically established similar and dissimilar pairs. In this way, common composition patterns of multiple images can be better summarized, which especially benefits the rare samples and endows our model with better generalizability to render nontrivial results. Extensive experimental results show the superiority of our model compared with prior arts.

## Introduction

Automatic image cropping is an effective technology to improve composition quality of the input images. As one of the key components of computational imaging, it has attracted wide research interest (Chen et al. 2003; Islam, Lai-Kuan, and Chee-Onn 2017). Provided human-annotated cropping boxes (Wang, Shen, and Ling 2018; Zeng et al. 2019), automatic cropping can be formalized as a cropping box regression problem. However, as regression models are easily affected by the bias of training set (Yang et al. 2021), recent regression-based approaches (Hong et al. 2021; Pan et al. 2021b) are dominated by the normal samples and tend to render quasi-trivial cropping boxes (Fig. 1(a)). The quasi-trivial cropping boxes can be described as the central cropping boxes whose boundaries locate near the average location of training samples. When the quasi-trivial cropping

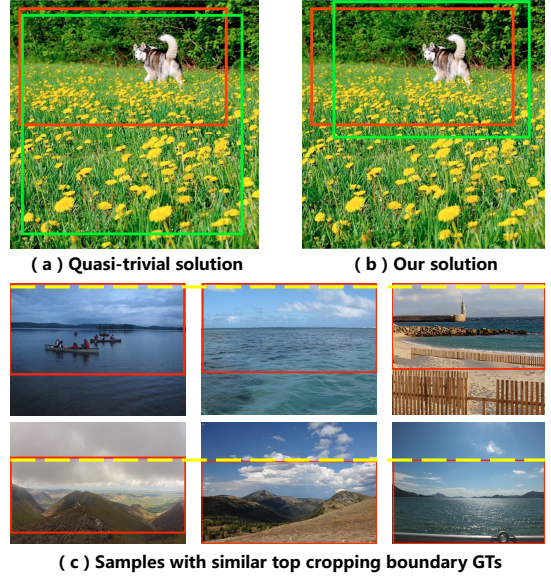


Figure 1: An illustration of quasi-trivial solution and the examples of our observation. (a) The cropping result from CACNet (Hong et al. 2021). (b) The cropping result from our approach. (c) Samples whose top cropping boundaries locate relatively nearly at the original image coordinate. Red and green boxes indicate the annotated and predicted cropping results respectively; Yellow dotted line indicates an identical relative reference position.

boxes just inherit the location bias of the annotated cropping boundaries and ignore the composition patterns of the input images, the improvement of composition quality is little, which limits the real-world application. Trying to generate nontrivial cropping boxes, previous methods take extra data with additional annotations, *e.g.*, composition classification dataset (Lee et al. 2018), into account. However, the predicament is hardly addressed. Therefore, we try to reveal the intrinsic cause of quasi-trivial solution. As shown in Fig. 2, by analysing the CACNet baseline (Hong et al. 2021), we illustrate the cause of quasi-trivial solution: *dominated by the normal samples, the composition patterns of rare samples are barely exploited*. Hence, we argue to find the beauty of composition by taking full advantage of rare samples rather

\*These authors contributed equally.

†Corresponding author.

than by introducing extra data.

In this paper, we probe into the problem: *how to better exploit the composition patterns of the rare training samples?* Similar problems have been studied for decades by the imbalanced learning methods (Branco, Torgo, and Ribeiro 2017; Lin et al. 2017; Yang et al. 2021), whose main idea is re-weighting. However, directly adapting these methods into the cropping box regression task can barely improve the composition quality. Therefore, we delve into this problem from another view, which is to enhance the composition representations of the rare by looking for the connections between samples. As illustrated in Fig. 1(c), compared in the same row, when the annotated top cropping boundaries tend to locate at a close relative location, the visual focuses locate relatively nearby; compared with two different rows, the focuses locate relatively differently. Based on this observation, we make a hypothesis: *cropping boundaries with close relative location annotations share the similar composition patterns and vice versa*. Based on this, we make full use of the relationships between the samples, which especially benefits the rare ones.

Inspired by the contrastive learning (Hadsell, Chopra, and LeCun 2006; Chen et al. 2020; He et al. 2020) and deep clustering works (Zhou et al. 2022; Wu et al. 2022), we propose Contrastive Composition Clustering (C2C) to regularize the composition features based on their regression target distance. Due to that our hypothesis is based on the cropping boundary, each image has four composition features to depict each cropping boundary respectively. The optimization targets of C2C are: (1) draw features of the samples with similar annotations closer (2) widen the feature distance of negative pairs where annotated boundary locates far away from each other. When C2C is performed in the mini-batch, we calculate the distance map to select positive and negative pairs dynamically. Based on the distance map, after breaking the contrastive loss into alignment and uniformity loss (Wang and Isola 2020), the contrastive clustering procedure can successfully work in our regression scenario. As illustrated in Fig. 1(a) and (b), benefited from the proposed C2C, our solution can render nontrivial cropping boxes.

Extensive experiments on public benchmarks show that our solution outperforms state-of-the-art regression-based approaches even without extra data. By visualizing the cropping results, with C2C, our model hardly collapses into quasi-trivial solutions. Compared with other imbalanced learning baselines, C2C also shows superiority in both performance and time efficiency. For the first time, without external data, we demonstrate state-of-the-art image cropping with only a deep clustering procedure.

## Related Work

### Image Cropping

Conventional image cropping algorithms (Chen et al. 2003; Suh et al. 2003; Marchesotti, Cifarelli, and Csurka 2009; Cheng et al. 2010; Greco and La Cascia 2013) formulate the composition patterns with predefined rules (Zhang et al. 2005; Nishiyama et al. 2009; Yan et al. 2013; Fang et al. 2014). However, performance of the classic algorithms are

far from satisfactory. Recently, data-driven algorithms have significantly improved. They are driven by two main ideas: candidate box selection and cropping box regression.

**Candidate box selection algorithms** follow a two-stage pipeline. First, candidate boxes are generated according to prior knowledge. Then, the candidate boxes are ranked based on the learned composition aesthetic knowledge. The knowledge can be modeled by saliency detection (Wang, Shen, and Ling 2018; Tu et al. 2020), teacher-student architecture (Wei et al. 2018), region of interest and discard (Zeng et al. 2019, 2020), view-wise mutual relation (Li et al. 2020), visual elements dependencies (Pan et al. 2021a; Wang et al. 2022) and view-wise difference (Pan et al. 2022). But the results are dependent on the candidates. To overcome this limitation, some works attempt to find end-to-end solutions.

**Cropping box regression algorithms** are end-to-end solutions that imitate the process of human cropping. Reinforcement learning algorithms estimate the final cropping box step by step (Li et al. 2018, 2019). Other algorithms predict the cropping results according to a salient cluster (Pan et al. 2021b), or a composition class (Hong et al. 2021). The common deficiencies of these algorithms are the need for extra data and the problem of degenerating into quasi-trivial solutions. In contrast to these approaches, our work can render nontrivial cropping boxes without extra data.

## Contrastive Learning and Deep Clustering

Motivated by the InfoMax principle (Linsker 1988), contrastive learning aims to learn representations by contrasting positive pairs against negative pairs (Hadsell, Chopra, and LeCun 2006; Dosovitskiy et al. 2014). The pioneer works introduce the memory bank to store the representation vectors and update them by contrasting (Wu et al. 2018; Zhuang, Zhai, and Yamins 2019; He et al. 2020). Some other works contrast in batches (Doersch and Zisserman 2017; Ji, Henriques, and Vedaldi 2019; Ye et al. 2019; Chen et al. 2020). Recently, Wang *et al.* (Wang and Isola 2020) has characterized the contrastive learning by alignment and uniformity. When deep clustering (Zhou et al. 2022) concerns the joint optimization of clustering and deep representation learning, some works (Ling et al. 2022; Deng et al. 2022; Chen et al. 2022) introduce the idea of contrasting into deep clustering. In this work, we also conduct composition feature clustering based on the idea of contrasting and combine the clustering procedure into regular training process, which helps to get a better composition representation.

## Imbalanced Learning

In confront of the training data bias, a naive idea is to re-weight the loss with the frequency of samples (Lin et al. 2017). Some other works also try to create new samples to balance the original distribution (Torgo et al. 2013; Branco, Torgo, and Ribeiro 2017) or to decouple the encoder and decoder (Kang et al. 2019). Recently, Yang *et al.* (Yang et al. 2021) proposes the label and feature distribution smoothing to train the regressor unbiasedly. We directly adapt these methods into our problem as baselines and compare our C2C with them in the experiment section.

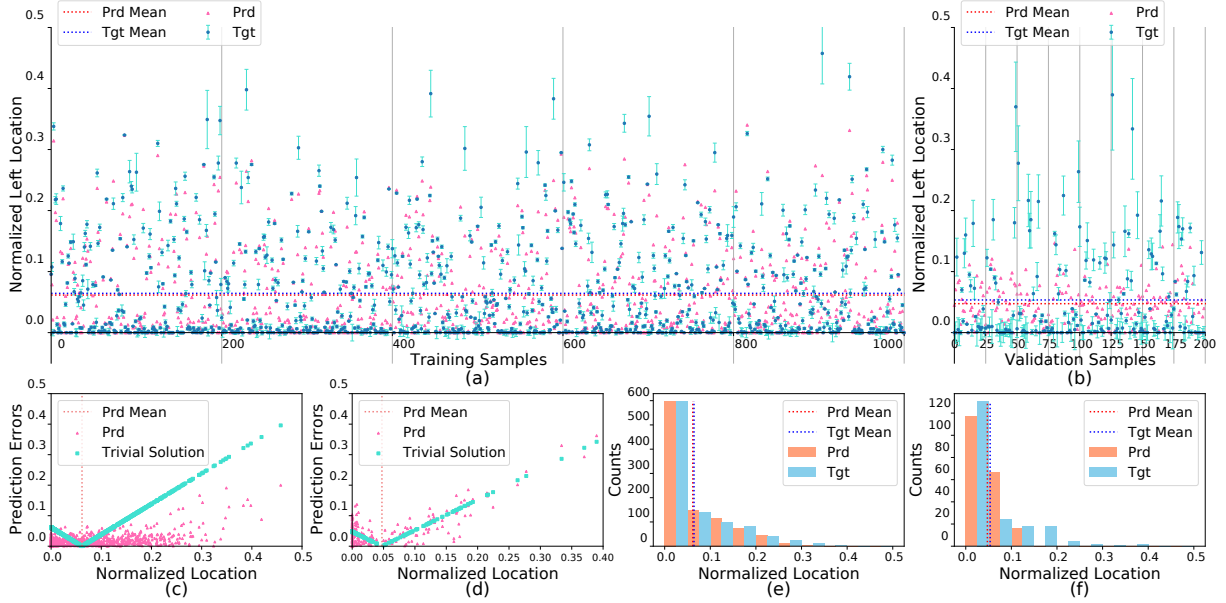


Figure 2: Statistical analysis of the left cropping boundaries generated from CACNet vanilla regression baseline. (a) The test results on the training set. (b) The test results on the validation set. (c) The error distribution on the training set. (d) The error distribution on the validation set. (e) The location distribution on the training set. (f) The location distribution on the validation set. “prd” means the prediction, “tgt” denotes the target. Best viewed in color.

## Problem and Approach

### Quasi-Trivial Solution

We explain what is the quasi-trivial cropping solution and illustrate the cause of it by analyzing the recomposed results of the vanilla regression baseline CACNet (Hong et al. 2021). To make the analysis unaffected by extra data, the composition classification module is dropped. Statistic of the left boundary is shown in Fig. 2 (results of other boundaries are in similar distributions). The red and blue points in Fig. 2(a-b) show the predicted and target left boundary location respectively. The location is normalized into  $[0, 1]$ , where 0 represents the leftmost location and 1 represents the rightmost location. The ordinal number of samples are indicated by the horizontal coordinate. The error bar on the target point colored in light blue indicates the absolute value of the prediction error on this sample. From Fig. 2(b), on the validation set, we observe that the samples whose target is far away from the mean location tend to have larger prediction error, which means that the baseline tends to locate the boundary at an approximate average location for an unseen image. Therefore, we compare the baseline with the mean location which can serve as a trivial solution. As illustrated in Fig. 2(c-d), when the vertical and horizontal coordinates of each point represent the prediction error and the corresponding target location respectively, the trivial solution behaves as two symmetrical lines with the mean location as the axis of symmetry. From Fig. 2(d), on the validation set, we find that the baseline acts extremely similar to the trivial solution. We call this the quasi-trivial solution of this boundary. By plotting the location distribution of the targets and predicted results, we find that normal training samples with the most amount locate near the mean location

(Fig. 2(e)), hence the baseline tends to predict around the mean location for an unseen image (Fig. 2(f)), which means the quasi-trivial solution comes from that the normal samples overwhelm the rare samples. Hence, we try to exploit the composition patterns of the rare samples better.

### Cropping Boundary Locating Network

As the imbalanced distributions of four cropping boundaries have different levels and forms, the whole cropping box regression problem is divided into four boundary locating sub-tasks. Inspired by SABL (Wang et al. 2020), we build the cropping boundary locating network (CBLNet). As shown in Fig. 3, CBLNet contains three main components: a convolutional backbone, a boundary feature encoder which builds the boundary features, and a group of regression heads that predict the location of four cropping boundaries.

**Backbone.** A MobileNetV2-based (Sandler et al. 2018) backbone that fuses multi-scale features (Zeng et al. 2019, 2020) is employed. It generates multi-scale feature maps  $\mathcal{F} \in \mathbb{R}^{C \times H \times W}$  from the original image, where  $C$  is the channel dimensionality,  $H$  and  $W$  represents the height and width of the feature maps.

**Boundary feature encoder.** The extracted features  $\mathcal{F}$  are processed by two multi-head attention modules to encode the horizontal and vertical information. The multi-head attention module is implemented by  $m$   $1 \times 1$  convolutions, where  $m$  is the number of heads. The horizontal feature map  $\mathcal{H} \in \mathbb{R}^{m \times C \times H \times W}$  and vertical feature map  $\mathcal{V} \in \mathbb{R}^{m \times C \times H \times W}$  can be computed by

$$\mathcal{H} = \sigma_y(\xi(\tau_0^h(\mathcal{F}) * \mathcal{F}, \dots, \tau_{m-1}^h(\mathcal{F}) * \mathcal{F})), \quad (1)$$

$$\mathcal{V} = \sigma_x(\xi(\tau_0^v(\mathcal{F}) * \mathcal{F}, \dots, \tau_{m-1}^v(\mathcal{F}) * \mathcal{F})), \quad (2)$$

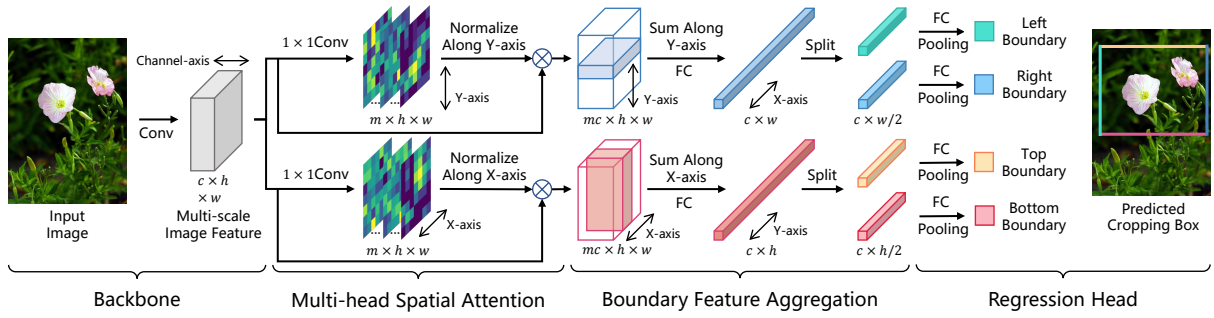


Figure 3: The pipeline of the cropping boundary locating network. The multi-scale features extracted from the input image are used to encode the horizontal and vertical features. Then, the horizontal and vertical features are split into boundary features. The final cropping boxes are predicted based on the four boundary features.

where  $\sigma_y(\cdot)$  and  $\sigma_x(\cdot)$  are the normalizing operators along Y and X axis, respectively,  $\xi$  is the concatenating operator along the channel axis,  $\tau_i^h(\cdot)$  and  $\tau_i^v(\cdot)$  represent the  $i$ -th head  $1 \times 1$  convolution of the horizontal and vertical attention modules, respectively. After the dimension reduction and the sum operations along the Y and X axis,  $\mathcal{H}$  and  $\mathcal{V}$  are transformed into the horizontal vector  $\mathbf{h} \in \mathbb{R}^{C \times W}$  and vertical vector  $\mathbf{v} \in \mathbb{R}^{C \times H}$ , respectively. Then,  $\mathbf{h}$  and  $\mathbf{v}$  are split into boundary features  $\mathbf{x}^l, \mathbf{x}^r \in \mathbb{R}^{C \times W/2}$  and  $\mathbf{x}^t, \mathbf{x}^b \in \mathbb{R}^{C \times H/2}$ , which are composition features for left, right, top, and bottom cropping boundaries, respectively.

**Regression head.** The normalized locations w.r.t. the input image of four cropping boundaries are predicted by a pooling operator and a 3-layer fully-connected network with the sigmoid activation function. All regression heads are supervised with the  $\ell_1$  loss.

### Contrastive Composition Clustering

After dividing the whole cropping box regression into four sub-regression, composition features w.r.t. four cropping boundaries are obtained. Based on our hypothesis, when similar cropping boundaries have common knowledge and dissimilar ones tend to have different composition patterns, it is intuitive to minimize the distance of similar samples and to maximize the distance of dissimilar samples in the composition feature space. As illustrated in Fig. 4, inspired by the contrastive learning algorithms (Wu et al. 2018; Zhuang, Zhai, and Yamins 2019; He et al. 2020), we propose the Contrastive Composition Clustering (C2C).

The C2C introduces the idea of contrastive learning into our regression scenario by setting the positive and negative pair thresholds, *i.e.*,  $z_p$  and  $z_n$ . The positive and negative pair refers to the composition feature pair whose boundary annotations locate nearly and far away from each other, respectively. In this setting, for each batch, the positive and negative pairs are not fixed. Hence, we structure the pairs with a distance map, which is used to calculate the positive and negative pair matrix  $\mathcal{P}$  and  $\mathcal{N}$ . Inspired by (Wang and Isola 2020), which identifies two contrastive loss properties: alignment and uniformity, we adapt the alignment and uniformity loss into our C2C. The alignment loss is to cluster positive pairs, and the uniformity loss is to discriminate between the elements in negative pairs. In our approach, the

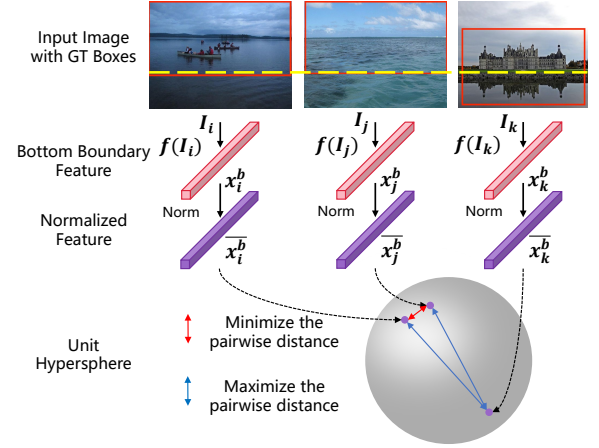


Figure 4: Contrastive composition clustering. C2C regularizes the features of the cropping boundaries. Taking the learning process of bottom boundary features as an example: the boundary features are  $\ell_2$ -normalized to the unit hypersphere. The pairwise distances on the unit hypersphere are maximized or minimized conditioned on the pairwise distances of their ground truth boundaries. After training,  $\ell_2$  normalization is removed.

regression alignment loss and uniformity loss take the form

$$L_{align} = \sum_{d=l,r,t,b} \mathbb{E}_{0 \leq i < j \leq N} \mathcal{P}_{i,j}^d \max(\|\bar{\mathbf{x}}_i^d - \bar{\mathbf{x}}_j^d\|_2^\alpha - \epsilon, 0) \quad (3)$$

$$L_{uniform} = \sum_{d=l,r,t,b} \mathbb{E}_{0 \leq i < j \leq N} \mathcal{N}_{i,j}^d e^{-t\|\bar{\mathbf{x}}_i^d - \bar{\mathbf{x}}_j^d\|_2^2}, \quad (4)$$

where  $\bar{\mathbf{x}}$  is the  $\ell_2$  normalized boundary composition features,  $l, r, t, b$  indicates left, right, top, and bottom respectively,  $N$  denotes the batch size,  $\epsilon$  controls the expected similarity level of the features with similar regression targets,  $\mathcal{P}_{i,j}^d$  and  $\mathcal{N}_{i,j}^d$  can be calculated as

$$\mathcal{P}_{i,j}^d = \mathbb{1}(y_i^d - y_j^d < z_p) \quad (5)$$

$$\mathcal{N}_{i,j}^d = \mathbb{1}(y_i^d - y_j^d > z_n), \quad (6)$$






Figure 5: Qualitative comparison with image cropping models. Redundancy in blue boxes violates the composition.

where  $y$  is the normalized ground truth location of the cropping boundary, and  $\mathbb{1}(\cdot)$  is the sign function.

Details about the derivation process of the optimization target and the pseudo-code of C2C is described in the supplementary. The alignment and uniformity loss are jointly optimized with the  $\ell_1$  loss for the cropping boundary locating. The total loss of the cropping boundary regression pipeline can be computed by

$$L = L_1 + \beta L_{align} + \gamma L_{uniform}, \quad (7)$$

where  $\beta$  and  $\gamma$  are the hyper-parameters used to control the relative importance of the alignment and uniformity loss. With C2C, we get more informative and discriminative composition features from the rare samples which prevents the model collapse into the quasi-trivial solution.

## Experiments

Here we demonstrate that, CBLNet with C2C achieves better performance than other cropping box regression algorithms with extra data. When imbalanced learning approaches can serve as naive baselines to exploit the rare samples, we compare C2C with these baselines. The results illustrate the superiority of C2C. A detailed ablation study is also conducted to illustrate the effect of each component. Further analyses and visualizations also prove the rationale of our hypothesis and design.

### Implementation Details

The input images are resized to  $224 \times 224$ . After processed by the backbone pre-trained on ImageNet, the channel dimensionality of multi-scale feature map  $\mathcal{F}$  of size  $H = W = 14$  is reduced to  $C = 256$ . The head number of the multi-head spatial attention module is set to  $m = 6$ . In the training stage, 32 images are batched as the input. In the alignment loss,  $z_l$  is set to 0.05,  $\alpha$  is set to 1, and  $\epsilon$  is set to 0.5. In the uniformity loss,  $z_h$  is set to 0.7 and  $t$  is set to 1. Only random cropping is used in data augmentation. By setting  $\beta = \gamma = 0.025$ , the network is optimized by Adam with the learning rate of  $5 \times 10^{-4}$  for 80 epochs.

## Datasets and Evaluation Metrics

FCDB (Chen et al. 2017) dataset is used for evaluation. The FCDB contains 1743 images, and each of them has a ground truth cropping box. In the performance comparison, 1395 images of the FCDB are used for training and 348 images for testing. We follow the setting of (Hong et al. 2021) to randomly choose 200 images as the validation set and the other images are used for training. To demonstrate whether an approach can render nontrivial cropping boxes, according to their frequency, the test set or the validation set is divided into three disjoint subsets: many-shot, medium-shot, and few-shot samples. The frequency of samples is related to their ground truth box size ratios w.r.t. the original image. Hence, in our setting, the ground truth box size ratios of the many-shot, medium-shot, and few-shot samples are in the range of 65% ~ 100%, 40% ~ 65%, and 0% ~ 40%, respectively. The performance of intersection over union (IoU) and the boundary displacement error (BDE) are reported.

## Performance Comparison

**Comparison with image cropping models.** Quantitative results on the FCDB dataset are illustrated in Table 1. For the selection-based algorithms, the recalled top-1 cropping boxes are used to calculate the IoU and BDE metrics. Based on the results, we can make the following observations: (a) *The gap between two types of algorithms lies in the few-shot samples.* This might be a manifestation of the quasi-trivial solution problem. The CBLNet trained with the C2C outperforms the previous regression algorithms especially on the few-shot samples without extra data, which bridges the performance gap between regression-based and selection-based algorithms. (b) *C2C is more helpful than extra data.* Compared to the vanilla CACNet without the composition classification to the final CACNet, it is obvious that the performance boost is from the extra data with the annotation about composition classification. When the performance of CBLNet without C2C is comparable with vanilla CACNet, after trained with C2C, the CBLNet can perform better than the CACNet, which shows the superiority of the proposed C2C. Qualitative comparison is shown in Fig. 5. We observe that the selection-based algorithms, *i.e.*, the VPN, can gen-

Algorithm	Metric	IoU $\uparrow$			BDE $\downarrow$			
	All	Many	Med.	Few	All	Many	Med.	Few
candidate box selection algorithms								
VPN $^\dagger$	0.665	<b>0.735</b>	0.631	<b>0.435</b>	0.085	<b>0.068</b>	0.092	<b>0.147</b>
GAIC $^\dagger$	0.666	0.721	<b>0.655</b>	0.408	0.084	0.072	<b>0.086</b>	<b>0.147</b>
TransView	0.682	-	-	-	0.080	-	-	-
VEN	<b>0.735</b>	-	-	-	<b>0.072</b>	-	-	-
cropping box regression algorithms								
A2RL $^\ddagger$	0.636	0.734	0.577	0.349	0.097	0.070	0.111	0.186
A3RL	0.696	-	-	-	0.077	-	-	-
vanilla-CACNet $^\ddagger$	0.700	0.799	0.650	0.369	0.075	0.050	0.086	0.167
CACNet $^\ddagger$	0.716	<b>0.809</b>	0.673	0.386	0.070	<b>0.047</b>	0.079	0.159
regression baseline	0.692	0.776	0.655	0.389	0.077	0.056	0.085	0.158
CBLNet - Ours	0.700	0.787	0.659	0.391	0.075	0.052	0.083	0.158
CBLNet+C2C - Ours	<b>0.718</b>	0.805	<b>0.680</b>	<b>0.418</b>	<b>0.069</b>	<b>0.047</b>	<b>0.078</b>	<b>0.146</b>

Table 1: Quantitative comparison with other state-of-the-art image cropping models on the FCDB dataset. Best performance is in boldface. The performances tagged by  $\dagger$  are our reproduced results, and  $\ddagger$  denotes the results from the original authors.

Algorithm	Metric	IoU $\uparrow$				BDE $\downarrow$				Training Time $\downarrow$
		All	Many	Med.	Few	All	Many	Med.	Few	
CBLNet baseline		0.700	0.787	0.659	0.391	0.075	0.052	0.083	0.158	0.720
SMOEN		0.702	0.796	0.658	0.374	0.075	0.051	0.085	0.166	0.956
RRT		0.701	0.791	0.653	0.416	0.074	0.051	0.086	0.149	0.730 $\times$ 2
FOCAL-R		0.708	0.797	0.663	0.408	0.072	0.050	0.082	0.151	0.747
INV		0.707	0.798	0.660	0.404	0.073	0.050	0.083	0.153	<b>0.734</b>
LDS&FDS		0.709	0.798	0.666	0.408	0.072	0.049	0.082	0.151	1.187
C2C		<b>0.718</b>	<b>0.805</b>	<b>0.680</b>	<b>0.418</b>	<b>0.069</b>	<b>0.047</b>	<b>0.078</b>	<b>0.146</b>	0.788

Table 2: Quantitative comparison of imbalanced learning algorithms on the FCDB dataset. Best performance is in boldface. The metric of training time shows how many seconds an algorithm needs for training a batch of data. Note that the regressor re-training (RRT) algorithm needs an extra training round, hence the training time is doubled.

erate nontrivial cropping results but can not obey the symmetry (first row) or visual balance (third row) composition rules well. This shows the limitation of selection-based algorithms. While previous regression-based algorithms tend to render quasi-trivial cropping results, our approach can render nontrivial cropping results and align better with composition rules.

**Comparison with imbalanced learning baselines.** The imbalanced learning baselines to be compared:

- *SMOEN* (Branco, Torgo, and Ribeiro 2017) defines the rare samples and creates synthetic samples by interpolating the feature and the regression target. During the training stage, annotated boundary may shift in the range of 10% of the original length or width.
- *FOCAL-R* is the regression version of the focal loss (Lin et al. 2017). Following (Yang et al. 2021), we rewrite the focal loss as  $1/n \sum_{i=1}^n \rho(|\mu e_i|)^\psi e_i$ , where  $e_i$  is the  $\ell_1$  error for  $i$ -th sample,  $\rho(\cdot)$  is the sigmoid function, and  $\mu$  and  $\psi$  are set to  $\mu = \psi = 2$ .
- *Inverse reweighting (INV)* is a traditional method reweighting the loss according to inverse frequency of the classes. (Yang et al. 2021) adopts the scheme based on the

regression target distribution in the regression scenario.

- *Regressor re-training (RRT)* (Kang et al. 2019) decouples the training process of feature and classifier. (Yang et al. 2021) transforms it into the regression version, re-training the regression head with inverse reweighting.
- *LDS&FDS* (Yang et al. 2021) propose to smooth the distribution of regression target to reweight the loss and recalibrate the statistics of the feature distribution.

The results are illustrated in Table 2. It shows that previous methods have notable limitations: (a) RRT can boost the performance on the few-shot samples but harms that on the medium-shot samples. (b) LDS&FDS can improve the performance on many-, medium-, few-shot samples, but at a cost of additional time consumption. C2C obtains significant gains on all groups of samples with low time cost.

## Ablation Study

**Boundary feature encoder and C2C.** The performance contributions of the boundary feature and C2C are shown in Table 1. We can observe that the regression baseline tends to output the quasi-trivial cropping results. By only adding the boundary feature encoder to build the CBLNet, the performance cannot be significantly improved. After training with

Batch Size	IoU $\uparrow$				BDE $\downarrow$			
	All	Many	Med.	Few	All	Many	Med.	Few
4	0.734	0.811	0.685	0.424	0.064	<b>0.041</b>	0.076	0.161
8	0.738	0.805	0.698	0.436	0.061	0.043	0.071	0.145
16	0.741	0.801	0.710	0.438	0.060	0.043	0.066	0.144
32	<b>0.753</b>	0.809	0.715	0.460	<b>0.059</b>	0.042	<b>0.065</b>	0.139
64	0.751	0.809	0.714	0.463	<b>0.059</b>	0.042	0.066	0.142
128	0.749	<b>0.813</b>	0.704	0.465	0.062	0.042	0.073	0.142
256	0.750	0.799	<b>0.718</b>	<b>0.471</b>	0.061	0.046	0.068	<b>0.128</b>

Table 3: Ablation study on the batch size. Best performance is in boldface.




Figure 6: Distributions of the boundary features from the CBLNet with/without C2C on the validation set. The color of the points represents the location of their ground truth boundary, the lighter the color is, the closer the ground truth boundary locates to the original image boundary. Best viewed in color.

C2C, the cropping results align better with the ground truth. This shows that the main contribution of this work is from the C2C rather than the CBLNet design.

**Batch size.** The proposed C2C contrasts the samples in the batch without the memory bank. We vary the batch size from 4 to 256. Larger batch size provides more positive and negative pairs. From the results in Table 3, it can be seen that the change of the performance on the many-shot samples is relatively smooth. As for the few-shot samples, with larger batch size, the performance is better. The results reveal that the proposed C2C benefits from large batch sizes (Chen et al. 2020; He et al. 2020). However, for a better overall performance, we adopt batch size of 32.

## Analysis

As illustrated in Fig. 6, we visualize the distribution of the composition features of left and right boundaries (the results of top and bottom boundaries are in similar distributions)

from the CBLNet trained with and without C2C by reducing the features into two dimensions by t-SNE (Van der Maaten and Hinton 2008). The location of the corresponding ground truth boundaries is denoted by color. The darker color means that the annotated boundary locates farther from the original input boundary. Without C2C, the features of the rare samples, whose ground truth is far away from the input image boundary, are mixed up with that of the normal samples. This means that the model without C2C can not properly depict the unseen rare samples, and seems to just project these samples randomly in the feature space. Hence, these samples follow the normal distribution of training set, which is the cause of quasi-trivial solutions. With C2C, firstly, the rare samples with darker color clustered and the feature distribution aligns well with the location distribution of the regression targets, which shows the success of C2C. Secondly, the expected feature distribution leads to a better performance, which demonstrates the rationale of our observation and hypothesis that cropping boundaries locating nearby share similar composition patterns.

## Conclusions

In this work, we study why existing cropping box regression algorithms tend to render quasi-trivial solution. We find that the crux lies in the limitation of representational capacity. Based on our observation, we make a hypothesis that cropping boundaries with similar annotations share similar composition patterns. Therefore, we propose to conduct composition feature clustering to enhance the composition representation. Firstly, by presenting CBLNet, the composition features of four cropping boundaries are obtained. Inspired by the idea of contrastive learning and deep clustering, we propose C2C to cluster the samples with similar annotations jointly with the normal optimization procedure. By training with C2C, our CBLNet can outperform other state-of-the-art regression-based image cropping models on the FCDB dataset even without extra data. Compared with other unbiased training baselines, the C2C also achieves superior performance with the little time cost. Our work reveals that cropping box regression networks actually can directly learn from the cropping box annotations, but how to train a network matters. Our work provides a new perspective on the task of cropping box regression and states that composition feature clustering can make a difference.

## Acknowledgements

This work was funded by the DigiX Joint Innovation Center of Huawei-HUST.

## References

- Branco, P.; Torgo, L.; and Ribeiro, R. P. 2017. SMOGN: a pre-processing approach for imbalanced regression. In *1st Intell. Workshop Learn. Imbalanced Domains: Theory Appl.*, 36–50.
- Chen, L.-Q.; Xie, X.; Fan, X.; Ma, W.-Y.; Zhang, H.-J.; and Zhou, H.-Q. 2003. A visual attention model for adapting images on small displays. *Multimedia Syst.*, 9(4): 353–364.
- Chen, T.; Kornblith, S.; Norouzi, M.; and Hinton, G. 2020. A simple framework for contrastive learning of visual representations. In *Int. Conf. Mach. Learn.*, 1597–1607.
- Chen, Y.; Pan, Z.; Shi, M.; Lu, H.; Cao, Z.; and Zhong, W. 2022. Design What You Desire: Icon Generation from Orthogonal Application and Theme Labels. In *Proc. ACM Int. Conf. Multimedia*, 2536–2546.
- Chen, Y.-L.; Huang, T.-W.; Chang, K.-H.; Tsai, Y.-C.; Chen, H.-T.; and Chen, B.-Y. 2017. Quantitative analysis of automatic image cropping algorithms: A dataset and comparative study. In *Proc. IEEE Winter Conf. Appl. Comput. Vis.*, 226–234.
- Cheng, B.; Ni, B.; Yan, S.; and Tian, Q. 2010. Learning to photograph. In *Proc. ACM Int. Conf. Multimedia*, 291–300.
- Deng, X.; Huang, D.; Chen, D.-H.; Wang, C.-D.; and Lai, J.-H. 2022. Strongly Augmented Contrastive Clustering. *arXiv preprint arXiv:2206.00380*.
- Doersch, C.; and Zisserman, A. 2017. Multi-task self-supervised visual learning. In *Proc. Int. Conf. Comput. Vis.*, 2051–2060.
- Dosovitskiy, A.; Springenberg, J. T.; Riedmiller, M.; and Brox, T. 2014. Discriminative unsupervised feature learning with convolutional neural networks. *Adv. Neural Inf. Process. Syst.*, 27.
- Fang, C.; Lin, Z.; Mech, R.; and Shen, X. 2014. Automatic image cropping using visual composition, boundary simplicity and content preservation models. In *Proc. Int. Conf. Multimedia*, 1105–1108.
- Greco, L.; and La Cascia, M. 2013. Saliency based aesthetic cut of digital images. In *Proc. Int. Conf. Image Anal. Process.*, 151–160.
- Hadsell, R.; Chopra, S.; and LeCun, Y. 2006. Dimensionality reduction by learning an invariant mapping. In *Proc. Conf. Comput. Vis. Pattern Recog.*, volume 2, 1735–1742.
- He, K.; Fan, H.; Wu, Y.; Xie, S.; and Girshick, R. 2020. Momentum contrast for unsupervised visual representation learning. In *Proc. Conf. Comput. Vis. Pattern Recog.*, 9729–9738.
- Hong, C.; Du, S.; Xian, K.; Lu, H.; Cao, Z.; and Zhong, W. 2021. Composing Photos Like a Photographer. In *Proc. Conf. Comput. Vis. Pattern Recog.*, 7057–7066.
- Islam, M. B.; Lai-Kuan, W.; and Chee-Onn, W. 2017. A survey of aesthetics-driven image recomposition. *Multimedia Tools Appl.*, 76(7): 9517–9542.
- Ji, X.; Henriques, J. F.; and Vedaldi, A. 2019. Invariant information clustering for unsupervised image classification and segmentation. In *Proc. Int. Conf. Comput. Vis.*, 9865–9874.
- Kang, B.; Xie, S.; Rohrbach, M.; Yan, Z.; Gordo, A.; Feng, J.; and Kalantidis, Y. 2019. Decoupling representation and classifier for long-tailed recognition. *Int. Conf. Learn. Representations*.
- Lee, J.-T.; Kim, H.-U.; Lee, C.; and Kim, C.-S. 2018. Photographic composition classification and dominant geometric element detection for outdoor scenes. *J. Vis. Commun. Image Represent.*, 55: 91–105.
- Li, D.; Wu, H.; Zhang, J.; and Huang, K. 2018. A2-RL: Aesthetics aware reinforcement learning for image cropping. In *Proc. Conf. Comput. Vis. Pattern Recog.*, 8193–8201.
- Li, D.; Wu, H.; Zhang, J.; and Huang, K. 2019. Fast a3rl: Aesthetics-aware adversarial reinforcement learning for image cropping. *IEEE Trans. Image Process.*, 28(10): 5105–5120.
- Li, D.; Zhang, J.; Huang, K.; and Yang, M.-H. 2020. Composing Good Shots by Exploiting Mutual Relations. In *Proc. Conf. Comput. Vis. Pattern Recog.*, 4213–4222.
- Lin, T.-Y.; Goyal, P.; Girshick, R.; He, K.; and Dollár, P. 2017. Focal loss for dense object detection. In *Proc. Int. Conf. Comput. Vis.*, 2980–2988.
- Ling, H.-B.; Zhu, B.; Huang, D.; Chen, D.-H.; Wang, C.-D.; and Lai, J.-H. 2022. Vision Transformer for Contrastive Clustering. *arXiv preprint arXiv:2206.12925*.
- Linsker, R. 1988. Self-organization in a perceptual network. *Comput.*, 21(3): 105–117.
- Marchesotti, L.; Cifarelli, C.; and Csurka, G. 2009. A framework for visual saliency detection with applications to image thumbnailing. In *Proc. Int. Conf. Comput. Vis.*, 2232–2239.
- Nishiyama, M.; Okabe, T.; Sato, Y.; and Sato, I. 2009. Sensation-based photo cropping. In *Proc. ACM Int. Conf. Multimedia*, 669–672.
- Pan, Z.; Cao, Z.; Wang, K.; Lu, H.; and Zhong, W. 2021a. TransView: Inside, Outside, and Across the Cropping View Boundaries. In *Proc. Int. Conf. Comput. Vis.*, 4218–4227.
- Pan, Z.; Cao, Z.; Xian, K.; Lu, H.; and Zhong, W. 2022. Discriminate Clearer To Rank Better: Image Cropping By Amplifying View-Wise Differences. In *Proc. Int. Conf. Image Process.*, 3176–3180.
- Pan, Z.; Xian, K.; Lu, H.; and Cao, Z. 2021b. Robust Image Cropping by Filtering Composition Irrelevant Factors. In *Proc. Int. Conf. Image Graphics*, 277–289.
- Sandler, M.; Howard, A.; Zhu, M.; Zhmoginov, A.; and Chen, L.-C. 2018. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proc. Conf. Comput. Vis. Pattern Recog.*, 4510–4520.
- Suh, B.; Ling, H.; Bederson, B. B.; and Jacobs, D. W. 2003. Automatic thumbnail cropping and its effectiveness. In *Proc. Annu. ACM Symp. User Interface Softw. Technol.*, 95–104.
- Torgo, L.; Ribeiro, R. P.; Pfahringer, B.; and Branco, P. 2013. Smote for regression. In *Portuguese Conf. Artif. Intell.*, 378–389.



- Tu, Y.; Niu, L.; Zhao, W.; Cheng, D.; and Zhang, L. 2020. Image cropping with composition and saliency aware aesthetic score map. In *Proc. of AAAI Conf. Artif. Intell.*, volume 34, 12104–12111.
- Van der Maaten, L.; and Hinton, G. 2008. Visualizing data using t-SNE. *J. Mach. Learn. Res.*, 9(11).
- Wang, J.; Zhang, W.; Cao, Y.; Chen, K.; Pang, J.; Gong, T.; Shi, J.; Loy, C. C.; and Lin, D. 2020. Side-aware boundary localization for more precise object detection. In *Eur. Conf. Comput. Vis.*, 403–419.
- Wang, K.; Du, S.; Liu, C.; and Cao, Z. 2022. Interior Attention-Aware Network for Infrared Small Target Detection. *IEEE Trans. Geosci. Remote Sens.*
- Wang, T.; and Isola, P. 2020. Understanding contrastive representation learning through alignment and uniformity on the hypersphere. In *Int. Conf. Mach. Learn.*, 9929–9939.
- Wang, W.; Shen, J.; and Ling, H. 2018. A deep network solution for attention and aesthetics aware photo cropping. *IEEE Trans. Pattern Anal. Mach. Intell.*, 41(7): 1531–1544.
- Wei, Z.; Zhang, J.; Shen, X.; Lin, Z.; Mech, R.; Hoai, M.; and Samaras, D. 2018. Good view hunting: Learning photo composition from dense view pairs. In *Proc. Conf. Comput. Vis. Pattern Recog.*, 5437–5446.
- Wu, Y.; Shi, M.; Du, S.; Lu, H.; Cao, Z.; and Zhong, W. 2022. 3D Instances as 1D Kernels. In Avidan, S.; Brostow, G.; Cissé, M.; Farinella, G. M.; and Hassner, T., eds., *Eur. Conf. Comput. Vis.*, 235–252. Cham: Springer Nature Switzerland.
- Wu, Z.; Xiong, Y.; Yu, S. X.; and Lin, D. 2018. Unsupervised feature learning via non-parametric instance discrimination. In *Proc. Conf. Comput. Vis. Pattern Recog.*, 3733–3742.
- Yan, J.; Lin, S.; Bing Kang, S.; and Tang, X. 2013. Learning the change for automatic image cropping. In *Proc. Conf. Comput. Vis. Pattern Recog.*, 971–978.
- Yang, Y.; Zha, K.; Chen, Y.; Wang, H.; and Katabi, D. 2021. Delving into deep imbalanced regression. In *Intell. Conf. on Mach. Learn.*, 11842–11851.
- Ye, M.; Zhang, X.; Yuen, P. C.; and Chang, S.-F. 2019. Unsupervised embedding learning via invariant and spreading instance feature. In *Proc. Conf. Comput. Vis. Pattern Recog.*, 6210–6219.
- Zeng, H.; Li, L.; Cao, Z.; and Zhang, L. 2019. Reliable and efficient image cropping: A grid anchor based approach. In *Proc. Conf. Comput. Vis. Pattern Recog.*, 5949–5957.
- Zeng, H.; Li, L.; Cao, Z.; and Zhang, L. 2020. Grid Anchor based Image Cropping: A New Benchmark and An Efficient Model. *IEEE Trans. Pattern Anal. Mach. Intell.*
- Zhang, M.; Zhang, L.; Sun, Y.; Feng, L.; and Ma, W. 2005. Auto cropping for digital photographs. In *Proc. IEEE Int. Conf. Multimedia Expo*, 4–7.
- Zhou, S.; Xu, H.; Zheng, Z.; Chen, J.; Bu, J.; Wu, J.; Wang, X.; Zhu, W.; Ester, M.; et al. 2022. A Comprehensive Survey on Deep Clustering: Taxonomy, Challenges, and Future Directions. *arXiv preprint arXiv:2206.07579*.
- Zhuang, C.; Zhai, A. L.; and Yamins, D. 2019. Local aggregation for unsupervised learning of visual embeddings. In *Proc. Int. Conf. Comput. Vis.*, 6002–6012.