

Repurposing existing deep networks for caption and aesthetic-guided image cropping

Horanyi, Nora; Xia, Kedi; Yi, Kwang Moo; Bojja, Abhishake Kumar; Leonardis, Ales; Chang, Hyung Jin

DOI:

[10.1016/j.patcog.2021.108485](https://doi.org/10.1016/j.patcog.2021.108485)

License:

Creative Commons: Attribution-NonCommercial-NoDerivs (CC BY-NC-ND)

Document Version

Peer reviewed version

Citation for published version (Harvard):

Horanyi, N, Xia, K, Yi, KM, Bojja, AK, Leonardis, A & Chang, HJ 2022, 'Repurposing existing deep networks for caption and aesthetic-guided image cropping', *Pattern Recognition*, vol. 126, 108485, pp. 108485. <https://doi.org/10.1016/j.patcog.2021.108485>

[Link to publication on Research at Birmingham portal](#)

General rights

Unless a licence is specified above, all rights (including copyright and moral rights) in this document are retained by the authors and/or the copyright holders. The express permission of the copyright holder must be obtained for any use of this material other than for purposes permitted by law.

- Users may freely distribute the URL that is used to identify this publication.
- Users may download and/or print one copy of the publication from the University of Birmingham research portal for the purpose of private study or non-commercial research.
- User may use extracts from the document in line with the concept of 'fair dealing' under the Copyright, Designs and Patents Act 1988 (?)
- Users may not further distribute the material nor use it for the purposes of commercial gain.

Where a licence is displayed above, please note the terms and conditions of the licence govern your use of this document.

When citing, please reference the published version.

Take down policy

While the University of Birmingham exercises care and attention in making items available there are rare occasions when an item has been uploaded in error or has been deemed to be commercially or otherwise sensitive.

If you believe that this is the case for this document, please contact UBIRA@lists.bham.ac.uk providing details and we will remove access to the work immediately and investigate.

Repurposing Existing Deep Networks for Caption and Aesthetic-Guided Image Cropping

Nora Horanyi^a, Kedi Xia^b, Kwang Moo Yi^c, Abhishake Kumar Bojja^c,
Aleš Leonardis^a, Hyung Jin Chang^{a,*}

^aUniversity of Birmingham, United Kingdom

^bZhejiang University, China

^cUniversity of Victoria, Canada

Abstract

We propose a novel optimization framework that crops a given image based on user description and aesthetics. Unlike existing image cropping methods, where one typically trains a deep network to regress to crop parameters or cropping actions, we propose to *directly* optimize for the cropping parameters by repurposing pre-trained networks on image captioning and aesthetic tasks, without any fine-tuning, thereby avoiding training a separate network. Specifically, we search for the best crop parameters that minimize a combined loss of the initial objectives of these networks. To make the optimization stable, we propose three strategies: (i) multi-scale bilinear sampling, (ii) annealing the scale of the crop region, therefore effectively reducing the parameter space, (iii) aggregation of multiple optimization results. Through various quantitative and qualitative evaluations, we show that our framework can produce crops that are well-aligned to intended user descriptions and aesthetically pleasing.

Keywords: Image cropping, Aesthetics, Deep network re-purposing, Image captioning

1. Introduction

With the advent of social networks, it is now common that images are provided with captions and tags – for example via Instagram or Twitter – where captions them-

*Corresponding author

Email address: h.j.chang@bham.ac.uk (Hyung Jin Chang)



Figure 1: Our proposed framework reliably crops visually pleasing images by following natural descriptions of users. This can produce different crops on the same image based on different users’ descriptions.

selves are highly tied in with the user’s intentions regarding these images. Therefore, an automated process for enhancing images, for example providing artistic crops or making thumbnail images that respect user intentions, would be useful for these social networking platforms (SNS). Besides SNS applications, more tied in with computer vision applications such as semi-automated image dataset generation [1], tracking target object initialization [2], story-based automatic image transition generation [3], can also benefit from text-based image cropping. In this paper, as illustrated in Fig. 1, we focus on a novel image cropping task, which aims to crop an image automatically based on user intent – expressed through a natural text-based description – and the aesthetics of the cropping outcome.

Because of the usefulness of an automated image cropping system, various methods have been suggested. However, existing image cropping methods [4, 5, 6, 7] are typically designed to be purely automatic, leaving the user out of the loop. For example, [6] automatic cropping is based on the maximization of the saliency inside the cropping region. Attention-based methods like this try to preserve the most salient part of the image during cropping. One major downside of them is that the user cannot influence their behaviour. Recent works [8, 9, 10, 11] have focused on making this process even easier through automatically cropping a photo-based on aesthetics. Despite being fully

automated, aesthetics-based methods leave no room for the user to intervene. Furthermore, they provide no guarantee that the initial content and the intent of the images are preserved.

Efforts have also been taken toward methods that take user intention as input. Description-based object detection [12] and localization [13] have recently been proposed for this purpose. However, these description-based methods do not consider how natural images are created, and their behaviours are far from how humans would crop. Most distinctively, these methods provide very tight cropping around an object.

As in many other areas of computer vision, towards this goal, one could apply an end-to-end deep learning framework to find crops that fit the descriptions given an image [12]. However, training such a network in a typical supervised deep learning setup would require an immense amount of labelled data, with captions for multiple sub-regions of the images, as well as their respective “*ground-truth*” crops, which may be subjective depending on the creator of the dataset. This is a challenging task.

Therefore, in this paper, we take an alternative approach that exploits existing networks trained for related tasks – image captioning and aesthetics estimation – and repurpose them to automatically crop the images, thereby avoiding the hardships of training a separate network.

Our contributions are four-fold:

- We propose a new deep networks repurposing framework to optimize crop parameters *directly* using a bilinear sampler [14], a pre-trained image captioning network [15], and a pre-trained aesthetic estimation network [11].
- We optimize to find the crop region that best fits the provided caption in terms of the image captioning network losses, as well as maximizes the aesthetics network scores.
- We generate a new dataset with multiple ground truth bounding box annotations for each caption.
- With approaches above, we were able to not only outperform state-of-the-art methods but also produce more visually pleasing image crops with well-reflecting user’s intention

Utilizing the two networks requires special attention. Since we are, in fact, repurposing the two pre-trained networks for a different purpose, we keep them intact to minimize the change inflicted upon them, and propose a new loss term that we minimize instead of their original ones: learning to generate image captions and measuring how much image looks good. For the image captioning network we propose to ignore the order of the words as we simply want the contents to be accurate. For the aesthetic network, we directly try to maximize the aesthetic score.

Since this optimization process is, in its basic form, highly unstable due to the nature of the image gradients, we further propose a new optimization strategy based on *scale annealing* and *multiple restarts*. Instead of directly optimizing for the position and the scale of the crop, we optimize only for the position and anneal the scale throughout the optimization. We use a *multiple restart* technique, where for each scale we start from a random location and use the average of the optimization outcomes. We further take advantage of the fact that image captions should not differ drastically as the image is blurred, and use a multi-scale representation of the image for the captioning network.

2. Related Works

The image cropping research can be categorized into *attention-based*, *aesthetics-based*, and *description-based* approaches.

Attention-based methods. This class of methods exploits visual saliency models or salient object detectors to find the most visually important regions in the original image [16, 17]. The majority of the attention-based cropping methods rank candidates based on their attention score [18, 19]. Thus these methods are able to identify the most important and attractive regions of the image. Recent methods choose the important area based on certain attention scores [20, 21, 22, 23]. One of the most recent methods [6] focuses on those image regions which most probably attract the human gaze at first sight.

Aesthetics-based methods. Methods based on aesthetics crop images by relying on the attractiveness of the cropped image with the help of a quality classifier [4, 24, 25].

The goal of these methods is to extract the optimal rectangular sub-region of a given image to produce an image with a high aesthetic score [7, 10]. Recently, Li *et al.* [9] formulated the automatic image cropping problem as a sequential decision-making process and proposed a weekly supervised approach which only uses aesthetic information as supervision. This method was the first to use reinforcement learning for automatic image cropping and was able to overcome the disadvantages of the sliding window method. The existing aesthetics-based cropping methods can be further categorized into supervised [26, 27] or weekly supervised [5, 11, 19, 28] methods. In particular, the weekly supervised methods, which do not include bounding box supervision, have been researched actively as producing cropping box annotations for the training is expensive [29].

Description-based methods. These methods aim to localize a region described by a given referring expression. Most of these methods treat comprehension as bounding box localization, which is similar to our cropping task. A recent method by Rohrbach *et al.* [13] uses joint embedding to find the object directly by selecting the best region based on an input expression. Yu *et al.* [12] proposed a modular network for referring expression comprehension. Similarly to our application, MAttNet [12] and Align2Ground [30] are able to locate the image region described by a general referring expression. Their downfall is that these methods rely strictly on the input expressions and do not consider aesthetics at all. In addition, MAttNet heavily relies on three specific decompositions of the expression – subject appearance, location, and relationship to other objects – these may not exist in natural descriptions of images.

3. Methodology

We currently have deep networks designed to do many traditional computer vision tasks, and they perform in many cases even better than the traditional ones. It is also quite common to use pre-trained networks as backbones for performing a certain task. Here, we are proposing an alternative, and use these existing networks as building blocks. We are proposing to go beyond the paradigm that deep networks should give a solution in a single shot, and instead perform inference through optimization, as was a

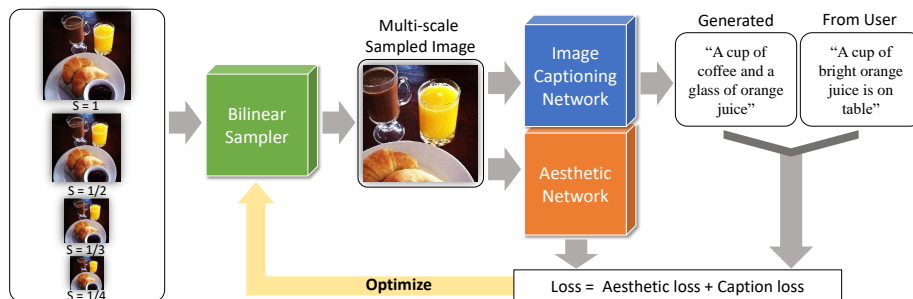


Figure 2: Overall framework of the proposed method (CAGIC). The framework takes an image as input, which goes through multi-scale bilinear sampling to produce a cropped image. Note that the parameters for this sampling do not come from a network as in other existing works, but rather are the parameters that we will directly optimize for later. We then use this cropped region as input to both the image captioning network and the aesthetic network.

common strategy before deep learning.

We first describe the overall architecture, including our *multi-scale sampling strategy*, then explain how we perform inference by *optimizing* the framework instead of training networks to obtain the desired crop region. We further detail how we can stabilize this optimization process through *scale annealing* and *multiple restarts*.

3.1. Framework

Fig. 2 shows the overall framework of the proposed method. Our framework comprises three major components: a) the bilinear sampler that operates on a multi-scale, b) the image captioning network, and c) the aesthetic network.

The image captioning network automatically generates a natural language expression describing the given image content. There is a large body of work on this problem [31]. Among those for the image captioning network, we use the method from [15], with the models pre-trained with the MS-COCO [32] dataset. Visual aesthetic preference can be described either as a single score or as a distribution of scores. We follow the definition in [11], where the aesthetic score is provided by professional photographers. For the aesthetic network we use [11] with the pre-trained models¹.

¹<https://github.com/yiling-chen/view-finding-network>

Note that we take special care that none of the images that were used in the training of any of the pre-trained models is included in our evaluation later. *Despite the fact that these two networks were trained on entirely different datasets, we found that the pre-trained models are good enough for our purpose.*

Multi-scale bilinear sampling. As we are directly optimizing for the location and scale of the crop region, it is essential that the gradients that the bilinear sampling gives are robust. To ensure this, we propose to use a multi-scale strategy, inspired by the observation that, even when an image becomes blurry, its content does not change. Therefore, if we denote the bilinear sampling process as $\text{Sample}(\mathbf{I}, \boldsymbol{\theta})$, where \mathbf{I} is the image and $\boldsymbol{\theta}$ is the crop parameters composed of the center coordinates of the crop x and y , and its scale s , and $\text{Resize}(\cdot)$ is the resizing operation, for the cropped image $\mathbf{I}_{crop}(\boldsymbol{\theta}_s)$ we can write

$$\mathbf{I}_{crop}(\boldsymbol{\theta}_s) = \frac{1}{|\mathbf{S}|} \sum_{s \in \mathbf{S}} \text{Sample}(\text{Resize}(\mathbf{I}, s), \boldsymbol{\theta}), \quad (1)$$

where $\mathbf{S} \in \{\frac{1}{4}, \frac{1}{3}, \frac{1}{2}, 1\}$ is the set of scales, and $|\mathbf{S}|$ is the cardinality of this set, which is four. We also omit \mathbf{I} on the left-hand side for brevity. Here, we set the sampling process to always consider the source image coordinates to be between -1 and 1 , thus removing the need of adjusting the sampling parameters per image. After each resizing operation, we also apply Gaussian image blurring filters.

3.2. Inference

With the framework we infer the parameters of the crop $\boldsymbol{\theta}$ by optimizing the network *w.r.t.* the image captioning networks loss and the aesthetic network output. However, as we are not using the two networks with their original purpose—learning to generate image captions or compare two images to find the one that looks better—, here, we formally introduce our optimization objective.

If we denote the two objectives as $\mathcal{L}_{caption}$ for the image captioning part and $\mathcal{L}_{aesthetic}$ for the aesthetic part, for a given pair of image \mathbf{I} and caption \mathbf{y} , our objective is therefore to find $\hat{\boldsymbol{\theta}}$ such that

$$\hat{\boldsymbol{\theta}} = \arg \min_{\boldsymbol{\theta}} \mathcal{L}_{total}(\mathbf{I}, \mathbf{y}, \boldsymbol{\theta}), \quad (2)$$

where

$$\mathcal{L}_{total}(\mathbf{I}, \mathbf{y}, \boldsymbol{\theta}) = \mathcal{L}_{caption}(\mathbf{I}, \mathbf{y}, \boldsymbol{\theta}) + \lambda \mathcal{L}_{aesthetic}(\mathbf{I}, \boldsymbol{\theta}), \quad (3)$$

and λ is the hyper-parameter that balances the two loss terms. In all our experiments, we empirically set $\lambda = 0.01$. The two losses we choose to optimize are closely related to the networks’ original formulation but modified to our need.

Image Caption Loss $\mathcal{L}_{caption}(\mathbf{I}, \mathbf{y}, \boldsymbol{\theta})$. When training a network to output a reasonable caption, the order of the words is important. However, in our task, this is not necessarily so, as we only want the contents to be what the caption describes. We need not regenerate the sentence that the user inputs. In fact, in our earlier experiments, we found that when the order of the words was considered, depending on the user, the network focused too much on the order of the words, not on the contents, and the performance was poor. Thus, for the image caption loss we choose to ignore the order of the words coming out of the captioning network.

Specifically, if we denote the ground truth one-hot encoded vector representation for the t -th word of the user caption as \mathbf{y}_t , the captioning network as $f(\cdot)$, and cross-entropy as H , we write

$$\mathcal{L}_{caption}(\mathbf{I}, \mathbf{y}, \boldsymbol{\theta}) = H\left(\frac{1}{T_u} \sum_{t=1}^{T_u} \mathbf{y}_t, \frac{1}{T_c} \sum_{t=1}^{T_c} f(\mathbf{I}_{crop}(\boldsymbol{\theta}))_t\right), \quad (4)$$

where T_u and T_c represent the number of words in the user caption and the captioning network f generating caption respectively. Note that we average the word vectors, effectively removing the order information.

Image Aesthetic Loss $\mathcal{L}_{aesthetic}(\mathbf{I}, \boldsymbol{\theta})$. For the aesthetic term we simply aim to maximize the aesthetic score output from the network. If we denote the aesthetic network as $g(\cdot)$, we therefore write

$$\mathcal{L}_{aesthetic}(\mathbf{I}, \boldsymbol{\theta}) = -g(\mathbf{I}_{crop}(\boldsymbol{\theta})). \quad (5)$$

where, as above, \mathbf{I}_{crop} is from Eq. (1) and \mathbf{y} is the user caption. This loss helps our crops to be realistic crops close to how humans would crop images as the aesthetic network learned any photographic rules implicitly encoded in professional photographs [11].

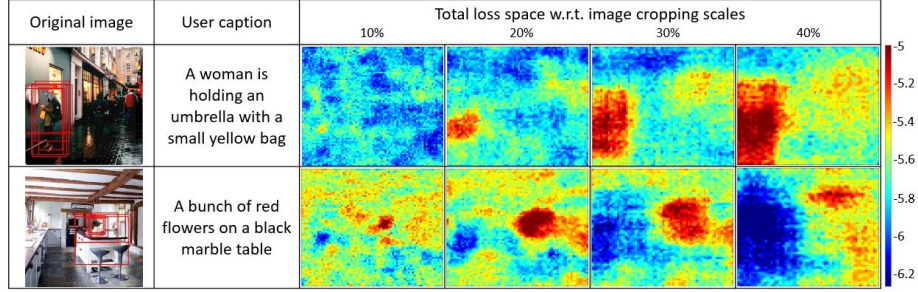


Figure 3: We visualize the total loss space for different image cropping parameters. The percentages indicate the image cropping scales compared to the original image size.

3.3. Stabilizing the Optimization

We design a new cost function considering the two networks' outputs together searching among various bilinear samples.

In Fig. 3 we visualized the total loss space (The extension of this figure can be found in Section 5. of the Supplementary material). During the optimization process, the cropping parameter is shrinking by 2% in every iteration (See more details in Scale Annealing). The percentages in this figure correspond to the cropping sizes *w.r.t.* the original image size. We can see that choosing the correct scale is important in order

Algorithm 1 Optimization with multiple restart.

Require: \mathbf{I} : input image, \mathbf{y} : user caption

```

1: function OPTIMIZE( $\mathbf{I}$ ,  $\mathbf{y}$ )
2:   for  $i = 1$  to  $S$  do                                     ▷ For each scale level
3:     for  $k = 1$  to  $K$  do                                     ▷ Optimize  $K$  times
4:        $x_0 \sim \mathcal{N}(x^i, \sigma)$                              ▷ Restart initialization
5:        $y_0 \sim \mathcal{N}(y^i, \sigma)$ 
6:        $\hat{x}_k^i, \hat{y}_k^i \leftarrow \text{Eq. (7)}$                    ▷ Find optimum point
7:     end for
8:      $x^{i+1} \leftarrow \frac{1}{K} \sum_{k=1}^K \hat{x}_k^i$                  ▷ Gather result
9:      $y^{i+1} \leftarrow \frac{1}{K} \sum_{k=1}^K \hat{y}_k^i$ 
10:  end for
11: end function

```

to ensure that the crop will include all the relevant parts of the image based on the user’s description. However, the image cropping parameter space is too large and non-convex as shown in Fig. 3, so unstable converge is inevitable. Our initial attempts indirectly optimizing Eq. (2) were not very successful, even with the help of the scale-space bilinear sampling in Eq. (1). We, therefore, propose two additional methods that stabilize the optimization process, leading to better final results. We explain these methods below and summarize the entire optimization algorithm in Alg. 1.

Scale Annealing. One hardship when directly optimizing for the crop parameters, is that once determined to be outside of the crop region, pixels outside the crop region have no means to affect the optimization process. This leads to instability when optimizing also for scale, as for example, when the crop accidentally shrinks, it will be difficult for the system to recover from it. We, therefore, choose to exclude the scale parameter from optimization and to anneal the scale to become smaller throughout the optimization process. Specifically, we set the scale to be s^i , where i is the optimization iteration. We empirically set $s = 0.98$. In other words, the scale is reduced by 2% at each iteration. During optimization, we keep tracking which crop parameter gave the lowest loss so far and return that crop region as our final result.

Optimizing with Multiple Restarts. To further stabilize the optimization process and escape local optima, we employ a multiple restart technique [33] as shown in Fig. 4. Based on the research of [34], for each optimization iteration, we apply random noise $\Delta x \sim \mathcal{U}(0, 1)$ and $\Delta y \sim \mathcal{U}(0, 1)$, where $\mathcal{U}(0, 1)$ is the uniform distribution, to the outcome of the previous iteration. After adding the noise, we further clip them to prevent the cropped region from going out of the image border. We then run our optimization to find the optimal crop for this given scale, and repeat the process K times, and average their results to obtain our final solution for this scale (in our experiments, we use the $K = 10$).

If we denote the crop center estimates at iteration i as x^i and y^i , we write

$$(x^{i+1}, y^{i+1}) = \frac{1}{K} \sum_{k=1}^K (\hat{x}_k^i, \hat{y}_k^i), \quad (6)$$

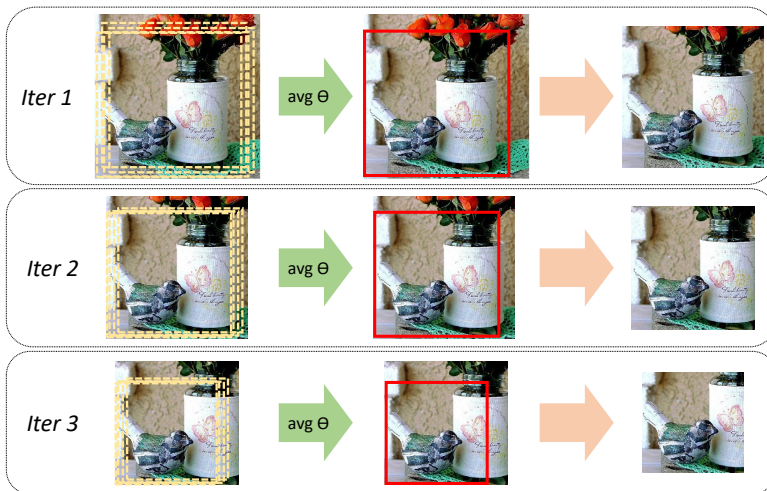


Figure 4: For each iteration, we employ the multiple restart strategy which starts random location and takes the average of the optimization outcomes. In the next iteration, the scale is annealed as scheduled.

where

$$(\hat{x}_k^i, \hat{y}_k^i) = \arg \min_{x, y | x_0 \sim \mathcal{U}(0,1), y_0 \sim \mathcal{U}(0,1)} \mathcal{L}_{total}(\mathbf{I}, \mathbf{y}, (x, y, s^i)), \quad (7)$$

and x_0 and y_0 are the starting points of optimization for x and y respectively, and x_i and y_i are the final converged locations for this scale level. To find the $\hat{\theta}$, we apply limited-memory Broyden-Fletcher-Goldfarb-Shanno (L-BFGS) [35], as we are searching for the optimal location and not simply seeking to perform gradient descent for each scale. The L-BFGS approach is one of the most popular quasi-Newton methods that construct positive definite Hessian approximations. It is a local search algorithm that is intended for convex optimization problems with a single optimum. Thus, the L-BFGS is suitable for our method, as we apply scale annealing and prefer local optimizations.

Although the multiple restart strategy does not theoretically guarantee that the average location is lower in terms of the final objective, it ensures that we are always optimizing towards the general improving direction.

4. Experiments

We implement our method in TensorFlow [36]². All experiments are run on an Intel i7- CPU @ 3.40GHZ, 16 GB RAM, and two NVIDIA TITAN Xp GPU.

4.1. Dataset and baselines

As the task we aim for – cropping based on natural descriptions – is different from what existing datasets can offer, we create a novel dataset. Before we discuss our dataset in detail, we first briefly review existing datasets.

Deficiencies of existing datasets. A widely used dataset for natural images with captions is the MS-COCO dataset [32], which contains 123k images, with predefined splits for training, validation, and testing. Each image is annotated with five captions by Amazon Mechanical Turkers. However, most of the captions describe the whole image, and is not suitable for caption-guided cropping. Therefore, based on MS-COCO, RefCOCO [37] dataset was proposed, where parts of images are described with annotated bounding boxes. This dataset, however, is heavily skewed towards tight detection of certain objects, and are not suitable for creating natural crops. They are aimed more at text-guided object detection.

Dataset	Caption
MS-COCO	A puppy dog sitting in an office with a computer in it
refCOCO	Chair to the left
Proposed	A computer and a lamp on the top of the desk

Figure 5: Caption comparison between MS-COCO, refCOCO, and our new dataset on ImgId=38353 of the MS-COCO train2014 set.

In Fig. 5 we show example captions on the same image from different datasets to demonstrate that captions from refCOCO [37], MS-COCO [32] are not adequate for

²To guarantee reproducibility, we will release our code if the paper is accepted

our tasks. As one can see, the caption from MS-COCO describes the entire image, while the captions from refCOCO is very short and related to a single object only.

The Visual Genome dataset [38] is perhaps the closest to what we need for our task. This large dataset consists of 108k images from the YFCC100M [39] and MS-COCO datasets. Different from the original captions of the MS-COCO dataset, the Visual Genome provides about 50 descriptions of different image regions. The annotations for this dataset is also geared towards more traditional localization and description tasks. However, there are two problems in this dataset that makes it not feasible for evaluating natural crops: (i) annotations are tight bounding boxes focused on objects – that do not represent how people typically take photos; (ii) identical captions can denote multiple regions – for example “red sky” could correspond to any bounding box within the sky – demonstrating that annotations are only part of the possible “ground truth”, and not overlapping with any of these do not mean that a crop is wrong.

As shown in Fig. 6, the captions and annotations would not look at all like natural crops, if they were to be cut out – they are either too tight and sometimes not even capturing the entire object, as shown by the examples. Also notice that the captions are very short, almost as if they are descriptions of a single object. There is also an issue of non-specific regions as in the case of Fig. 6 right, where the descriptions all talk about the sky and the cloud – in addition, none of the crops would actually be recognizable to a human being on what they are about. In the supplementary appendix, we show qualitative and quantitative results on this dataset.




Figure 6: Example captions and corresponding annotations from the Visual Genome dataset.

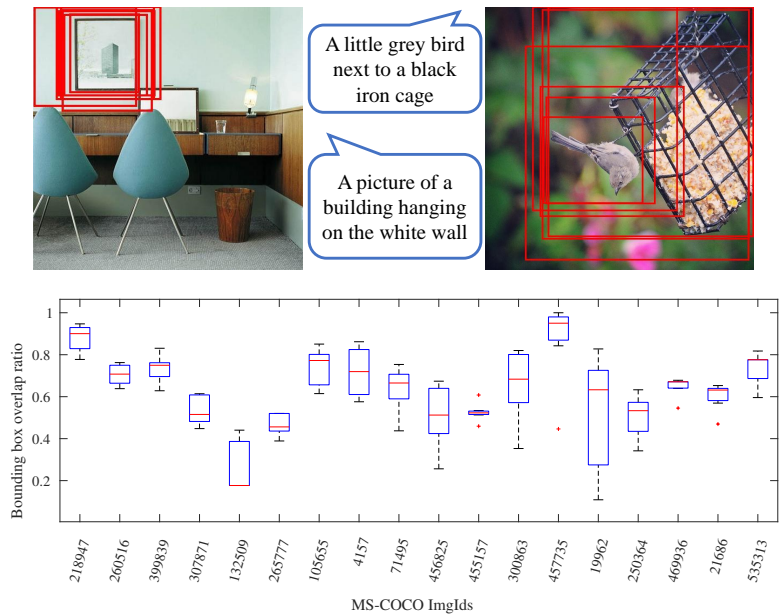


Figure 7: (Top) Illustration of the diversity of the bounding box annotations for different images. (Bottom) Box plots of the distribution of ground truth overlaps for different MS-COCO images.

4.1.1. A novel dataset

Therefore, we create a novel dataset based on MS-COCO, similarly to the refCOCO and Visual Genome datasets. We select 100 images randomly from the MS-COCO test set to avoid the images being ever seen by any of our pre-trained networks and create our own captions for each image.

Ground-truth captions. In our new dataset, we are interested in evaluating the ability to automatically generate crops related to a given caption. We therefore manually generate descriptive expressions, focusing on distinctive parts of the image. Our captions are roughly 10 words long on average to ensure that they are specific and clear.

Ground-truth crops. One of the tricky parts in creating a dataset for caption-based image cropping is the definition of ground truth. The concept of “which image region fits the description well” is subjective and can wildly differ from person to person. Therefore, to remove the subjective nature of our ground truth annotations as much as possible, we carefully select regions of the image that are unique and distinctive

for a human annotator to create a “natural” crop out of. We then asked 8 participants to generate ground truth crops based on our captions individually. We further asked the participants to consider the aesthetics of these crops. We use these ground truth annotations to perform quantitative comparisons. As an evaluation metric we use the Intersection over Union (IoU), a standard metric for evaluating bounding box-based tasks [11, 40].

Even with care, it is inevitable that the cropping task is subjective and dependent on the annotator. We show an example of the ground-truth annotations in Fig. 7 (top), as well as the agreement between annotators in Fig. 7 (bottom). Note that even when disagreeing on the exact crop, they are all overlapping – the main content is shared.

4.1.2. Baselines

We compare our method against the following eight different methods: **GradCAM** [41] – a naive untrained baseline where we apply GradCAM with the captioning network to extract regions in the image corresponding to the user caption. We then threshold the activation map with a threshold of 0.2 of the maximum value, which we empirically set; **A2-RL** [9], **VPN** [29], **Anchor** [42] – full automatic methods to demonstrate that these crops do not necessarily correspond to the users’ intentions; **GradCAM+A2-RL/VPN/Anchor** – GradCAM and the auto-crop networks combine sequentially; **MAttNet** [12] – the state-of-the-art network for tight localization from text descriptions.

4.2. Comparison with the state-of-the-art

Qualitative results. We first show qualitative highlights in Fig. 9. As shown, our method provides crops of the highest quality. In Fig. 8, we show that our method can deliver entirely different results on the same image, when different captions are provided. Notice how our results are not only corresponding well to the provided captions, but also good looking, demonstrating that our method successfully considers both aspects. In particular, as shown in Fig. 8 (b) and (f), our method well crops regions when the description is about parts of the image that are small or are located near the edge.

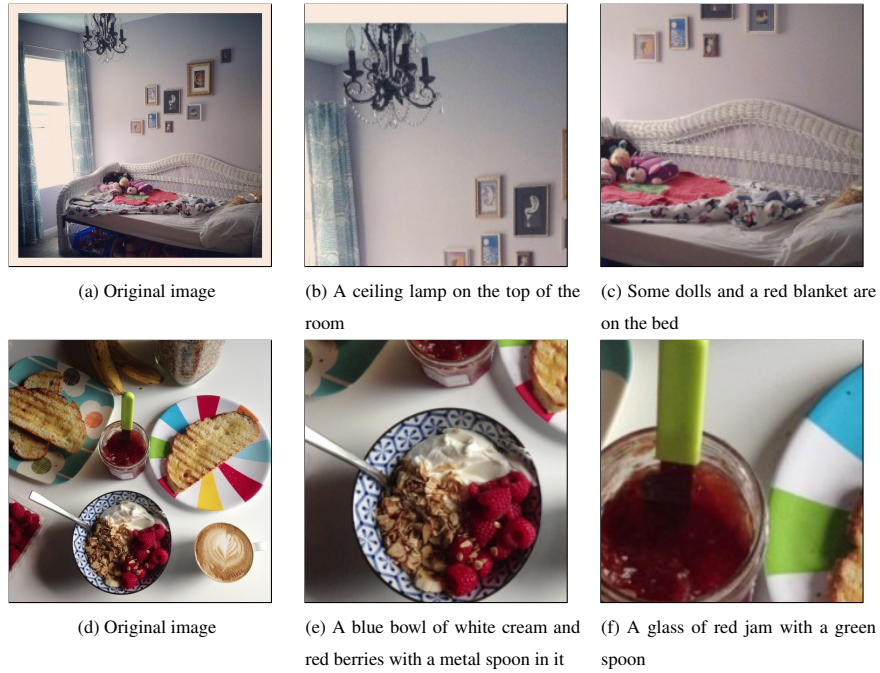


Figure 8: Different captions can lead to entirely different crops even on a single image.

Quantitative results – IoU. We report the average IoU between produced crops and all eight of the ground-truth annotations in Table 1. As shown, our method provides the highest IoU value among the compared methods. Note that we have higher numbers than even those which were trained specifically for referencing from captions. This is mainly due to the fact that existing methods perform tight crops. However, when a human is asked to perform the same task, they tend to include context.

User study – is the crop really what we want? As one of the main goals here is to faithfully follow user guidance, we evaluate it by asking users to re-annotate the crop outcomes. We asked four users, who were not exposed previously to the original images, to caption cropped images with natural descriptions. We then evaluate how similar these new descriptions are to the target descriptions of our dataset. We evaluate only the top three methods from the IoU evaluation. We report results in terms of

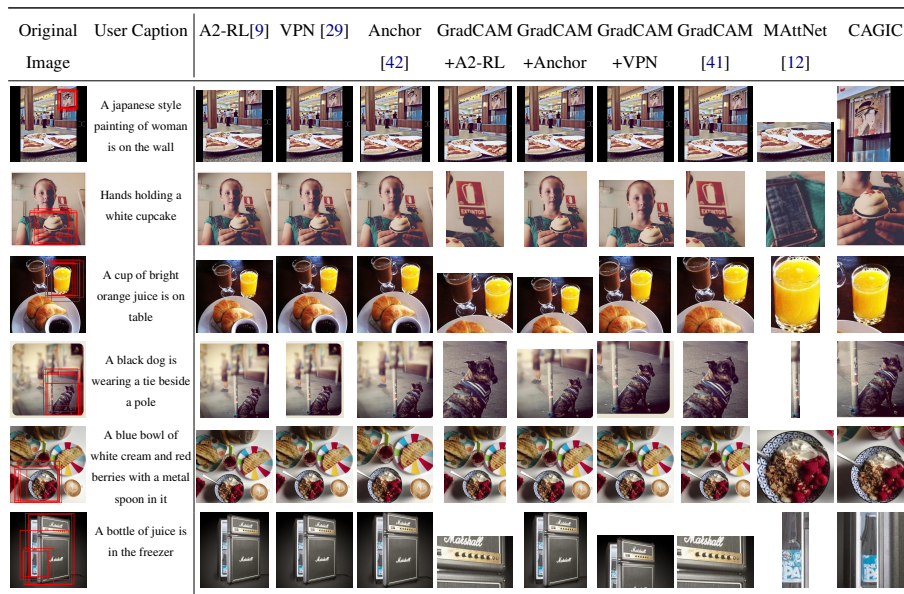


Figure 9: The cropped images obtained by the proposed method and the eight baseline methods. The user defined ground truth bounding box annotations are shown on the original images with red. The proposed method well crops the images as user described.

Method	Mean \pm Std.
Original	0.2869 \pm 0.0280
A2-RL[9]	0.2975 \pm 0.0238
VPN[29]	0.3151 \pm 0.0233
Anchor [42]	0.3325 \pm 0.0236
GradCAM+A2-RL	0.3465 \pm 0.0108
GradCAM+Anchor	0.3554 \pm 0.0199
GradCAM+VPN	0.3561 \pm 0.0126
GradCAM[41]	0.3597 \pm 0.2017
MAttNet[12]	0.3851 \pm 0.2607
CAGIC	0.4160 \pm 0.0129

Table 1: Quantitative comparison of the different methods using IoU measure on the output bounding boxes.

Method	Blue-1	Blue-2	Blue-3	Blue-4	METEOR	ROUGE.L	CIDEr
GradCAM[41]	0.2728	0.1410	0.0874	0.0531	0.1182	0.2790	0.6973
MAttNet[12]	0.1718	0.0937	0.0603	0.0355	0.1132	0.2947	0.7154
CAGIC	0.3424	0.1876	0.1017	0.0631	0.1702	0.2970	0.9054

Table 2: Comparison on user intention presence. We ask users to caption cropped images and compare with natural language metrics how similar they are with the original desired caption.

	Original Image	MAttNet[12]	GradCAM[41]	CAGIC
Aggregated percentage (%)	21.04	23.93	25.51	29.52

Table 3: Top three methods of the qualitative comparison along with original image were compared through human survey and evaluated by aggregation.

metrics used for natural language processing³ in Table 2. Our method provides the best results for all metrics, demonstrating that the user intention is best preserved.

User study – which crop is better?

Due to the subjective nature of our task, we further perform a user study, where we ask users to select which image is preferred over the top-3 methods, as well as the original image as the baseline. Specifically, we ask users to “*Select the crop described by the caption which looks the best,*” given the four images in a graphical user interface. We ask a total of 13 users, resulting in more than 1000 decisions.

In the user study, for each user, we show a randomly selected subset of our dataset. We report the probability of being selected for each method in Table 3. Our method is the most preferred among compared methods. Our user study shows that the subjects preferred the cropped images over the original image. Users preferred the methods with wider output crops that were not mainly focused on the subject of the image caption;

³<https://github.com/Maluuba/nlg-eval>




Figure 10: Qualitative comparison between our methods which consist of the combination of Caption Network (CN), Scale Anneal (SA), Multiple Restart (MR) and Multi-Scale Bilinear Sampling (MS). The results show that the combination of all of these elements (proposed full method) together can provide the best output image.

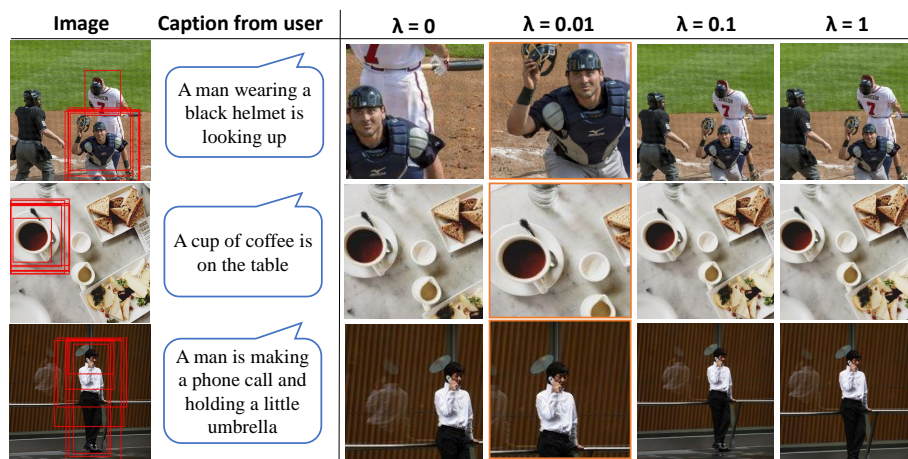


Figure 11: Comparison of different aesthetics ratios.

therefore, they were able to preserve the contextual information. While the output of GradCAM was contextually correct, users tend to select the proposed method when the aesthetics of the crop was not pleasing. This highlights that aesthetics need to be considered when emulating what humans would do.

Method	Mean \pm Std.
CN	0.3338 \pm 0.0693
CN+SA	0.3406 \pm 0.0584
CN+SA+MR	0.3562 \pm 0.0565
CN+SA+MR+MS	0.3816 \pm 0.0615
CAGIC (Full)	0.4160 \pm 0.0129

Table 4: Quantitative result of the ablation study using IoU measure on the output bounding boxes

4.3. Ablation Study

To motivate our design choices, we present a qualitative comparison of our method with various components of our full pipeline disabled; see Fig. 10. We compare against five variants. The full method uses the caption network (CN) along with scale anneal (SA), multiple restart (MR), multi-scale bilinear sampling (MS) as well as aesthetics. As shown, as we introduce aesthetics scores to the system we are able to produce more visually appealing crops of the original image. In the case of CN+SA+MR+MS, the method is able to find the relevant part of the original image, but the contents are not placed in the centre, as can be noticed in the examples shown in the last two rows. Multiple restart help the approach to centre the content better, while scale annealing prevents drastic scale changes. In short, the full method delivers the best outcome in all cases.

The quantitative results of the ablation study are in Table 4. show that every component of the proposed method is effective and that using aesthetics on top of the different optimization and search stabilization methods improved the results. In addition, we further demonstrate the effectiveness of the aesthetics loss in our framework in Fig. 11. As shown, the produced crops are most similar to the ground truth regions when the aesthetics loss is enabled.

4.4. Runtime

Currently, our un-optimized implementation is unable to run real-time — requires an average of 2.06 seconds per single optimization iteration. The runtime of the other

deep learning-based methods where the input goes through one pre-trained network is shorter compared to our optimization-based method. However, despite our method not being real-time, it was clearly demonstrated that it performs better and it is flexible. Our future work would be to resolve this, similar to how style transfer [43] started off taking minutes per image, but is now able to run real-time [44].

5. Conclusion

In this paper, we have proposed a novel optimization framework that produces image crops that follow users’ descriptions and aesthetics criteria. The main idea behind this study was to demonstrate that this can successfully be achieved without training a specialized network, but instead utilizing two pre-trained networks on related tasks, namely image captioning and aesthetics measuring. We designed a new cost function considering the two networks’ outputs together and performed a search among various bilinear sampling parameters. However, the parameter space is very large and non-convex, so to achieve a stable and efficient solution, we designed a new scale annealing and multiple restarts search strategy. We have shown that our proposed method outperforms other existing approaches which are based on either solely saliency-based or caption-based cropping methods.

Our ongoing research is extending in two directions. The description- and aesthetics-driven captioning will further be coupled with attention mechanisms and gaze estimation. On a more conceptual level, we are extending the ideas of solving complex tasks in optimization frameworks by repurposing existing modular networks trained on auxiliary tasks, thereby seeking efficient alternatives to extensive learning of networks from scratch.

Acknowledgements

This work was partially supported by the Natural Sciences and Engineering Research Council of Canada (NSERC) Discovery Grant “Deep Visual Geometry Machines” (RGPIN-2018-03788), and by systems supplied by Compute Canada. We acknowledge MoD/Dstl and EPSRC for providing the grant to support the UK academics

involvement in a Department of Defense funded MURI project through EPSRC grant EP/N019415/1, and IITP grant funded by the Korea government (MSIT) (IITP-2020-0-01789, ITRC support program). We thank NVIDIA for sponsoring one GeForce Titan Xp GPU. This work was supported by the Institute of Information and communications Technology Planning and evaluation (IITP) grant funded by the Korean government (MSIT) (2021-0-00537, Visual common sense through self-supervised learning for restoration of invisible parts in images).

References

- [1] W. Huang, C. L. Tan, J. Zhao, Generating ground truthed dataset of chart images: Automatic or semi-automatic?, in: GREC, 2007.
- [2] I. A. Iswanto, B. Li, Visual object tracking based on mean-shift and particle-kalman filter, *Procedia computer science* 116 (2017) 587–595.
- [3] W.-T. Chu, C.-H. Yu, H.-H. Wang, Optimized comics-based storytelling for temporal image sequences, *IEEE Transactions on Multimedia* 17 (2) (2014) 201–215.
- [4] J. Chen, G. Bai, S. Liang, Z. Li, Automatic image cropping: A computational complexity study, 2016, pp. 507–515. [doi:10.1109/CVPR.2016.61](https://doi.org/10.1109/CVPR.2016.61).
- [5] Y. Kao, R. He, K. Huang, Automatic image cropping with aesthetic map and gradient energy map, *IEEE*, 2017, pp. 1982–1986.
- [6] M. Cornia, S. Pini, L. Baraldi, R. Cucchiara, Automatic image cropping and selection using saliency: An application to historical manuscripts, in: *Italian Research Conference on Digital Libraries*, Springer, 2018, pp. 169–179.
- [7] G. Guo, H. Wang, C. Shen, Y. Yan, H. M. Liao, Automatic image cropping for visual aesthetic enhancement using deep neural networks and cascaded regression, *IEEE Transactions on Multimedia* 20 (8) (2018) 2073–2085. [doi:10.1109/TMM.2018.2794262](https://doi.org/10.1109/TMM.2018.2794262).
- [8] N. Shan, D. S. Tan, M. S. Deneke, Y.-Y. Chen, W.-H. Cheng, K.-L. Hua, Photobomb defusal expert: Automatically remove distracting people from photos,

- IEEE Transactions on Emerging Topics in Computational Intelligence (99) (2018) 1–11.
- [9] D. Li, H. Wu, J. Zhang, K. Huang, A2-RL: Aesthetics Aware Reinforcement Learning for Image Cropping, 2018, pp. 8193–8201.
- [10] Y.-L. Chen, T.-W. Huang, K.-H. Chang, Y.-C. Tsai, H.-T. Chen, B.-Y. Chen, Quantitative analysis of automatic image cropping algorithms: A dataset and comparative study, IEEE, 2017, pp. 226–234.
- [11] Y.-L. Chen, J. Klopp, M. Sun, S.-Y. Chien, K.-L. Ma, Learning to compose with professional photographs on the web, in: ACM International Conference on Multimedia, ACM, 2017, pp. 37–45.
- [12] L. Yu, Z. Lin, X. Shen, J. Yang, X. Lu, M. Bansal, T. L. Berg, Mattnet: Modular attention network for referring expression comprehension, in: CVPR, 2018, pp. 1307–1315.
- [13] A. Rohrbach, M. Rohrbach, R. Hu, T. Darrell, B. Schiele, Grounding of textual phrases in images by reconstruction, in: ECCV, Springer, 2016, pp. 817–834.
- [14] M. Jaderberg, K. Simonyan, A. Zisserman, K. Kavukcuoglu, Spatial transformer networks, in: NIPS, 2015.
- [15] K. Xu, J. Ba, R. Kiros, K. Cho, A. Courville, R. Salakhudinov, R. Zemel, Y. Bengio, Show, attend and tell: Neural image caption generation with visual attention, in: International conference on machine learning, 2015, pp. 2048–2057.
- [16] J. Huang, H. Chen, B. Wang, S. Lin, Automatic thumbnail generation based on visual representativeness and foreground recognizability, 2015, pp. 253–261.
- [17] N. Jaiswal, Y. K. Meghrajani, Saliency based automatic image cropping using support vector machine classifier, in: ICIIIECS, IEEE, 2015, pp. 1–5.
- [18] J. Choi, C. Kim, Object-aware image thumbnailing using image classification and enhanced detection of roi, Multimedia Tools and Applications 75 (23) (2016) 16191–16207.

- [19] C. Fang, Z. Lin, R. Mech, X. Shen, Automatic image cropping using visual composition, boundary simplicity and content preservation models, in: ACM international conference on Multimedia, ACM, 2014, pp. 1105–1108.
- [20] W. Wang, J. Shen, F. Porikli, Saliency-aware geodesic video object segmentation, 2015, pp. 3395–3402.
- [21] W. Wang, J. Shen, L. Shao, Consistent video saliency using local gradient flow optimization and global refinement, *IEEE Transactions on Image Processing* 24 (11) (2015) 4185–4196.
- [22] W. Wang, J. Shen, L. Shao, F. Porikli, Correspondence driven saliency transfer 25 (11) (2016) 5025–5034.
- [23] W. Wang, J. Shen, R. Yang, F. Porikli, Saliency-aware video object segmentation (1) (2018) 20–33.
- [24] M. B. Islam, W. Lai-Kuan, W. Chee-Onn, A survey of aesthetics-driven image recomposition, *Multimedia Tools and Applications* 76 (7) (2017) 9517–9542.
- [25] P. Lu, H. Zhang, X. Peng, X. Peng, Aesthetic guided deep regression network for image cropping, *Signal Processing: Image Communication* (2019).
- [26] E. Hong, J. Jeon, S. Lee, Cnn based repeated cropping for photo composition enhancement, 2017.
- [27] P. Wang, Z. Lin, R. Mech, Learning an aesthetic photo cropping cascade, *IEEE*, 2015, pp. 448–455.
- [28] L. Zhang, M. Song, Y. Yang, Q. Zhao, C. Zhao, N. Sebe, Weakly supervised photo cropping, *IEEE Transactions on Multimedia* 16 (1) (2014) 94–107.
- [29] Z. Wei, J. Zhang, X. Shen, Z. Lin, R. Mech, M. Hoai, D. Samaras, Good view hunting: learning photo composition from dense view pairs, in: *CVPR*, 2018, pp. 5437–5446.

- [30] S. Datta, K. Sikka, A. Roy, K. Ahuja, D. Parikh, A. Divakaran, Align2ground: Weakly supervised phrase grounding guided by image-caption alignment, in: Proceedings of the IEEE International Conference on Computer Vision, 2019, pp. 2601–2610.
- [31] S. Bai, S. An, A survey on automatic image caption generation, *Neurocomputing* 311 (2018) 291–304.
- [32] T.-Y. Lin, M. Maire, S. J. Belongie, L. D. Bourdev, R. B. Girshick, J. Hays, P. Perona, D. Ramanan, P. Dollár, C. L. Zitnick, Microsoft coco: Common objects in context, 2014.
- [33] G. Song, B. B. Avants, J. C. Gee, Multi-start method with prior learning for image registration, 2007, pp. 1–8. [doi:10.1109/ICCV.2007.4409159](https://doi.org/10.1109/ICCV.2007.4409159).
- [34] W. Kwedlo, A new random approach for initialization of the multiple restart em algorithm for gaussian model-based clustering, *Pattern Analysis and Applications* 18 (4) (2015) 757–770.
- [35] R. H. Byrd, P. Lu, J. Nocedal, C. Zhu, A limited memory algorithm for bound constrained optimization, *SIAM Journal on Scientific Computing* 16 (5) (1995) 1190–1208.
- [36] M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard, M. Kudlur, J. Levenberg, R. Monga, S. Moore, D. Murray, B. Steiner, P. Tucker, V. Vasudevan, P. Warden, M. Wicke, Y. Yu, X. Zheng, Tensorflow: A System for Large-Scale Machine Learning, in: USENIX Conference on Operating Systems Design and Implementation, 2016, pp. 265–283.
- [37] S. Kazemzadeh, V. Ordonez, M. Matten, T. Berg, Referitgame: Referring to objects in photographs of natural scenes, in: EMNLP, 2014, pp. 787–798.
- [38] R. Krishna, Y. Zhu, O. Groth, J. Johnson, K. Hata, J. Kravitz, S. Chen, Y. Kalantidis, L.-J. Li, D. A. Shamma, et al., Visual genome: Connecting language and

- vision using crowdsourced dense image annotations, *International journal of computer vision* 123 (1) (2017) 32–73.
- [39] B. Thomee, D. A. Shamma, G. Friedland, B. Elizalde, K. Ni, D. Poland, D. Borth, L.-J. Li, Yfcc100m: The new data in multimedia research, *Communications of the ACM* 59 (2) (2016) 64–73.
- [40] M. Kristan, J. Matas, A. Leonardis, T. Vojř, R. Pflugfelder, G. Fernandez, G. Nebehay, F. Porikli, L. Čehovin, A novel performance evaluation methodology for single-target trackers, *IEEE transactions on PAMI* 38 (11) (2016) 2137–2155.
- [41] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, D. Batra, Grad-cam: Visual explanations from deep networks via gradient-based localization (2017) 618–626.
- [42] H. Zeng, L. Li, Z. Cao, L. Zhang, Reliable and efficient image cropping: A grid anchor based approach, in: *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [43] L. A. Gatys, A. S. Ecker, M. Bethge, *Image Style Transfer Using Convolutional Neural Networks*, 2016.
- [44] J. Johnson, A. Alahi, L. Fei-fei, *Perceptual Losses for Real-Time Style Transfer and Super-Resolution*, 2016, pp. 694–711.