# A Generic Visualization Approach for Convolutional Neural Networks

Ahmed Taha, Xitong Yang, Abhinav Shrivastava, and Larry Davis

University of Maryland, College Park

Abstract. Retrieval networks are essential for searching and indexing. Compared to classification networks, attention visualization for retrieval networks is hardly studied. We formulate attention visualization as a constrained optimization problem. We leverage the unit <u>L2</u>-Norm <u>constraint</u> as an <u>attention filter (L2-CAF)</u> to localize attention in both classification and retrieval networks. Unlike recent literature, our approach requires neither architectural changes nor fine-tuning. Thus, a pre-trained network's performance is never undermined

L2-CAF is quantitatively evaluated using weakly supervised object localization. State-of-the-art results are achieved on classification networks. For retrieval networks, significant improvement margins are achieved over a Grad-CAM baseline. Qualitative evaluation demonstrates how the L2-CAF visualizes attention per frame for a recurrent retrieval network. Further ablation studies highlight the computational cost of our approach and compare L2-CAF with other feasible alternatives. Code available at https://bit.ly/3iDBLFv

### 1 Introduction

Both classification and retrieval neural networks need attention visualization tools. These tools are important in medical and autonomous navigation to understand and interpret networks' decisions. Moreover, attention visualization enables weakly supervised object localization (WSOL) which reduces the cost of data annotation. WSOL avoids bounding-box labeling required by fully supervised approaches. Attention visualization and WSOL have been intensively studied for classification architectures [44, 48, 34, 32, 45, 7]. However, these approaches do not address retrieval networks. In this paper, we leverage the unit L2-Norm constraint as an attention filter (L2-CAF) that works for both classification and retrieval neural networks, as shown in Figure 1.

For classification networks, Zhou *et al.* [48] propose class activation maps (CAM) for attention visualization and WSOL. Further research [34, 45, 46, 7] improved WSOL by augmenting the most discriminative region with other less discriminative parts, *e.g.*, augment a cat's head with its legs. This improvement comes at the cost of few drawbacks: (1) They impose architectural constraints, *e.g.*, global average pooling (GAP) layer; (2) While fine-tuning boosts localization efficiency, it degrades classification accuracy. Grad-CAM [32] avoids these limitations, but it is originally formulated for classification networks.



**Fig. 1.** L2-CAF enables both class-oblivious and class-specific visualizations. This separates our work from dominant literature that targets classification networks only. The supplementary video shows more vivid and challenging visualizations.



**Fig. 2.** An overview of the proposed unit L2-Norm constrained attention filter (L2-CAF). Given a pre-trained CNN with an auxiliary head (Aux NN), feed an input frame through a normal feed-forward pass (green solid path) to generate the network output logits/embedding NT(x). Then, feed the same input again but multiply the last convolutional feature map by a constrained attention filter f (orange dashed path) to generate a new filtered output FT(x, f). Optimize the filter's weights through gradient descent to minimize the difference between NT(x) and FT(x, f). In standard CNN architectures, the L2-CAF is typically  $7 \times 7$ , *i.e.*, a cheap optimization problem  $\in \mathbb{R}^{49}$ .

Retrieval networks are essential for visual search [25, 19], zero-shot learning [3, 43, 47], and fine-grained retrieval [36, 26]. The large metric learning [26, 40, 6] and product quantization [4, 23, 9] literature reflect their importance. Despite that, attention visualization for retrieval networks has not been evaluated quantitatively. It is more challenging compared to classification due to the network's output – a class-oblivious embedding.

The main contribution of this paper is to leverage the L2-CAF as a visualization filter to identify key features of both classification and retrieval networks' output. Figure 2 illustrates the approach. Given a pre-trained CNN, feeding the same input x through the network (green solid path) will always generate the same output NT(x). If the final convolutional feature map is multiplied by a constrained attention filter f in an element-wise manner (orange dashed path), the network generates a filtered output FT(x, f). Through gradient descent, we optimize f to minimize the L2 loss  $L = ||NT(x) - FT(x, f)||^2$ . The optimized filter f reveals key spatial regions, e.g., the cat's head. The filter size  $(f_w, f_h)$ depends on the convolution layer size, e.g., the last convolution layer in standard CNNs  $\in \mathbb{R}^{7\times7}$ .

This approach imposes no constraints on the network architecture besides having a convolution layer. The input can be a regular image or a pre-extracted convolutional feature. The network output can be logits trained with softmax or a feature embedding trained with a ranking loss. Furthermore, this approach neither changes the original network weights nor requires fine-tuning. Thus, network performance remains intact. The visualization filter is applied only when an attention map is required. Thus, it poses no computational overhead during inference. L2-CAF visualizes the attention of the last convolutional layer of GoogLeNet within 0.3 seconds.

Section 3 describes two variants of the L2-CAF and their mathematical optimization details. The first is the class-oblivious variant illustrated in Figure 2. The second is the class-specific variant for classification networks to localize objects of a specific class. We also present a technique to reduce the computational cost of the L2-CAF's optimization formulation. We benchmark our approach quantitatively using WSOL for both classification and retrieval architectures.

In summary, the key contributions of this paper are:

- 1. A novel attention visualization approach for *both* classification and retrieval networks (Sec. 3). This approach achieves state-of-the-art WSOL results using classification architectures (Sec. 4.1).
- A modified Grad-CAM to better support WSOL on retrieval networks (Sec. 4.2); L2-CAF achieves significant localization improvement margins, up to an absolute 36%, compared to the vanilla Grad-CAM.
- 3. A method to visualize attention for video frames that are temporally fused using a recurrent network (Sec. 4.3).

# 2 Related Work

This section briefly reviews weakly supervised object localization (WSOL) for classification networks. Grad-CAM is reviewed in the WSOL retrieval evaluation Section 4.2. The supplementary material reviews different ranking losses (*e.g.*, N-pair). Figure 3 presents a high-level categorization of WSOL approaches in terms of (1) supported architectures; (2) whether fine-tuning is required or not? The experiment section provides further one-to-one comparisons.

Classification networks' attention visualization increases interpretability and enables WSOL. CAM [48] and Grad-CAM [32] identify the most discriminative spatial region. To boost WSOL performance, [34, 45, 46, 7] propose architectural modifications to augment the most discriminative region with less-discriminative object regions. This is achieved by fine-tuning a pre-trained network while hiding the most discriminative region stochastically. This forces the network to recognize other informative regions and thus improve WSOL. To detect and hide the most discriminative region while fine-tuning, a network is assumed to use global average pooling (GAP) [48, 34, 7] or an equivalent  $1 \times 1$  feature reduction convolution layer [45]. This fine-tuning paradigm tends to degrade classification performance.



Fig. 3. An overview of weakly supervised object localization (WSOL) approaches for classification and retrieval networks. Some approaches impose architectural constraints and require fine-tuning, *e.g.*, CAM and ADL.

## **3** Constrained Attention Filter (CAF)

This section presents two variants for optimizing the proposed L2-CAF. The first variant, *class-oblivious*, works for both classification and retrieval CNNs. It generates a single heatmap per frame. The second variant, *class-specific*, works for classification CNNs and generates class-specific heatmaps per frame. Both variants impose no architectural constraints in terms of spatial pooling (GAP, FCN) or temporal fusing components (RNN, LSTM).

#### 3.1 Class-Oblivious Variant

Given a pre-trained network and an input  $x \in \mathbb{R}^{W \times H \times 3}$ , the last convolution layer provides a feature map  $A \in \mathbb{R}^{w \times h \times k}$ , with size  $w \times h$  and k channels. The network's output NT(x), logits or embedding, depends on discriminative features in A. We optimize an L2 normalized filter f to identify the discriminative features to the network's output NT(x). After multiplying A by the filter  $f(A \odot f = A' \in \mathbb{R}^{w \times h \times k})$ , the network generates a filtered output FT(x, f). While fixing the network's weights and input, we optimize f to minimize

$$L = ||NT(x) - FT(x, f)||^2, \text{ subject to } ||f||_2 = 1,$$
(1)

FT(x, f) equals NT(x) if and only if  $f = f_I = \{1\}^{w \times h}$  which is infeasible due to the unit L2-Norm constraint.

**Intuition:** An ideal heatmap can be regarded as a filter that approximates NT(x) by blocking irrelevant features in A. Accordingly, we seek a filter f that spatially prioritizes convolutional features and flexibly captures irregular (e.g., discontinuous) shapes or multiple different agents in a frame. The L2-Norm, a simple *multi-mode* differentiable filter, satisfies these requirements. On account of irrelevant features, the  $||f||_2 = 1$  constraint assigns higher weights to relevant features. Figure 11 qualitatively emphasizes the intuition behind the L2-Norm constraint.

This formulation (Eq. 1) is oblivious to the nature of the network's output (logits or embedding), architecture, and input format (RGB image or preextracted features). For a given input x, the class-oblivious formulation generates a single heatmap. This can be a limitation if the input x contains objects from different classes. The next subsection tackles this limitation by offering an alternative class-specific optimization formulation.

#### 3.2 Class-Specific Variant

To support class-specific heatmaps per input, we first assume a classification CNN architecture with class-specific logits. We learn the attention for class c by optimizing the L2-CAF f using the following loss function

$$L_{c} = -FT_{c}(x, f) + \sum_{i=0, i \neq c}^{N} FT_{i}(x, f), \text{ subject to } ||f||_{2} = 1,$$
(2)

where  $FT_c(x, f)$  is the filtered output's logit for class c and N is the total number of classes. This loss maximizes the output logit for the intended class c while minimizing the output logits for all other classes.

Figure 1 presents a qualitative comparison between the class-oblivious and class-specific variants. For example, the first example shows an image of a butterfly standing on a mastiff's nose. The first image shows the resulting heatmap from optimizing Eq. 1. The following two images show the result heatmaps from optimizing Eq. 2 for the mastiff and butterfly classes, respectively. In these examples, the L2-CAF is applied to the last convolutional layer.

**Technical Details:** To compute the class-oblivious heatmap for an input x, we utilize gradient descent for l iterations. At iteration i,  $L^i$  is computed using the filter  $\frac{f^i}{||f^i||_2}$ . The filter f is initialized randomly, *i.e.*,  $f^1 \in [0,1]^{w \times h}$ . Gradient descent iteratively updates f to minimize L. We terminate when L converges and remains approximately the same for d iterations. Concretely, we terminate the gradient descent at i = l when  $|L^l - L^{l-d}| < \epsilon$  where  $\epsilon = 10^{-5}$  and d = 50. This constrained minimization formulation is non-convex, so we also impose a maximum number of iterations  $L_{max}$  to avoid oscillating between local minima. After termination, the heatmap is generated by resizing  $\frac{|f|^l}{||f^l||_2}$  to the input's size. The same procedure is used for class-specific heatmaps with  $L_c$  (Eq. 2). For more details, please refer to our released code.

**Timing:** To optimize the small (e.g.,  $7 \times 7$ ) L2-CAF using gradient descent, the vanilla L2-CAF requires multiple feed-forward and backpropagation passes through the network. This is affordable for lightweight networks like MobileNet [16] and GoogLeNet (InceptionV1) [37] but computationally expensive for bulky networks like VGG [33] and DenseNet [17]. We propose a technique to reduce this cost through (1) making a single feed-forward pass through the whole network to compute the network's output at every layer, (2) optimizing the L2-CAF fusing a small subset of layers.



**Fig. 4.** Reduce the computational optimization cost of the L2-CAF f by solving an equivalent sub-problem (blue-dashed). Instead of using the network's endpoints (x, NT(x)), use (V(x), V'(x)) to optimize f.

Figure 4 illustrates this technique. Instead of optimizing the filter f through the network's endpoints (x, NT(x)), it is equivalent to use the outputs of the direct pre and post layers (V, V') to the attention filter. For a given input x, these layers' outputs (V(x), V'(x)) require a single feedforward pass through the whole network. Once computed, the loss function from Eq. 1 becomes

$$L = ||V'(x) - FT(V(x), f)||^2, \text{ subject to } ||f||_2 = 1.$$
(3)

To generate class-specific heatmaps for a classification network, V'(x) must be the network's logits NT(x). Since it is typical to visualize the attention of the last convolution layer, this formulation skips the overhead of a network's trunk and significantly reduces the computational cost. The speed-up of this technique is quantified through an ablation study.

The fast L2-CAF approach is a computationally cheaper alternative to sampling [27, 35] and masking [11, 10] approaches. In addition, the L2-CAF has a smaller set of hyper-parameters. For instance, while both L2-CAF and masking [11, 10] approaches require a stopping criterion for an optimization problem, Fong *et al.* [10] evaluate multiple mask-sizes per image. Furthermore, the fast L2-CAF works on a small subset of network layers, *i.e.*, independent of the network backbone. Thus, it compares favorably for video processing. For 3D volumes (*e.g.*, medical images), our optimization problem remains independent of the network size, *i.e.*,  $\in \mathbb{R}^{7 \times 7 \times 7}$ .

## 4 Experiments

The next two subsections present L2-CAF's quantitative evaluation using classification and retrieval networks, respectively. Then, a recurrent retrieval network qualitatively illustrates L2-CAF's potential for video applications. Finally, we present our ablation studies.

#### 4.1 WSOL Using Classification Networks

The L2-CAF is quantitatively evaluated using WSOL on both *standard* and *finetuned* classification architectures. We leverage the ImageNet validation set [8]

		Classification		Localization		
Method	Backbone	Top $1\uparrow$	Top $5\uparrow$	Top $1\uparrow$	Top $5\uparrow$	
Grad-CAM L2-CAF (ours)	GoogLeNet [37] GoogLeNet [37]	$71.17 \\ 71.17$		44.43 <b>45.48</b>	57.50 <b>59.32</b>	
Grad-CAM L2-CAF (ours)	ResNetV2-50 [14] ResNetV2-50 [14]	$71.51 \\ 71.51$	$rac{86.56}{86.56}$	46.57 <b>48.18</b>	59.96 62.38	
Grad-CAM L2-CAF (ours)	DenseNet-161 [17] DenseNet-161 [17]	$78.20 \\ 78.20$	$91.39 \\ 91.39$	49.28 <b>49.68</b>	<b>66.57</b> 65.28	

 
 Table 1. Classification and localization accuracies on the ImageNet (ILSVRC) validation set using standard architectures – no fine-tuning required.

for evaluation on standard architectures. For fined-tuned architectures, we follow ADL [7] evaluation procedure and utilize both ImageNet [8] and CUB-200-2011 [39] datasets. In all experiments, we use the fast L2-CAF technique. The loss in Eq. 2 is minimized using the last convolution layer and the network's logits (before softmax) as endpoints.

**Evaluation using standard architectures** is performed using both the top-1 and top-5 predictions. Similar to [32], we obtain the top predictions for every image, then, optimize our filter f to learn the corresponding heatmap for every prediction. Following Zhou *et al.* [48], we segment the heatmap using a simple thresholding technique. This generates connected segments of pixels; we draw a bounding box around the largest segment. Localization is correct if the predicted class is correct and the intersection over union (IoU) between the ground truth and estimated bounding boxes is  $\geq 50\%$ . Table 1 compares L2-CAF and Grad-CAM using three architectures. Both approaches are applied to the last  $7 \times 7$  convolution layer. We fix the architecture and evaluate different localization approaches – same classification but different localization performance.

**L2-CAF versus Grad-CAM:** Grad-CAM is 5 times faster than L2-CAF on DenseNet-161 (7 times on GoogLeNet). Both approaches support a large variety of architectures. In terms of localization accuracy, L2-CAF compares favorably to Grad-CAM. Fong and Vedaldi [11] explain why gradient-based approaches like Grad-CAM are not optimal for visualization. They show that neural networks' gradients  $\frac{\partial y}{\partial x}$  are independent of the input image x for linear classifiers ( $y = wx + b; \frac{\partial y}{\partial x} = w$ ). For non-linear architectures, this problem is reduced but not eliminated. They also show qualitatively that gradient saliency maps contain strong responses in irrelevant image regions. We hypothesize that DenseNet-161's better classification accuracy and, accordingly, better gradient closes the localization performance gap.

**Evaluation using fine-tuned architectures** is performed using the top-1 accuracy on fine-tuned architectures (*e.g.*, VGG-GAP [48]); this follows the evaluation procedure in attention-based dropout layer (ADL) [7]. ADL is the current state-of-the-art method for WSOL. During fine-tuning, ADL applies dropout at multiple network stages. It is not straightforward to determine where to plug these extra dropout layers – it is network dependent. Therefore, we leverage

			CUB-200-2011	ImageNet
Method	Backbone	Tuning	$\overline{\mathrm{CLS} \uparrow \mathrm{LOC} \uparrow}$	$\mathrm{CLS} \uparrow \mathrm{LOC} \uparrow$
CAM L2-CAF (ours) Fuse 2 CAMs CAM L2-CAF (ours)	VGG-GAP VGG-GAP VGG-GAP VGG-GAP VGG-GAP	GAP GAP ACoL [45] ADL ADL	$\begin{array}{cccccccccccccccccccccccccccccccccccc$	69.9643.4669.9644.0967.50 <b>45.83*</b> 69.5842.9369.5843.27
CAM L2-CAF (ours)	ResNet50-SE ResNet50-SE	ADL ADL	78.94 <b>61.71</b> 78.94 61.16	76.218 49.90 76.218 <b>50.49</b>

**Table 2.** Classification and localization accuracies on the CUB-200-2011 test and ImageNet validation split using fine-tuned architectures. The accuracy with an asterisk<sup>\*</sup> indicates that the score is from the original paper.

their publicly released VGG-GAP and ResNet50-SE implementations to evaluate our approach. The ACoL performance is reported from the original paper.

Table 2 presents a quantitative evaluation using CUB-200-2011 and ImageNet datasets. The first column denotes the object localization approach, *e.g.*, CAM versus L2-CAF. Grad-CAM is dropped because it is equivalent to CAM when a GAP layer is utilized [32]. In the second column (backbone), all the architectures utilize a global average pooling or an equivalent surrogate [45]. The third column denotes the fine-tuning approaches considered: GAP [48], ACol [45], ADL [7]. We fine-tune the VGG-GAP architecture with both GAP and ADL. L2-CAF consistently outperforms CAM's localization on the ImageNet validation set.

Relation with WSOL approaches (e.g., ADL): To generate class activation maps (CAMs), WSOL approaches employ a global average pooling layer (GAP) [48, 34, 46, 7], or equivalent [45]. L2-CAF relaxes this architectural requirement. Thus, while supporting previous WSOL approaches, L2-CAF introduces a new degree of freedom. This flexibility is vital to explore attention visualization and WSOL beyond standard supervised classification networks.

#### 4.2 WSOL Using Retrieval Networks

Weakly supervised object localization provides a quantitative evaluation metric for attention visualization approaches. The ability to localize attention for various architectures is a core advantage of L2-CAF. In this subsection, we quantitatively evaluate L2-CAF against Grad-CAM. We employ the class oblivious formulation (Eq. 1) using the last convolution layer and the raw embedding (before unit-circle normalized) as endpoints.

**Datasets:** We employ CUB-200-2011 birds [39] and Stanford CARS196 [21] retrieval datasets, *i.e.*, standard retrieval datasets [26, 40, 6, 41]. Both datasets provide the ground truth bounding box annotation. They pose several challenges for foreground objects' localization. Birds are not naturally rectangular; discriminative parts (*e.g.*, head [5]) occupy a small part of the body. Cars pose a similar challenge in terms of relatively smaller discriminative parts (*e.g.*, wheel) relative to the whole body. Figure 5 depicts the ratio of the ground truth bounding box to the whole image size for both datasets.



Fig. 5. Histogram of the foreground objects' bounding box size relative to the whole image in CUB and CARS196 datasets. CUB birds tend to occupy less than 50% of the whole image (left-skewed), while the Stanford cars are normally distributed.

**Evaluation metrics:** For retrieval, we utilize both Recall@1 (R@1) and the Normalized Mutual Information (NMI) metrics. For localization, we follow the same evaluation procedure in [48, 32] for classification networks. We replace the top-1 by R@1 metric to decide if the network's output is correct or not. The same IoU > 50% criterion is used to evaluate localization.

Vanilla Grad-CAM baseline: To evaluate L2-CAF quantitatively, we extend the classification Grad-CAM to deal with retrieval networks. The Grad-CAM class-discriminative localization map  $M^c$  has been proposed as follows

$$M^{c} = RELU\left(\sum_{k} \alpha_{k}^{c} A^{k}\right) \tag{4}$$

$$\alpha_k^c = \frac{1}{w \times h} \sum_{i=0}^w \sum_{j=0}^h \frac{\partial y^c}{\partial A_{i,j}^k},\tag{5}$$

where  $M^c \in R^{w \times h}$  for any class  $c, A \in R^{w \times h \times k}$  is a convolutional feature map with k channels.  $\alpha^c \in R^k$  quantifies the  $k^{th}$  channel's importance for a target class c. Basically,  $\sum_k \alpha_k^c A^k$  provides a weighted sum of the feature maps (A) for class c.  $\alpha_k^c$  is computed using the gradient of the score for class  $y^c$  with respect to the feature maps  $A^k$ .

To support a retrieval network, we utilize the gradient of the output embedding  $\frac{\partial y}{\partial A_{i,j}^k}$  instead of the class score  $\frac{\partial y^c}{\partial A_{i,j}^k}$  as follows

$$\alpha_k^y = \frac{1}{w \times h} \sum_{i=0}^w \sum_{j=0}^h \frac{\partial y}{\partial A_{i,j}^k},\tag{6}$$

we denote this formulation as vanilla Grad-CAM for retrieval. We compute  $\frac{\partial y}{\partial A_{i,j}^k}$  using tf.gradients [1].

**Grad-CAM-abs baseline:** The Vanilla Grad-CAM is largely inferior for retrieval networks because of the RELU in Eq. 4. RELU is introduced for classification networks to emphasize feature maps that have a positive influence on

**Table 3.** Triplet (TL) and N-pair (NP) losses' quantitative retrieval evaluation using NMI and Recall@1 on CUB-200-2011 and CARS196. Quantitative localization accuracy evaluation using the 0.5 intersection over union (IoU) criterion.  $\triangle$  column indicates the absolute localization improvement margin relative to the vanilla Grad-CAM.

			CUB-200-2011			CARS196			
			Retrieval Localization		Retrieval		Localization		
Method	Backbone	Loss	$\overline{\text{NMI}\uparrow \text{ R}@1\uparrow}$	$\mathrm{LOC}\uparrow$	Δ	NMI↑	R@1↑	$\mathrm{LOC}\uparrow$	$\triangle$
Grad-CAM Grad-CAM-abs L2-CAF (ours)	GoogLeNet GoogLeNet GoogLeNet	TL TL TL	$\begin{array}{ccccccc} 0.582 & 47.75 \\ 0.582 & 47.75 \\ 0.582 & 47.75 \end{array}$	20.63 22.96 <b>29.63</b>	+2.33 +9.00	$\begin{array}{c} 0.532 \\ 0.532 \\ 0.532 \end{array}$	$54.55 \\ 54.55 \\ 54.55$	32.38 46.21 <b>54.10</b>	+13.82 +21.72
Grad-CAM Grad-CAM-abs L2-CAF (ours)	ResNet ResNet ResNet	TL TL TL	$\begin{array}{cccc} 0.601 & 50.06 \\ 0.601 & 50.06 \\ 0.601 & 50.06 \end{array}$	16.15 29.49 <b>39.28</b>	+13.34 +23.13	$\begin{array}{c} 0.565 \\ 0.565 \\ 0.565 \end{array}$	$\begin{array}{c} 61.55 \\ 61.55 \\ 61.55 \end{array}$	31.77 56.32 <b>61.27</b>	$^{-}$ +24.55 +29.50
Grad-CAM Grad-CAM-abs L2-CAF (ours)	GoogLeNet GoogLeNet GoogLeNet	NP NP NP	$\begin{array}{cccc} 0.583 & 48.95 \\ 0.583 & 48.95 \\ 0.583 & 48.95 \end{array}$	14.13 19.87 <b>30.50</b>	+5.74 +16.37	$\begin{array}{c} 0.597 \\ 0.597 \\ 0.593 \end{array}$	$\begin{array}{c} 65.23 \\ 65.23 \\ 65.23 \end{array}$	28.29 55.01 <b>64.85</b>	+26.72 +36.56
Grad-CAM Grad-CAM-abs L2-CAF (ours)	ResNet ResNet ResNet	NP NP NP	$\begin{array}{cccc} 0.580 & 47.92 \\ 0.580 & 47.92 \\ 0.580 & 47.92 \end{array}$	11.92 26.67 <b>38.69</b>	+14.75 +26.77	$\begin{array}{c} 0.609 \\ 0.609 \\ 0.609 \end{array}$	$\begin{array}{c} 67.61 \\ 67.61 \\ 67.61 \end{array}$	32.42 61.62 <b>67.35</b>	+29.20 +34.93



Fig. 6. Qualitative attention evaluation for different visualization approaches on retrieval networks. Both Grad-CAM variants suffer near images' corners.

the class of interest  $y^c$ , assuming pixels with negative gradient belong to other classes. This assumption is valid for classification but invalid for retrieval. Therefore, we further modify the Grad-CAM formulation by replacing the *RELU* with the absolute function *abs*. This Grad-CAM-abs baseline is defined as follows

$$M_{abs}^{y} = abs\left(\sum_{k} \alpha_{k}^{y} A^{k}\right).$$
<sup>(7)</sup>

Implementation details are reported in the supplementary material.

**Results:** Table 3 presents a quantitative evaluation for both retrieval and localization performance. ResNet-50 has more parameters than GoogLeNet; and is marginally better in terms of retrieval. Generally, N-pair loss outperforms triplet loss. Cars are rectangular and thus simpler than CUB birds for bounding box



Fig. 7. Qualitative localization evaluation on CUB-200-2011 and CARS196 using a retrieval network trained with a triplet loss. The green and blue bounding boxes indicate the ground truth and the L2-CAF bounding boxes, respectively.



**Fig. 8.** A convolution architecture to embed autonomous navigation videos. This network employs a ranking loss to learn a feature embedding and a recurrent layer for temporal modeling. The CNN layer is shared across the three frames. The attention filters  $(f_1, f_2, f_3)$  are used during attention visualization only.

localization. The localization error is highly correlated and upper-bounded by the retrieval performance (R@1). Grad-CAM-abs outperforms the vanilla Grad-CAM for retrieval. L2-CAF brings further localization improvement.

Figure 6 qualitatively compares different localization approaches. We found that feature maps at the images' corners can have a high positive gradient, while the feature maps at the foreground object can have a high negative gradient. It is a common practice to embed images into the unit-circle, *i.e.*, some images are embedded in the negative space. When this happens, the vanilla Grad-CAM ignores the foreground objects. Grad-CAM-abs handles negative gradient better but still suffers around the corners. Grad-CAM inferior behavior around the corner is qualitatively reported in [10]. This undesirable behavior degrades Grad-CAM's WSOL performance. Figure 7 shows a qualitative localization evaluation on both datasets. For CUB-200-2011, our estimated bounding box (blue) tends to be centered around the birds' heads.

#### 4.3 Recurrent Networks' Attention

This subsection illustrates how to visualize attention for temporally fused video frames through the Honda driving dataset (HDD) [29]. HDD is a video dataset for reasoning about drivers' actions (events) like crossing intersections, making left and right turns. A key objective is modeling the subtle intra-action (events) variations without explicit fine-grained labeling. For instance, an autonomous navigation application with a left-turn query video should differentiate smooth left-turns maneuvers from those interrupted by crossing pedestrians. A retrieval network models these intra-action variations through a feature embedding.



Fig. 9. The recurrent network's attention visualization at different time steps using heatmaps. Each row depicts three frames sampled from an action video. Contours highlight regions with higher attention. The network attends to the spatial locations of traffic lights and road signs. This figure is best viewed on a screen (color and zoom).

Figure 8 presents a recurrent retrieval network for video embedding. Given a trimmed video event, three frames are sampled at  $t_1$ ,  $t_2$ , and  $t_3$ . To enable a large training mini-batch for triplet loss, a pre-trained ResNet is employed to extract convolutional features for every frame. The extracted ResNet features are fed into a trainable shallow CNN. The resulting convolutional features are temporally fused using an LSTM [12, 15].

After training, we employ three L2-CAF filters  $(f_1, f_2, \text{ and } f_3)$  to visualize attention, *i.e.*, one filter per frame. These filters are inserted between the shallow CNN and the LSTM layers *during inference only*. To ground attention in each frame, we optimize each filter independently. Concretely, we pass the first frame's features through  $f_1$  and optimize  $f_1$  while feeding the second and third frames' features normally, *i.e.*,  $f_2$  and  $f_3$  are inactive. After  $f_1$  converges, we deactivate it and optimize the next filter  $f_2$  and so on. After optimizing all filters, each filter provides an attention map for the corresponding frame.

Figure 9 presents our qualitative evaluation. In the first row at  $t_3$ , the network attention is drawn to the traffic lights and double yellow lane marks. Similarly, the second row shows attention drawn toward the traffic light at  $t_{1,2}$ . The final row shows an interesting case at  $t_3$  where the attention is drawn to the stop sign and also to the frame's top center, which is the typical location for a traffic light. Through visualization, we can see that the network uses traffic lights, signs, and road signs as discriminative features. We found that a Mast Arm (L-shaped) traffic light is easier to detect by a neural network compared to a straight pole traffic light. The variable height of a straight pole traffic light poses a challenge for neural networks. For instance, the network attends to the right side of the frame multiple times at different heights in the second row at  $t_{1,2}$ .



**Fig. 10.** Sanity checks [2]. First column visualizes attention using a pretrained network– nothing random. Columns two to five visualize attention when logits and weights (alllayers) are randomized. Different random initializations generate different heatmaps.



Fig. 11. Qualitative evaluation for alternative constraints. Softmax offers a sparse result while the Gaussian filter assumes a single mode. L2-CAF supports multi-mode. The supplementary material shows more vivid visualizations for the L2-CAF, randomly initialized, converging in slow motion.

#### 4.4 Ablation Study

This subsection provides sanity checks for saliency maps [2], then evaluates alternative attention constraints, and finally presents a timing analysis.

**Sanity Checks:** Figure 10 shows how L2-CAF is affected by randomly initialized logits-layer or weights (all layers). Sanity checks [2] emphasize a high dependency between the optimized L2-CAF (heatmap) and the network's weights. **Alternative Attention Constraints:** We qualitatively compare the L2-CAF with both softmax and Gaussian constraints. These are selected for their differentiability, simplicity, and usability in recent literature. Other filtering alternatives (e.g., L1-Norm) are also feasible. Softmax is a typical attention mechanism for image captioning [42, 18] and machine translation [38]. In these problems, the softmax attention module is employed *recurrently* on a single image frame or an input sentence for every output word. This fits the softmax's sparse nature. Gaussian filters have been utilized for temporal action localization [24, 28]. They are denser (more relaxed) compared to softmax but also assume a single mode (elliptical shape). To localize objects in images using a Gaussian constraint, we optimize the filter's mean  $\mu \in \mathbb{R}^2$  while fixing the covariance matrix  $\sigma \in \mathbb{R}^{2\times 2}$ to the identity matrix.  $\sigma$  must be constrained to avoid a degenerate solution

14 A. Taha *et al.* 



Fig. 12. Time analysis for the L2-CAF. The fast L2-CAF brings a significant speed-up while solving the same optimization problem.

where the Gaussian becomes a uniform distribution, *i.e.*,  $\sigma \to \infty$ . All filters are optimized using the class-oblivious formulation (Sec 3.1).

Figure 11 provides a qualitative evaluation using GoogLeNet architecture and three attention constraints. The L2-CAF identifies key region(s) for a network's output with a single glimpse. The filter prioritizes these regions quantitatively. L2-CAF supports a large spectrum of neural networks as a post-training inspection tool. It supports complex architectures including, but not limited to, encoder-decoder [20, 30], generative [13], and U-shaped architectures [31, 22]. It neither undermines the performance nor raises the inference cost. L2-CAF is not the fastest attention visualization approach but is computationally cheap.

**Timing Analysis:** Speed is the main limitation of our iterative formulation. Figure 12 presents a timing analysis for L2-CAF. The y-axis denotes the processing time per frame in seconds. The vanilla L2-CAF uses the default endpoints (x, NT(x)), while the fast L2-CAF uses the output of the last convolution layer and the logits NT(x) as endpoints. The vanilla and fast L2-CAF are equivalent optimization problems but the fast L2-CAF provides a significant speed-up. The three fully connected layers in VGG, between the last convolution layer and the logits, limit the fast optimization technique. VGG-GAP replaces these fully connected layers with an average pooling layer, so its speed is similar to GoogLeNet. Fast L2-CAF takes  $\approx 0.4$  and 0.3 seconds per frame on VGG-GAP and GoogLeNet, respectively. The DenseNet-161's speed-up is maximum because the fast L2-CAF skips all dense blocks before the last convolution layer.

### 5 Conclusion

We have introduced the unit L2-Norm constrained attention filter (L2-CAF) as a visualization tool that works for a large spectrum of neural networks. L2-CAF neither requires fine-tuning nor imposes architectural constraints. Weakly supervised object localization is utilized for quantitative evaluation. State-of-the-art results are achieved on both standard and fine-tuned classification architectures. For retrieval networks, L2-CAF significantly outperforms Grad-CAM baselines. Ablation studies highlight L2-CAF's superiority to alternative constraints and analyze L2-CAF's computational cost.

Acknowledgments: This work was partially funded by independent grants from Office of Naval Research (N000141612713) and Facebook AI.

# References

- 1. tf.gradients, https://www.tensorflow.org/api\_docs/python/tf/gradients
- Adebayo, J., Gilmer, J., Muelly, M., Goodfellow, I., Hardt, M., Kim, B.: Sanity checks for saliency maps. In: NIPS (2018)
- 3. Bucher, M., Herbin, S., Jurie, F.: Improving semantic embedding consistency by metric learning for zero-shot classification. In: ECCV (2016)
- Cao, Y., Long, M., Wang, J., Zhu, H., Wen, Q.: Deep quantization network for efficient image retrieval. In: AAAI (2016)
- 5. Chen, C., Li, O., Tao, C., Barnett, A.J., Su, J., Rudin, C.: This looks like that: deep learning for interpretable image recognition. NeurIPS (2018)
- Chen, W., Chen, X., Zhang, J., Huang, K.: Beyond triplet loss: a deep quadruplet network for person re-identification. In: CVPR (2017)
- 7. Choe, J., Shim, H.: Attention-based dropout layer for weakly supervised object localization. In: CVPR (2019)
- Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: CVPR (2009)
- Eghbali, S., Tahvildari, L.: Deep spherical quantization for image search. In: CVPR (2019)
- 10. Fong, R., Patrick, M., Vedaldi, A.: Understanding deep networks via extremal perturbations and smooth masks. In: ICCV (2019)
- 11. Fong, R.C., Vedaldi, A.: Interpretable explanations of black boxes by meaningful perturbation. In: ICCV (2017)
- 12. Funahashi, K.i., Nakamura, Y.: Approximation of dynamical systems by continuous time recurrent neural networks. Neural networks (1993)
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In: NIPS (2014)
- He, K., Zhang, X., Ren, S., Sun, J.: Identity mappings in deep residual networks. In: ECCV (2016)
- Hochreiter, S., Schmidhuber, J.: Long short-term memory. Neural computation (1997)
- Howard, A.G., Zhu, M., Chen, B., Kalenichenko, D., Wang, W., Weyand, T., Andreetto, M., Adam, H.: Mobilenets: Efficient convolutional neural networks for mobile vision applications. arXiv preprint arXiv:1704.04861 (2017)
- Huang, G., Liu, Z., Van Der Maaten, L., Weinberger, K.Q.: Densely connected convolutional networks. In: CVPR (2017)
- Kazemi, V., Elqursh, A.: Show, ask, attend, and answer: A strong baseline for visual question answering. arXiv preprint arXiv:1704.03162 (2017)
- 19. Kim, S., Seo, M., Laptev, I., Cho, M., Kwak, S.: Deep metric learning beyond binary supervision. In: CVPR (2019)
- Kingma, D.P., Welling, M.: Auto-encoding variational bayes. arXiv preprint arXiv:1312.6114 (2013)
- Krause, J., Stark, M., Deng, J., Fei-Fei, L.: 3d object representations for finegrained categorization. In: ICCVW (2013)
- Li, J., Lo, P., Taha, A., Wu, H., Zhao, T.: Segmentation of renal structures for image-guided surgery. In: MICCAI (2018)
- 23. Li, Q., Sun, Z., He, R., Tan, T.: Deep supervised discrete hashing. In: NIPS (2017)
- 24. Long, F., Yao, T., Qiu, Z., Tian, X., Luo, J., Mei, T.: Gaussian temporal awareness networks for action localization. In: CVPR (2019)

- 16 A. Taha et al.
- Oh Song, H., Jegelka, S., Rathod, V., Murphy, K.: Deep metric learning via facility location. In: CVPR (2017)
- Oh Song, H., Xiang, Y., Jegelka, S., Savarese, S.: Deep metric learning via lifted structured feature embedding. In: CVPR (2016)
- 27. Petsiuk, V., Das, A., Saenko, K.: Rise: Randomized input sampling for explanation of black-box models. arXiv preprint arXiv:1806.07421 (2018)
- Piergiovanni, A., Ryoo, M.S.: Learning latent super-events to detect multiple activities in videos. In: CVPR (2018)
- Ramanishka, V., Chen, Y.T., Misu, T., Saenko, K.: Toward driving scene understanding: A dataset for learning driver behavior and causal reasoning. In: CVPR (2018)
- 30. Rezende, D.J., Mohamed, S., Wierstra, D.: Stochastic backpropagation and approximate inference in deep generative models. PMLR (2014)
- Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: MICCAI (2015)
- Selvaraju, R.R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., Batra, D.: Gradcam: Visual explanations from deep networks via gradient-based localization. In: ICCV (2017)
- Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556 (2014)
- 34. Singh, K.K., Lee, Y.J.: Hide-and-seek: Forcing a network to be meticulous for weakly-supervised object and action localization. In: ICCV (2017)
- Smilkov, D., Thorat, N., Kim, B., Viégas, F., Wattenberg, M.: Smoothgrad: removing noise by adding noise. arXiv preprint arXiv:1706.03825 (2017)
- Sohn, K.: Improved deep metric learning with multi-class n-pair loss objective. In: NIPS (2016)
- Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., Rabinovich, A.: Going deeper with convolutions. In: CVPR (2015)
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. In: NIPS (2017)
- Wah, C., Branson, S., Welinder, P., Perona, P., Belongie, S.: The caltech-ucsd birds-200-2011 dataset (2011)
- Wang, J., Zhou, F., Wen, S., Liu, X., Lin, Y.: Deep metric learning with angular loss. In: ICCV (2017)
- Wu, C.Y., Manmatha, R., Smola, A.J., Krahenbuhl, P.: Sampling matters in deep embedding learning. In: CVPR (2017)
- 42. Xu, K., Ba, J., Kiros, R., Cho, K., Courville, A., Salakhutdinov, R., Zemel, R., Bengio, Y.: Show, attend and tell: Neural image caption generation with visual attention. arXiv preprint arXiv:1502.03044 (2015)
- Yuan, Y., Yang, K., Zhang, C.: Hard-aware deeply cascaded embedding. In: ICCV (2017)
- 44. Zeiler, M.D., Fergus, R.: Visualizing and understanding convolutional networks. In: ECCV (2014)
- 45. Zhang, X., Wei, Y., Feng, J., Yang, Y., S. Huang, T.: Adversarial complementary learning for weakly supervised object localization. In: CVPR (2018)
- 46. Zhang, X., Wei, Y., Kang, G., Yang, Y., Huang, T.: Self-produced guidance for weakly-supervised object localization. In: ECCV (2018)
- Zhang, Z., Saligrama, V.: Zero-shot learning via joint latent similarity embedding. In: CVPR (2016)
- Zhou, B., Khosla, A., Lapedriza, A., Oliva, A., Torralba, A.: Learning deep features for discriminative localization. In: CVPR (2016)