



## Metric-Guided Synthesis for Class Activation Mapping

Downloaded from: <https://research.chalmers.se>, 2025-12-15 18:42 UTC

Citation for the original published paper (version of record):

Luque Cerpa, A., Polgreen, E., Rajan, A. et al (2026). Metric-Guided Synthesis for Class Activation Mapping. Communications in Computer and Information Science, 2578: 351-375.  
[http://dx.doi.org/10.1007/978-3-032-08327-2\\_17](http://dx.doi.org/10.1007/978-3-032-08327-2_17)

N.B. When citing this work, cite the original published paper.



# Metric-Guided Synthesis for Class Activation Mapping

Alejandro Luque-Cerpa<sup>1(✉)</sup>, Elizabeth Polgreen<sup>2</sup>, Ajitha Rajan<sup>2</sup>,  
and Hazem Torfah<sup>1</sup>

<sup>1</sup> Chalmers University of Technology and University of Gothenburg, Gothenburg, Sweden

{luque,hazemto}@chalmers.se

<sup>2</sup> University of Edinburgh, Edinburgh, UK

{elizabeth.polgreen,arajan}@ed.ac.uk

**Abstract.** Class activation mapping (CAM) is a widely adopted class of saliency methods used to explain the behavior of convolutional neural networks (CNNs). These methods generate heatmaps that highlight the parts of the input most relevant to the CNN output. Various CAM methods have been proposed, each distinguished by the expressions used to derive heatmaps. In general, users look for heatmaps with specific properties that reflect different aspects of CNN functionality. These may include similarity to ground truth, robustness, equivariance, and more. Although existing CAM methods implicitly encode some of these properties in their expressions, they do not allow for variability in heatmap generation following the user's intent or domain knowledge. In this paper, we address this limitation by introducing SyCAM, a metric-based approach for synthesizing CAM expressions. Given a predefined evaluation metric for saliency maps, SyCAM automatically generates CAM expressions optimized for that metric. We specifically explore a syntax-guided synthesis instantiation of SyCAM, where CAM expressions are derived based on predefined syntactic constraints and the given metric. Using several established evaluation metrics, we demonstrate the efficacy and flexibility of our approach in generating targeted heatmaps. We compare SyCAM with other well-known CAM methods on three prominent models: ResNet50, VGG16, and VGG19.

**Keywords:** Explainability · Class activation mappings · Oracle-guided inductive synthesis

## 1 Introduction

Convolutional Neural Networks (CNNs) have enabled the development of efficient solutions for a wide range of challenging vision problems, such as object

This work was partly supported by the Wallenberg AI, Autonomous Systems and Software Program funded by the Knut and Alice Wallenberg Foundation, a Royal Academy of Engineering Research Fellowship and a Royal Society Industry Fellowship.

© The Author(s) 2026

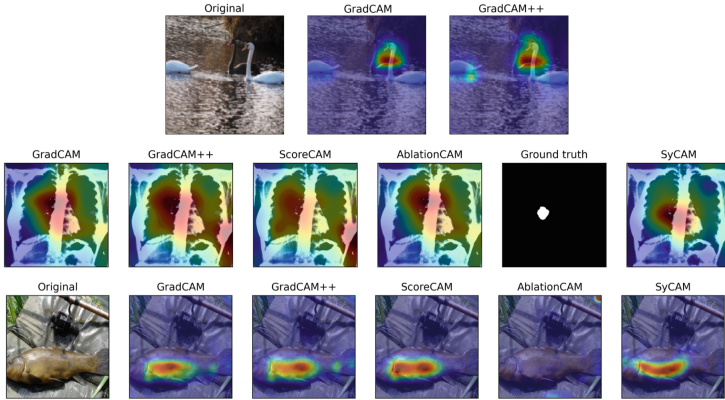
R. Guidotti et al. (Eds.): xAI 2025, CCIS 2578, pp. 351–375, 2026.

[https://doi.org/10.1007/978-3-032-08327-2\\_17](https://doi.org/10.1007/978-3-032-08327-2_17)

detection in autonomous vehicles [11], facial recognition through semantic segmentation [19], and medical image analysis [18, 28]. Despite their advantages, CNNs, like other neural models, suffer from an opaque decision-making process, making it challenging to build trust in their predictions. For instance, physicians using a CNN for X-ray classification require more than just a diagnosis; they need to understand which specific part of the X-ray led to that diagnosis. Similarly, in autonomous vehicles, debugging and employing CNN-based object detectors requires insight into which parts of an image triggered a classification. The growing need to explain CNN behavior has led to the development of various explainability techniques [4, 6, 32, 36]. However, with each new method, it is becoming increasingly clear that more systematic approaches are needed to generate explanations that adapt to the specific intents and needs of end users [14].

In this paper, we address the challenge of incorporating intent by proposing a metric-based approach to generating explanations for CNNs, i.e., where explanations are optimized for a predefined metric. Specifically, we study this problem in the context of explainability methods based on *class activation mappings* (CAM) [37]. CAM methods are one of the most adopted methods for generating saliency maps, i.e., heatmaps that highlight the regions of an input most relevant to the CNN’s prediction. The definition of relevance differs from one method to another, and thus, each method may result in different heatmaps. Figure 1, shows example heatmaps generated by a set of different CAM methods. In the figure, the heatmaps highlight slightly different image regions across the different techniques. Heatmaps produced by CAM are the result of computing a linear combination of the feature maps from the convolutional layers of a CNN. CAM methods generate heatmaps by weighting feature maps based on their contribution to the class score. They differ in how they calculate the weights, each offering a unique expression for computing them. Selecting the most appropriate CAM method depends on the specific application and the level of detail in the activation map.

For example, GradCAM [32] tends to favor larger activation regions in the activation maps, highlighting the most prominent part of an image that influences classification. In an image with multiple swans (see Fig. 1, first row), GradCAM primarily highlights the two more visible swans on the right. However, if these swans are removed, the model may still classify the image correctly based on the partial swan. In contrast, GradCAM++ [4] incorporates all activation regions, highlighting also the partial swans as influential in the CNN’s decision. Depending on the user’s intent, whether they seek to understand the model’s overall behavior or identify the most influential part of a specific image, they may prefer one method over the other. Choosing the right CAM method is, therefore, often a complex task, traditionally relying on human intuition and empirical experimentation. A non-expert user may not know which CAM expression is most suitable for their application, and a systematic approach to guide the generation of optimal CAM expressions, with respect to certain predefined metrics, is missing.



**Fig. 1.** Saliency maps generated using different CAM methods (GradCAM [32], GradCAM++ [4], ScoreCAM [36], AblationCAM [6], and one using an expression synthesized by our SYCAM framework) for three models different CNNs trained on three different data sets ImageNet [5], COVID-QU-Ex [34] and ImageNette [13]. The first row of images shows heatmaps for GradCAM and GradCAM++. The second row of images shows how SYCAM guided by a ground truth metric, captures the ground truth more accurately than the other methods. The last row shows how SYCAM guided by the insertion metric generates a heatmap that closely mimics that of the dominant CAM method, ScoreCAM in this case.

Our framework, SYCAM, presents a metric-based framework for the automatic synthesis of CAM expressions. From a given class of expressions, SYCAM can synthesize expressions tailored to specific properties captured via a fitting evaluation metric. For instance, metrics capturing the overlap between a heatmap and the ground truth mask [33] or the pixel intensity of the heatmap in such overlap [9] guide towards the generation of CAM expressions that generate heatmaps with higher similarity to ground truth. Metrics like the Deletion or Insertion metrics [27] lead to CAM expressions that generate heatmaps highlighting what the model is truly paying attention to (independent of the correctness of the prediction).

Consider the heatmaps shown in the second row in Fig. 1, where different saliency methods have been applied to explain why a CNN classified an X-ray image as a COVID case. Here, we used SYCAM to synthesize an expression optimized toward the ground truth, i.e., during synthesis, expressions are evaluated based on whether they highlight pixels that also match the ground truth. From the figure, we can see that the heatmap generated according to the SYCAM synthesized expression captured the ground truth more accurately. In contrast, all other methods failed to do so, producing significantly less accurate representations of the ground truth region.

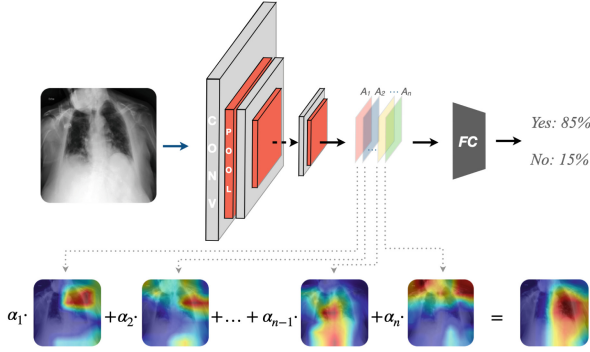
We applied SYCAM using different metrics to study their impact on heatmap generation. In Fig. 1, in the last row, we present another example where SYCAM

was applied using the insertion metric [27]. The insertion metric evaluates how the highest-scoring pixels in the heatmap influence image classification. In this case, ScoreCAM achieved the highest score compared to GradCAM, GradCAM++, and AblationCAM. SYCAM generated an expression that closely resembled ScoreCAM and achieved a very similar score. This demonstrates that when a particular method outperforms others based on a given evaluation metric, SYCAM produces expressions that align with the dominant approach.

Our results are based on an instantiation of SYCAM, adapting techniques from the syntax-guided synthesis literature (SyGuS) [2], which searches a space defined by a grammar of potential CAM expressions. This is done via an *oracle-guided inductive synthesis* (OGIS) approach [15]. In OGIS, a learner explores the space of possible solutions, guided by oracles that give feedback and evaluate the correctness of solutions generated by the learner. Here, the learner is a synthesis process that searches the space of possible CAM expressions, guided by two oracles: one that can remove equivalent solutions and another that evaluates the candidate solution according to a given evaluation metric. We show that, while a monolithic enumerative approach can be used to synthesize CAM expressions, it comes with the limitation of not taking into account any of the image properties, such as the image class. To overcome this limitation, we present an adaptation of the synthesis approach leveraging a class-based decomposition of the problem. Specifically, our approach allows for case splits, thereby synthesizing a set of CAM expressions for each class of properties. This is showcased in our experiments, where we use SYCAM to synthesize CAM expressions for three prominent models: ResNet50, VGG16, and VGG19 [22], and compare them with those of established methods like GradCAM, GradCAM++, ScoreCAM, and AblationCAM.

In summary, SYCAM provides a general framework for the systematic generation of CAM methods, with flexibility in two dimensions: first, the user can provide an evaluation metric to suit their use case; and second, the user can provide a syntactic template or grammar that defines the space of possible expressions, giving the potential to provide user intuition to the search algorithm. SYCAM will then find an expression that is dominant with respect to the given evaluation metric. Our primary contributions can be summarized as follows:

- We introduce the problem of synthesizing CAM expressions for CNNs optimized for specific evaluation metrics.
- We present a framework for solving the problem following the OGIS approach, adapting enumerative techniques from SyGuS.
- We present a thorough experimental evaluation demonstrating the efficacy of our framework in synthesizing CAM-expressions for the prominent classification models ResNet50, VGG16, and VGG19.
- We evaluate our framework over a COVID-19 benchmark, where the expressions incorporate metrics that favor expert knowledge.



**Fig. 2.** Overview of a CNN-based model that classifies X-ray images into COVID-19 positive or negative, and a CAM-based method that explains each classification.

## 2 Background

### 2.1 Convolutional Neural Networks

Convolutional Neural Networks (CNNs) have proven highly effective in pattern recognition tasks, especially in image processing [11, 17, 19]. The typical architecture of a CNN is illustrated in Fig. 2. The fundamental building blocks of a CNN are the convolutional layers, which are used for identifying features within an image. These layers work by iteratively applying filters, so-called kernels, to the input image. These filters are learned during the training phase. Each layer applies multiple filters to the output of the previous layer, starting with the input image. The output produced by applying a filter is known as a feature map. As the network deepens, the feature maps define increasingly complex features. While the initial layers may focus on basic elements, such as colors and edges, deeper layers recognize larger patterns. In classification tasks, these higher-level features are usually passed to a fully connected layer used for classifying the image. To reduce the size of these feature maps while preserving critical information, pooling layers are introduced after each convolutional layer. Common pooling methods include average pooling and max pooling, which reduce a portion of a feature map into its average or maximum value, respectively.

### 2.2 Class Activation Mapping

Saliency maps are visual artifacts that highlight the regions of an image most relevant to a model’s prediction. Class Activation Mapping (CAM) is a prominent set of methods for generating these saliency maps. The core idea of CAM is to compute a saliency map, the class activation map, via a weighted linear combination of the feature maps of the last convolutional layer [37]. This layer is chosen because it captures high-level features that are most relevant for the classification. The map is always generated for a certain class and points to the parts of an image the network focuses on when predicting that particular class.

The general concept of CAM is illustrated in Fig. 2. Formally, for an output class  $c$  and a convolutional layer  $l$ , the saliency map  $L^c$  can be defined as:  $L^c = \sum_k \alpha_k^c A_k^l$ , where  $A_k^l$  denotes the  $k$ -th feature map of the layer  $l$ , i.e., the two-dimensional array resulting from applying the  $k$ -th convolutional filter of the layer  $l$  to its input, and the weights  $\alpha_k^c$  depend on the CAM method chosen.

Each feature map highlights specific aspects of the image. High values in a feature map point to the presence of a feature; low values hint at its absence. Since each feature map contributes differently to the scores of individual classes, the linear combination is weighted to reflect the influence of each feature map in producing the final class activation map. These influences are represented by the weights  $\alpha_1, \dots, \alpha_n$ .

Different CAM methods are distinguished by how they calculate the weights. In the original CAM, the weights were defined via the learned weights corresponding to different classes. Methods like GradCAM or GradCAM++ [21] compute the weights by using the gradients of the score for a specific class with respect to the feature maps of a convolutional layer. Further gradient-based approaches building on the latter two include SMOOTH GRADCAM++ [26], AUGMENTED GRADCAM [24], and XGRADCAM [10]. Gradient-free approaches determine the importance of different regions in the input image for a specific class without relying on gradients. These include perturbation-based approaches like SCORECAM [36] and ABLATIONCAM [6], attention-based methods like Attention-Guided CAM [20], and methods like EIGENCAM [25], that apply principal component analysis to create the class activation map. For a comprehensive survey, we refer the reader to [14]. So far, no systematic approach has been used to determine the optimal weight expression for a given task. Our work addresses this challenge by introducing an automated approach capable of synthesizing both gradient-based and gradient-free expressions if the grammar allows it.

### 2.3 Heatmaps Evaluation Metrics

Several metrics have been introduced in the literature to evaluate saliency methods [16]. Some of these methods are perturbation-based, i.e., they evaluate the effect of masking regions highlighted by CAM methods on the model's performance. Examples of such metrics include Average Drop % [4], AOPC [31], ROAD [30], IROF [29], and the Deletion and Insertion [27] metrics. Another category of metrics is ground-truth-based. These metrics measure the distance of the explanation to the ground-truth explanation data. Some examples of these metrics are the  $m_{GT}$  metric defined in [33], the Segmentation Content Heatmap (SCH) metric [9], CEM [7], and CLEVR-AI [3]. SYCAM is agnostic to the chosen metric, and CAM-weight expressions can be synthesized by SYCAM with respect to any of the aforementioned metrics. Each evaluation metric focuses on different properties of the heatmaps, e.g., similarity to ground truth, robustness, and more. The synthesized expressions are, in consequence, based on well-founded criteria, guided by quantitative evaluation metrics. The resulting expressions generate

optimized saliency maps according to these metrics, thus allowing us to eliminate any human biases and to select a CAM method more suitable for a given context. An extensive study on the correctness of XAI techniques in generating the true explanations, the so-called fidelity, can be found in [23].

In the following, we provide a detailed introduction to key evaluation metrics that we also use later in our experiments.

*Average Drop %:* The intuition behind this metric is one that checks whether removing parts not highlighted by the heatmap of an image reduces the classification confidence of the model. Removing parts of an image that are not relevant shouldn't heavily impact the confidence drop. To check if the most relevant parts of an image  $i$  are preserved, the product of the heatmap and the image is computed. Then, the resulting image  $h$  is classified to measure the confidence drop. A low confidence drop implies that the heatmap contains the most relevant features of the image, so the lower the metric value, the better.

Given a dataset  $\mathcal{I}$ , the Average Drop % is expressed as  $\sum_{i \in \mathcal{I}} \frac{\max(0, y_i^c - h_i^c)}{y_i^c} \cdot 100$ , where  $y_i^c$  is the classification score for image  $i \in \mathcal{I}$  and class  $c$ , and  $h_i^c$  is the classification score for the product of image  $i$  and the generated heatmap for such image, and class  $c$ .

*Deletion and Insertion Metrics:* The Deletion metric measures the drop in the classification score when the most relevant pixels of the image are *gradually* removed, while the Insertion metric measures the rise in the classification score when they are iteratively added to a blank image. In this paper, we use modified versions of the Deletion and the Insertion metrics. Given a model  $M$ , an image  $x$ , a limit on the number of perturbations  $P$ , and a saliency map  $\mathcal{H}$  over  $x$ , the process recursively modifies the image according to the following formula:

$$x^{(0)} = x \quad \forall 1 \leq j \leq P : x^{(j)} = g(x^{(j-1)}, \bar{x}, r_j)$$

For the Deletion metric, the function  $g$  gradually replaces the most relevant pixels  $r_j$ , ordered by the saliency map  $L_M^c[e]$ , and their neighborhoods with the corresponding parts of a highly blurred version of  $x$ , denoted by  $\bar{x}$ . The neighborhood of a pixel is given by the feature maps used to compute the saliency map. If the feature maps  $A_k$  have size  $w \times h$ , we split the image into a grid of  $w \times h$  neighborhoods. For the Insertion metric, the process is analogous, but the most important pixels are added to the highly blurred version of the image instead. The metrics are then defined by:

$$\mu_{deletion} = \frac{1}{|I|} \sum_{x \in I} \frac{1}{P+1} \left( \sum_{j=0}^P M(x^{(0)})^{c_0} - M(x^{(j)})^{c_0} \right)$$

$$\mu_{insertion} = \frac{1}{|I|} \sum_{x \in I} \frac{1}{P+1} \left( \sum_{j=0}^P M(x^{(j)})^{c_0} - M(x^{(0)})^{c_0} \right)$$



where  $c_0$  is the output class of  $M(x^{(0)})$ ,  $M(x)^{c_0}$  is the output score for class  $c_0$ , and  $|I|$  is the number of images in the set  $I$ . Unlike for other similar metrics [12], no artifacts are added through this procedure, and  $\bar{x}$  is generated using deterministic filters, so the influence of random perturbations over the metric is also avoided. The purpose of using these versions instead of the originals [27] is to avoid adding artifacts, to make sure that the computed scores are relative to the initial score of the model  $M$  over the base image, and to make them work so that the higher the values, the better.

Intuitively, higher values of the metrics imply higher variations in the classification score during the perturbation process, i.e., higher relevance of the pixels highlighted by the saliency map. We may get some insights about the datasets employed using these metrics. For example, when the confidence drop for the Deletion metric is much lower for a subset than for the others, it means that the model is still confident in its classification even after masking portions of the objects. This may imply that the objects occupy most of the image or that the model identifies the objects through the background. In such cases, it is important to improve the dataset.

*Ground Truth Similarity Metrics:* Some evaluation metrics are focused on measuring the similarity of the heatmap to a ground-truth mask.

Given a heatmap  $H$  and a ground-truth mask  $H_{GT}$  with  $p$  pixels, the  $m_{GT}$  metric takes the  $p$  most relevant pixels of  $H$  and counts how many of those pixels are part of the ground-truth mask. This is, if  $n$  of the  $p$  most relevant pixels of  $H$  are part of the mask  $H_{GT}$ , then  $m_{GT}(H, H_{GT}) = n/p$ . While the  $m_{GT}$  metric measures the similarity between a heatmap and the ground truth, it doesn't provide information about the intensity of the heatmap pixels. The SCH metric, on the other hand, solves this problem. The SCH metric is given by:

$$SCH(H, H_{GT}) = \frac{\sum_{i,j} H_{i,j} \cdot M_{i,j}}{\sum_{i,j} H_{i,j}}$$

Intuitively, the  $SCH$  metric measures how concentrated the heatmap is in the ground-truth mask. The more concentrated the pixels are and the higher their relevance, the higher the metric will be, so the higher, the better.

Notice that a strong assumption of ground-truth similarity metrics is that the models are making classifications based on the ground-truth part of the image. If the models are not doing so, and are making decisions based on other information found in the input data, these metrics would return low values independently of how well the heatmaps are explaining what the models are paying attention to.

## 2.4 Syntax-Guided Synthesis

Syntax-Guided Synthesis (SyGuS) [2] is the problem of generating a function that both satisfies a semantic specification and is contained within a language described in a context-free grammar (CFG).

**Definition 1 (Context-Free Grammar).** A context-free grammar  $G$  is defined as a set of terminal symbols  $T$  and a set of nonterminal symbols  $NT$ , a start symbol  $S$ , and a set of production rules  $R \subseteq NT \times (T \cup NT)^*$  which describe the way programs may be constructed iteratively from the grammar.

**Definition 2 (SyGuS problems).** Formally, a SyGuS problem is a 4-tuple  $\langle T, G, \phi, F \rangle$  such that  $T$  is a first-order theory,  $G$  is a context-free grammar,  $\phi$  is a first-order formula, and  $F$  is a function symbol that may occur in  $\phi$ . A solution to a SyGuS problem  $\langle T, G, \phi, F \rangle$  is a function  $f$  such that  $T \models \phi[F \mapsto f]$  and  $f \in \mathcal{L}(G)$ , where  $\phi[F \mapsto f]$  denotes replacing all occurrences of  $F$  in  $\phi$  with  $f$ .

Later in the paper, we study a SyGuS instantiation of CAM expression synthesis, where the specification is that the resulting CAM method must perform better than a given threshold function. This is described in detail in Sect. 4. The most common approach for solving SyGuS problems is Oracle Guided Inductive Synthesis (OGIS) [15]; a family of algorithms that alternate between a learner, which attempts to learn a solution to the synthesis problem, and an oracle, which guides the learner via means of queries and responses, the simplest of which is a correctness query (the learner asks “is this candidate program correct?”, and the oracle replies with “yes” or “no”). There is a broad variety of learners in the literature, but the most common are enumerative techniques. We take inspiration from some of the most common enumerative learners [1] when implementing our approach.

### 3 CAM Expression Synthesis

#### 3.1 Problem Statement

Let  $t[?]$  define an expression with holes, known as a *template*, where holes in the expression are marked by the symbol  $?$ . For an expression  $e$ , the expression  $t[e]$  results from replacing every appearance of the symbol  $?$  in  $t$  with  $e$ . Following the definition of CAM as given in Sect. 2.2, we define a CAM template as  $L^c[?] = \sum_k ? \cdot A_k^l$ . The  $?$  is a placeholder for an expression that defines how the weights  $\alpha_k^c$  are computed for a class  $c$ . We refer to such expressions as CAM-weight expressions. We refer to the set of all instantiations of  $L^c[?]$  by  $\mathcal{F}_{CAM}$ .

**Problem 1 (CAM expression synthesis)** Let  $\mathcal{M} = (\mathcal{I} \rightarrow \mathbb{R}^{|\mathcal{C}|})$  be a set of CNN-based classifiers defined over a space of images  $\mathcal{I}$  and a set of classes  $\mathcal{C}$ . Given  $M \in \mathcal{M}$ , a set of images  $I \subseteq \mathcal{I}$ , a threshold function  $\lambda: \mathcal{M} \times \mathcal{I} \rightarrow \mathbb{R}$ , a set of CAM-weight expressions  $\mathcal{E}$ , and an evaluation function  $\mu: \mathcal{F}_{CAM} \times \mathcal{M} \times \mathcal{I} \rightarrow \mathbb{R}$ , synthesize an expression  $e \in \mathcal{E}$  s.t.  $\mu(L^c[e], M, I) > \lambda(M, I)$ .

In our problem statement, the role of the threshold function  $\lambda$  is to set a lower bound on the quality of synthesized expressions with respect to the evaluation function  $\mu$ . As we will see in our experiments, the threshold function can be given as a fixed number or as a function of any other CAM function. The evaluation function  $\mu$  defines a metric for evaluating CAM functions and can be realized by implementing known CAM evaluation metrics from the literature.

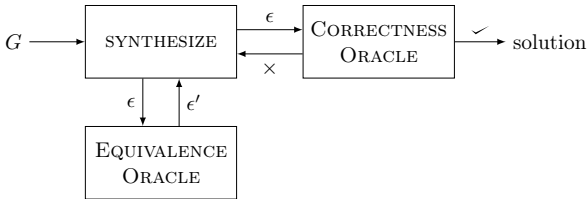
### 3.2 Oracle-Guided Synthesis of CAM-Weight Expressions

Now that we have a formally defined problem statement, we can discuss the SyCAM framework, an oracle-guided synthesis approach for the synthesis of CAM-weight expressions. Let us initially assume we have access to two oracles:

- an equivalence oracle: this oracle receives as input a set of expressions  $\epsilon$ , which may contain semantically equivalent expressions, and returns a set of expressions  $\epsilon' \subseteq \epsilon$  such that no two expressions in  $\epsilon'$  are semantically equivalent; and
- a correctness oracle: this oracle receives a single CAM-weight expression as input, and returns a boolean which is true if the CAM-weight expression results in a saliency map that scores above a pre-defined threshold on a pre-defined metric, for a given set of images and classification model.

We will define the SyCAM framework assuming access to these oracles, but we should note that these oracles are performing tasks that are, in general, undecidable (equivalence checking) or at least computationally expensive (evaluating a given CAM weight expression across a large set of images). We will address the practical implementation of each of these oracles in Sect. 4.

The general workflow of SyCAM is depicted below. SyCAM is composed of two main procedures, the synthesis phase (which is guided by the equivalence oracle) and the evaluation phase (the correctness oracle). Candidate expressions produced by the synthesis procedure are generated from a space of expressions defined by an input grammar  $G$ , as described in Sect. 4. The expressions are then forwarded one by one to the evaluation process. An expression is evaluated using a correctness oracle, defined in terms of the given evaluation metric  $\mu$  and a threshold function  $r$ , over a set of images  $I$ . If a candidate passes the evaluation process, it is returned as a solution to the overall synthesis process. If the evaluation fails, another expression is selected out of the current list of candidate expressions, and the same evaluation process is repeated for the new expression. If all candidate expressions have been evaluated with no success, the synthesizer is triggered again to generate a new set of candidates.



*Synthesis Phase:* The task of the synthesis phase is to enumerate expressions using the feedback given by the oracles. The algorithm we use is based on a classical program synthesis technique: the bottom-up search algorithm [1]. For the algorithm, we require an input grammar  $G$ , which defines an initial set of expressions (*terminals*) and production rules  $R$  that allow us to combine

**Algorithm 1.** Bottom-up Search

---

<b>function</b> SEARCH( $G, \lambda, M, I, \mu$ ) $exprs \leftarrow \emptyset$ $exprs' \leftarrow G.Term$ <b>while</b> True <b>do</b> $solution \leftarrow \text{EVAL}(exprs', \lambda, M, I, \mu)$ <b>if</b> $solution \neq \text{nil}$ <b>then</b> <b>return</b> $solution$ $exprs \leftarrow exprs \cup exprs'$ $exprs' \leftarrow \text{EXPAND}(expr, G)$ $exprs' \leftarrow \text{ELIMEQUIV}(exprs, exprs', I)$	<b>function</b> EXPAND( $E, G$ ) $E' \leftarrow \emptyset$ <b>for</b> $r \in G.R$ <b>do</b> <b>if</b> $r$ is 1-ary <b>then</b> <b>for</b> $e_1 \in E$ <b>do</b> $E' \leftarrow E' \cup r\{e_1\}$ <b>else if</b> $r$ is 2-ary <b>then</b> <b>for</b> $e_1, e_2 \in \text{PROD}(E, E)$ <b>do</b> $E' \leftarrow E' \cup r\{e_1, e_2\}$ <b>return</b> $E'$
---	--

---

expressions to synthesize new ones (*synthesize* block). The synthesis phase calls the equivalence oracle, ELIMEQUIV, in order to reduce the exponential growth in expressions. This bottom-up search algorithm is shown in Algorithm 1.

Initially, we populate  $exprs$  with all expressions for  $Term$ . At each iteration of the Algorithm 1, the search process deploys the EXPAND function to iterate through the production rules of the grammar, and generates all possible new expressions that use the elements in  $exprs$  to replace the non-terminals in each production rule and adds them to  $exprs$ . We use  $r\{e_1, \dots, e_n\}$  to indicate the result of taking the rule  $r \in R$  and replacing the first nonterminal symbol occurring in  $r$  with  $e_1$ , the second with  $e_2$ , and so on.

*Oracles:* The equivalence oracle ELIMEQUIV then reduces the set of expressions, by removing all semantically equivalent expressions. At each iteration, the correctness oracle EVAL checks the current list of programs to see if it contains a program that it deems correct (i.e., a program that performs above a given threshold on the evaluation metric). If so, the program is returned as a solution.

## 4 An Instantiation of SyCAM

Our approach is customizable to any grammar. In this section, we give an instantiation of our framework for a grammar defining gradient-based expressions, and expressions based on ScoreCAM and AblationCAM. We use a grammar with non-terminals  $NT = \{Expr, Term, Grads\}$ , and starting category  $S = \{Expr\}$ . The set of production rules  $R$  is defined as follows:

$$\begin{aligned}
 Expr &:= Term \mid Expr + Expr \mid 2 \cdot Expr + Expr \mid Expr \cdot Expr \mid ReLU(Expr) \\
 Term &:= Grads \mid top_5(Grads) \mid top_{10}(Grads) \mid top_{20}(Grads) \mid top_{50}(Grads) \\
 &\mid CICScores \mid AblScores \\
 Grads &:= GP\left(\frac{\partial Y^c}{A_k}\right) \text{ for any } c \in C \text{ and } k \in K.
 \end{aligned}$$

where  $ReLU$  is the element-wise function  $\max\{\cdot, 0\}$ . Inspired by algorithms like GradCAM, we use as terminals the *gradients*, denoted by  $Grads$ , of the score for

the predicted class  $c$  with  $Y^c = \mathcal{M}(x)$  with respect to the feature map activations  $A_k$  of the last convolutional layer, i.e.  $GP(\frac{\partial Y^c}{\partial A_k})$  where  $GP$  denotes the global average pooling operation, for each feature map  $k$  in the set of feature maps  $K$  and for each class  $c$  in the set of classes  $C$ . We also include as terminals the functions  $top_n(Grads)$  that nullify all but the highest  $n$  elements of  $Grads$  and are denoted by  $top_n$ .

Inspired by ScoreCAM, we include as terminals (*CICScores*) the *channel-wise Increase of Confidence (CIC)* [36]. For each feature map  $k$ , each image  $x$ , a baseline input image  $x_b$ , and a model  $\mathcal{M}$ , the CIC is defined as  $CIC(A^k) = \mathcal{M}(x \circ H^k) - \mathcal{M}(x_b)$  where  $H^k = s(Up(A^k))$ ,  $Up$  denotes the upsample operation that upsamples  $A^k$  into the input image size, and  $s(\cdot)$  is a normalization function to the range  $[0, 1]$ . We assume that  $\exists x_b : \mathcal{M}(x_b) = 0$ , and use  $C(A^k) = \mathcal{M}(x \circ H^k)$ .

*AblScores* represent the weights used in the definition of AblationCAM [6]. For each feature map  $k$  of a model  $\mathcal{M}$ , and each image  $x$ , these weights  $w$  are defined by  $w_k^c = \frac{y^c - y_k^c}{y^c}$ , where  $y^c = \mathcal{M}(x)$  and  $y_k^c$  is the result of setting all the activation cell values of  $A_k$  to zero and classifying again the  $x$ .

Notice that the terminals of the grammar are vectors with number of elements equal to feature maps in the last convolutional layer of  $\mathcal{M}$ , and every expression generated by this grammar produces vectors with the same size as the terminals.

#### 4.1 Equivalence Oracle

In general, determining the equivalence between two expressions is undecidable. We thus use an approximation, referred to as observational equivalence.

**Definition 3 (Observational Equivalence).** *Formally, two expression  $e_1$  and  $e_2$  are observational equivalent on a finite set of images  $I$ , according to an evaluation metric  $\mu$  and a model  $M$  iff  $\mu(L^c[e_1], M, I) = \mu(L^c[e_2], M, I)$ .*

If two expressions are observationally equivalent, we can remove one of these expressions from the pool of expressions used to build subsequent programs *provided* the semantics of the program fragment do not depend on context. That is, given two expressions  $e_1$  and  $e_2$ , and a set of images  $I$ , if  $e_1$  and  $e_2$  give the same result on the set of images, then so will  $C[e_1]$  and  $C[e_2]$  where  $C[e_1]$  is a program that uses  $e_1$  as a subexpression, and  $C[e_2]$  is the same program but with  $e_2$  in place of  $e_1$ .

CAM-weight expressions are arithmetic expressions so this property is true, *for the set of images  $I$* . However, since evaluating the expression on all the images in our dataset is time-consuming and impractical, our observational equivalence oracle uses a smaller subset of the full dataset. Thus, if this set is not representative of the full dataset, the observational equivalence oracle may remove expressions that we subsequently may need. As a consequence of this approximation of equivalence, SyCAM may fail to find some possible solutions, but this trade-off is worth it to prune the exponential growth of the search space.

## 4.2 Correctness Oracle

Recall that we wish to synthesize an expression  $e$  such that the following is true;  $\mu(L^c[e], M, I) > \lambda(M, I)$ , for a given  $I$  and  $M$ . Unlike many domains where OGIS is used, it is not possible for us to reason about this expression symbolically using techniques like SMT solvers, since this would require us to reason about the weights of the model and the pixels of each image. We must therefore use testing, executing each generated CAM method on the images in the set  $I$ . Expressions are then checked one by one for their correctness. If an expression passes the correctness test, the overall synthesis process terminates, returning this expression as a solution. As we will see in the experiments, this process is modified to run for a fixed amount of time and then return the best expression synthesized instead.

## 4.3 Class-Based Decomposition

A limitation of the algorithm presented in the previous section is that the grammar we use does not contain any expressions that can perform case splits, e.g., an “if” expression, or any logical expressions that can define when to apply a specific CAM method to a particular image. This means that the function we synthesize is applied uniformly to all images, regardless of the properties of that image. In this section, our goal is to extend this grammar to permit case splits. One obvious way of doing this would be to introduce an “if-then-else” statement into the grammar  $G$ , as well as expressions for identifying features or characteristics of different images. This, however, results in a significantly expanded search space, and an intractable synthesis problem. In this section, thus, we break down the synthesis method into two parts: a *classification* model  $\mathcal{M}$ , and a set of CAM-weight expressions, synthesized using the enumeration approach described previously, that should be applied to each class. Thus, for a given set of classes, the end CAM-weight expression will be an expression in the grammar given by  $G$  extended with the following production rule:

$$Expr := (Y^{c_i} = \max(Y^{c_1}, \dots, Y^{c_n})) ? Expr : Expr$$

where  $Y^{c_i}$  is the confidence score for class  $c_i$ , and  $\{c_1, \dots, c_n\}$  is the set of classes generated by the model  $\mathcal{M}$ . It would be possible to use any classifier in this step, but we take advantage of having a classifier that can choose the expression for each image: the model  $\mathcal{M}$ .

One advantage of this approach is that it prevents the rejection of good expressions that do not perform well only for small subsets of images. However, an important disadvantage is that the algorithm has to be executed once for each output class. If the number of output classes of a model is in the order of thousands, this method requires considerable computing power.

## 5 Experiments

In this section, we present three sets of experiments that show the efficacy of SYCAM:

- **E1 - Sanity check:** We show the ability of SyCAM to synthesize known dominant CAM expressions for a given metric, or better ones if the grammar allows for more expressive expressions. We apply SyCAM to one Pytorch model, VGG16 trained over the PASCAL VOC 2007 dataset [8].
- **E2 - Enumeration vs class-based decomposition:** We compare SyCAM to the classical approaches of GradCAM, GradCAM++, ScoreCAM, and AblationCAM. Here, we particularly compare the enumerative and class-based decomposition approaches and show how the latter improves over the former. SyCAM is applied to three PyTorch models: ResNet50, VGG16, and VGG19, trained over the Imagenette dataset [13].
- **E3 - Incorporating ground truth:** This experiment is an application of SyCAM to show that, in contrast to standard methods, SyCAM allows us to incorporate expert knowledge into the generation of saliency maps, resulting in better saliency maps. SyCAM is applied to a ResNet50 model trained over COVID-19 X-ray images from the COVID-QU-Ex dataset [34].

All experiments were run using an NVIDIA T4 with 16 GB RAM. The code used can be found in <https://github.com/starlab-systems/SyCAM>. The computations were enabled by resources provided by the National Academic Infrastructure for Supercomputing in Sweden (NAISS), partially funded by the Swedish Research Council through grant agreement no. 2022-06725.

## 5.1 Experimental Setup

As established in Sect. 3.1, the goal is to find an expression  $e$  such that  $\mu(L^c[e], M, I) > \lambda(M, I)$ , for a given image set  $I$  and a model  $M$  by following Algorithm 1. We adapt Algorithm 1, to one that instead of stopping the computation the moment we find an expression  $e$  that beats a threshold  $\lambda$ , to one that continues the search, always saving the so-far best expression found, and taking the value  $\mu(L^c[e], M, I)$  for the best expression as the new threshold  $\lambda$ . We start with  $\lambda = 0$ . We let the experiments run for a fixed amount of time and return the best expression synthesized during that time.

*Grammar:* The grammar used for the experiments is defined in Sect. 4. Computing the weights generated by terminals *CICScores* and *AblScores* is computationally expensive. In the case of ResNet50, VGG16, and VGG19, they are required to compute 512 classifications per image. To reduce the computation overhead, for each dataset and evaluation metric, we precompute the weights produced by these terminals beforehand for each image in the dataset.

*Synthesis Phase:* The set of expressions generated grows exponentially. Because we have 7 terminals in the grammar, and we are enumerating solutions that can be produced with the rules defined, more than 1000 expressions were generated in only three applications of the EXPAND function (see Algorithm 1) even after discarding equivalent expressions using observational equivalence (Sect. 4.1). By the fourth application of the EXPAND function, we would generate more than

$10^6$  expressions. To prevent spending too much time just discarding equivalent expressions without evaluating them, instead of generating a whole new set of expressions with EXPAND before evaluating, we generate and yield expressions one by one. Each expression is evaluated just after being generated by the equivalence oracle and, if not discarded, by the correctness oracle.

*Equivalence Oracle:* The function ELIMEQUIV of Algorithm 1 employs a subset with 10 images, one of each class of the dataset, to ensure this subset is as representative as possible. We also tested subsets of size 20 and 30 of the Imagenette dataset, and we confirmed that the number of expressions discarded decreases (by 1.5% and 7.2%, respectively) with the number of images in the subset. However, the expressions synthesized were the same, and there is a trade-off with the computation time dedicated to generating expressions and discarding those equivalents, so we maintained the original size of 10 images.

*Correctness Oracle:* It would be time-consuming to test every generated CAM method on every image in the dataset  $I$ . For example, it takes around 18 min to compute the Deletion metric for ResNet50 and a single expression over a dataset of 4000 images. To overcome this computation overhead, we implement the correctness oracle as follows, aiming to discard expressions as early as possible.

The evaluation is done by applying an evaluation procedure on the candidate expression defined in terms of an evaluation metric  $\mu$  and over increasingly large sets of images  $I_1 \subset I_2 \subset \dots \subset I$ . If a candidate expression  $e$  is evaluated over the set  $I_i$  and results in a score larger than that computed by a threshold function  $\lambda$  over at least half of the images in  $I_i$ , and the average score is better, i.e.  $\mu(L^c[e], M, I_i) > \lambda(M, I_i)$ , it is then evaluated over the next largest set of images,  $I_{i+1}$ . Otherwise, the candidate is discarded. If a candidate expression  $e$  is evaluated on the set  $I$  and achieves a score higher than that computed by the threshold function over at least half of the images, and the average score is higher, then  $e$  is the best expression found, and  $\lambda$  is updated with the new threshold defined by  $e$ . With this multi-layered approach, we can quickly eliminate programs that already fail evaluation on smaller sets of data, thus speeding up the search for solutions. It is possible that a candidate may perform poorly on a subset  $I_i$  and then perform better than the threshold function on a subset  $I_{i+1}$ , and thus be discarded early. This risk is small since our subsets are relatively large and are uniformly sampled from the full set of images, so likely to be representative.

The oracle checks if the candidate expression  $e$  is better than the best solution found for at least half of the images because the goal is to find an expression that works well for as many images as possible, i.e., that generalizes well, and at the same time has a higher score.

Lastly, we set a timeout of 6 h for experiment E1 and 24 h for E2 and E3, i.e., the returned expression is the best expression found within these time bounds. We note that within this timeout, we were able to reach the fourth iteration of the algorithm, but despite the optimizations, we are not able to cover all the  $10^6$  expressions of that iteration. This results in a limitation on the size of the expressions that can be synthesized. However, we show later that SYCAM



manages to synthesize better expressions than widely known CAM methods. Furthermore, once the SYCAM expression is synthesized, the computational cost of generating saliency maps is similar to that of other CAM expressions, so it is reasonable to spend time on generating the best possible expression. The cost depends on the terminals used. For example, the cost of generating saliency maps for an expression that computes the *CICScores* as part of the expression is similar to the cost of ScoreCAM.

## 5.2 Experiments E1: Sanity Check

*“Show the ability of SYCAM to synthesize known CAM expressions if they are dominant for a given metric, or even better ones if the grammar allows for more expressive expressions”*

In this experiment, we applied SYCAM to a VGG16 model trained over the training subset of the PASCAL VOC 2007 dataset, and using the Average Drop % metric: one of the settings described in the original GradCAM++ paper [4]. The dataset  $I$  contains 2510 images distributed in 20 classes. For the correctness oracle, we use three subsets of images,  $I \supset I_2 \supset I_1$ . The set  $I_2$  is a subset of 1000 randomly chosen images, and  $I_1$  contains 100 randomly chosen images.

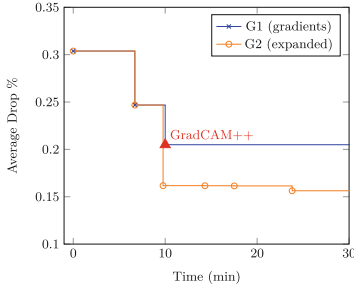
We applied SYCAM using two grammars for a fixed time of 6 h each:

- G1: A grammar defined by:

$$\begin{aligned} Expr &:= Term \mid Expr + Expr \mid 2 \cdot Expr + Expr \mid Expr \cdot Expr \mid ReLU(Expr) \\ Term &:= Grads \mid top_5(Grads) \mid top_{10}(Grads) \mid top_{20}(Grads) \mid top_{50}(Grads) \\ Grads &:= GP\left(\frac{\partial Y^c}{A_k}\right) \text{ for any } c \in C \text{ and } k \in K. \end{aligned}$$

- G2: The grammar G1 adding the terminals *CICScores* and *AblScores*. This is, the grammar described in Sect. 4.

Using grammar G1, SYCAM can only synthesize expressions that employ gradients, so SYCAM can’t synthesize ScoreCAM, AblationCAM, or other similar expressions. By incorporating the *CICScores* and *AblScores* terminals into G2, SYCAM can generate more diverse expressions. Figure 3 illustrates the results obtained by SYCAM using both G1 and G2. When G1 was used, SYCAM managed to synthesize GradCAM++, represented by  $ReLU(Grads)$  [21], guided by the Average Drop % metric, and did not synthesize a better expression before the timeout. However, when G2 was used, SYCAM was able to synthesize even better expressions, with a reduction of the Average Drop % metric of 5%. This shows that SYCAM can synthesize the dominant known expression, or even better expressions if there is a possibility of doing so, and the grammar allows it.



Grammar G1		Grammar G2	
SyCAM exp.	Avg.Drop	SyCAM exp.	Avg.Drop
GradCAM	0.3038	GradCAM	0.3038
$top_{50}$	0.2468	$top_{50}$	0.2468
<b>GradCAM++</b>	<b>0.2049</b>	ScoreCAM	0.1617
		$Grads + CISCores$	0.1616
		$2 \cdot Grads + CISCores$	0.1614
		<b>CISCores + AblScores</b>	<b>0.1563</b>

**Fig. 3.** SyCAM application to a VGG16 model trained over the PASCAL VOC 2007 dataset and the Average Drop % metric (lower is better). If only gradients-related terminals are included in the grammar, SyCAM synthesizes GradCAM++. Better CAM expressions are synthesized for an expanded grammar.

### 5.3 Experiments E2: Enumerate & Class-Based Decomposition

*“Show the efficacy of SyCAM, and that the class-based decomposition approach can perform better than the enumerate approach.”*

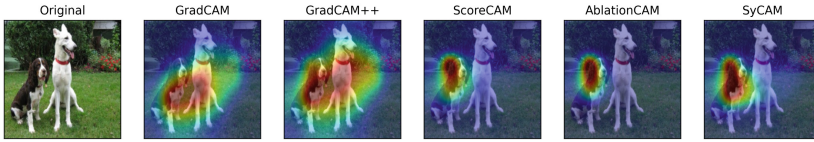
We use this set of experiments to show the efficacy of SyCAM in comparison to other well-known CAM methods: GradCAM, GradCAM++, ScoreCAM, and AblationCAM. We also show how the class-based decomposition allows us to obtain better results than the simple enumerative approach.

In general, whilst the average value-wise improvement in the score obtained by SyCAM in each experiment may look only marginally better than the base methods, this marginally better score already results in more targeted saliency maps for a significant number of images. In Fig. 4, while the SyCAM score shows minor improvement, SyCAM gets a better saliency map that does not highlight the right dog as GradCAM and GradCAM++ do and highlights the body of the left dog more than ScoreCAM and AblationCAM. This implies that the model is making the classification of the image mostly based on the left dog, specifically the head and the upper body.

In the following, we give more details about our findings for both the enumerative and class-based decomposition approaches.

*Results for the Enumerative Approach.* We use a reduced version of *ImageNet*, namely the *Imagenette* dataset [13] that includes images for 10 classes out of the 1000 classes of *ImageNet*. Specifically, we use the validation dataset, with 3925 images distributed evenly into the 10 output classes. The evaluation metric in all the experiments is the Deletion or the Insertion metric. The experiments run for a fixed time of 24 h. Then, they return the best expression found. For the correctness oracle, we use three subsets of images,  $I \supset I_2 \supset I_1$ . The set  $I_2$  is a subset of 1000 images: 100 from each class.  $I_1$  contains 100 images: 10 per class.

As explained in Sect. 5.1, the evaluation of each candidate over the whole dataset can lead to scalability issues. Around 1020 candidate expressions are



**Fig. 4.** Saliency maps generated by GradCAM, GradCAM++, ScoreCAM, AblationCAM, and the SyCAM expression for ResNet50, the class “2. English springer”, and the Deletion metric ( $P = 30$ , higher is better). The scores for each method are 0.2141, 0.2414, 0.2419, 0.2343, and 0.2441, respectively. SyCAM gets a better score and a saliency map that does not highlight the right dog as much as GradCAM and GradCAM++ do and highlights the body of the left dog more than ScoreCAM and AblationCAM.

generated in three iterations of Algorithm 1, and it takes  $\sim 0.3$ s to evaluate the Deletion or the Insertion metric ( $P = 10$ ) over a single image. The computation time of the evaluation of a candidate expression over each subset for the Deletion and Insertion metrics ( $P = 10$ ) is  $\sim 18$  m,  $\sim 5$  m, and  $\sim 27$ s, respectively.

In Table 1, we present our results for all three models. We show the expressions synthesized by SyCAM, and provide a comparison with GradCAM, GradCAM++, ScoreCAM, and AblationCAM, showing the average scores for each method when evaluated using the Deletion and the Insertion metrics. In general, SyCAM synthesizes expressions better than the base methods. In some cases, SyCAM synthesizes GradCAM (*Grads*), ScoreCAM (*CICScores*), or AblationCAM (*AblScores*). We point to the fact that AblationCAM performs better than the rest of the base methods for the Deletion metric, but not for the Insertion metric, for which it is almost always surpassed by ScoreCAM. This emphasizes the necessity of using a framework like SyCAM to synthesize the best expression for each context. For  $P = 15$  and  $P = 30$ , the computation time grows proportionally to  $P$ . There are some exceptions because of the correctness oracle. In the worst case, every expression is better than the threshold for  $I_1$  and  $I_2$  but worse than  $I$  and must be evaluated over the three sets. Note that both metrics depend on the ratio between  $P$  and the size of the feature maps in the last convolutional layer. Because the feature maps are bigger in VGG-16 and VGG-19 ( $14 \times 14$ ) than in ResNet50 ( $7 \times 7$ ), this dependency makes the values obtained by both metrics larger for ResNet50 than for VGG-19 and VGG-16 when comparing the same CAM methods. Larger values of  $P$  are correlated with better metric scores when comparing the same CAM methods.

*Results for the Class-based Decomposition Approach.* For each of the ten classes, the correctness oracle will use a subset  $I$ , which contains approximately 390 images each, and  $I_1 \subset I$ , which contains 100 images each. The time needed for evaluating an expression over each subset is  $\sim 2$  m 10 s and  $\sim 30$  s, respectively.

To make a fair comparison between the enumerate and the class-based decomposition approach, we established a timeout of 2.4 h per class. Although this limits the size of the expressions synthesized for each class, we can see in Table 2

**Table 1.** Enumerative approach for the models ResNet50, VGG-16, and VGG-19 over the Imagenette dataset using our variants of the Deletion and Insertion metrics (higher is better). G=GradCAM, G+=GradCAM++, S=ScoreCAM, A=AblationCAM.

Model	Deletion metric (P=10)						Insertion metric (P=10)					
	G	G+	S	A	SyCAM	SyCAM Exp.	G	G+	S	A	SyCAM	SyCAM Exp.
ResNet50	$5 \cdot Grads$											
	$+ top_{50}$						$top_{50}$					
	$+ ReLU(Grads)$						$* CICScores$					
	$+ AblScores$											
VGG-16	$2 \cdot Grads$											
	$+ top_5$											
	$+ CICScores$						$CICScores$					
	$+ AblScores$											
VGG-19	$top_{10}$											
	$+ AblScores$						$CICScores$					

Model	Deletion metric (P=30)						Insertion metric (P=30)					
	G	G+	S	A	SyCAM	SyCAM Exp.	G	G+	S	A	SyCAM	SyCAM Exp.
ResNet50	$Grads$											
	$+ top_{10}$						$Grads$					
	$0.6075$						$0.5435$					
VGG-16	$4 \cdot Grads$											
	$+ top_5$											
	$+ CICScores$						$CICScores$					
	$+ AblScores$											
VGG-19	$2 \cdot Grads$											
	$+ ReLU(Grads)$						$+ CICScores$					

that, for ResNet50, the average scores of SyCAM per class show improvement over the enumerative approach except for the Deletion metric for  $L = 10$  case. For certain images, SyCAM resulted in scores higher than those of GradCAM, GradCAM++, ScoreCAM, and AblationCAM. In other cases, SyCAM was outperformed. This is because other methods may get much higher scores for a small subset of the data, but not for the rest. As explained in Sect. 5.1, our goal is to find expressions that are better than the best one found in at least half of the images and that have a higher average score.

#### 5.4 Experiments E3: Incorporating Ground Truth

“Show how to incorporate expert knowledge into the generation of saliency maps.”

We use the COVID-QU-Ex dataset [34], which consists of 5826 chest X-ray images with infection segmentation data distributed between 1456 normal (healthy) images, 2913 COVID-19 images, and 1457 non-COVID images with other diseases. Ground-truth COVID-19 segmentation masks are provided. We fine-tuned a ResNet50 model to correctly classify the test dataset with accuracy higher than 95%, and precision and recall higher than 92% for each class.

We consider a set  $I$  with only the 2913 COVID-19 X-ray images and the associated ground-truth COVID-19 segmentation masks. For the correctness oracle,

**Table 2.** Class-based decomposition approach for ResNet50 over the Imagenette dataset using the Deletion and Insertion metrics (higher is better). G=GradCAM, G+=GradCAM++, S=ScoreCAM, A=AblationCAM.

Model	Deletion metric (P=10)						Insertion metric (P=10)					
	G	G+	S	A	SyCAM	SyCAM Exp.	G	G+	S	A	SyCAM	SyCAM Exp.
1.Tench	0.4329	0.4286	0.4203	0.4258	<b>0.4358</b>	$\text{Grads} + \text{top}_5$	0.2307	0.2209	0.2188	0.2346	<b>0.2367</b>	$\text{CICScores} + \text{AblScores}$
2.English springer	0.4293	0.4193	0.4413	<b>0.4513</b>	0.4406	$\text{Grads}^2 + \text{top}_{50}$	0.1741	0.1639	<b>0.1786</b>	0.1706	0.1741	$\text{Grads}$
3.Cassette player	0.2755	0.2719	0.2698	0.2577	<b>0.2766</b>	$\text{top}_{50}$	0.1584	0.1517	0.1514	0.1201	<b>0.1663</b>	$\text{top}_{10}$
4.Chain saw	0.4360	0.4301	<b>0.4439</b>	0.4432	0.4404	$\text{top}_{10}$	0.2524	0.2425	<b>0.2536</b>	0.2327	0.2524	$\text{Grads}$
5.Church	0.2634	0.2468	0.2505	<b>0.2799</b>	0.2640	$2 \cdot \text{Grads} + \text{top}_{20}$	0.0859	0.0803	0.0880	0.0862	<b>0.0891</b>	$\text{top}_{20}$
6.French horn	0.3891	0.3849	0.3831	<b>0.3972</b>	0.3891	$\text{Grads}$	0.2550	0.2441	0.2460	<b>0.2693</b>	0.2550	$\text{Grads}$
7.Garbage truck	0.3757	0.3643	0.3650	<b>0.3779</b>	0.3757	$\text{Grads}$	0.2411	0.2390	0.2412	0.2328	<b>0.2438</b>	$\text{Grads} + \text{top}_{20} + \text{Grads}^2$
8.Gas Pump	0.3405	0.3262	0.3369	<b>0.3470</b>	0.3405	$\text{Grads}$	0.0508	0.0472	<b>0.0570</b>	0.0464	0.0508	$\text{Grads}$
9.Golf ball	0.3309	0.3272	0.3206	0.3279	<b>0.3356</b>	$\text{top}_{10}$	0.4149	0.4047	0.3953	<b>0.4200</b>	0.4167	$\text{top}_{50}$
10.Parachute	0.2855	0.2805	<b>0.2918</b>	0.2897	0.2858	$2 \cdot \text{Grads} + \text{top}_{10}$	<b>0.4343</b>	0.4205	0.4231	0.4129	<b>0.4343</b>	$\text{Grads}$
Average	0.3559	0.3480	0.3523	<b>0.3598</b>	0.3584		0.2287	0.2215	0.2253	0.2226	<b>0.2308</b>	

Model	Deletion metric (P=30)						Insertion metric (P=30)					
	G	G+	S	A	SyCAM	SyCAM Exp.	G	G+	S	A	SyCAM	SyCAM Exp.
1.Tench	0.7215	0.7170	0.7120	0.7148	<b>0.7236</b>	$\text{top}_{50} + \text{CICScores}$	0.6415	0.6310	0.6297	0.6398	<b>0.6415</b>	$\text{Grads}$
2.English springer	0.6360	0.6290	0.6365	<b>0.6447</b>	0.6416	$\text{top}_{50} + \text{CICScores}$	0.4815	0.4672	0.4746	<b>0.4855</b>	0.4815	$\text{Grads}$
3.Cassette player	0.4374	0.4336	0.4323	0.4169	<b>0.4385</b>	$\text{top}_{10}$	0.3404	0.3256	0.3324	0.2932	<b>0.3427</b>	$\text{Grads} + \text{top}_{25}$
4.Chain saw	0.6711	0.6667	0.6694	0.6504	<b>0.6730</b>	$\text{top}_{10}$	0.5896	0.5698	0.5863	0.5678	<b>0.5896</b>	$\text{Grads}$
5.Church	0.4677	0.4548	0.4580	<b>0.4755</b>	0.4692	$2 \cdot \text{Grads} + \text{top}_{20}$	0.3488	0.3319	0.3442	<b>0.3632</b>	0.3501	$2 \cdot \text{Grads} + \text{top}_{20}$
6.French horn	0.6464	0.6412	0.6363	<b>0.6527</b>	0.6476	$\text{top}_{50}$	0.5968	0.5816	0.5867	<b>0.6126</b>	0.5968	$\text{Grads}$
7.Garbage truck	0.6676	0.6619	0.6622	<b>0.6681</b>	0.6676	$\text{Grads}$	0.6072	0.5972	0.6030	<b>0.6091</b>	0.6072	$\text{Grads}$
8.Gas Pump	0.6286	0.6223	0.6288	<b>0.6341</b>	0.6286	$\text{Grads}$	0.4049	0.3862	0.4008	0.4045	<b>0.4050</b>	$2 \cdot \text{Grads} + \text{top}_{25}$
9.Golf ball	0.5865	0.5791	0.5649	0.5875	<b>0.5895</b>	$\text{top}_{10}$	0.7158	0.7072	0.6987	<b>0.7224</b>	0.7158	$\text{Grads}$
10.Parachute	0.6041	0.5972	<b>0.6090</b>	0.5930	0.6041	$\text{Grads}$	0.7084	0.6924	0.6928	0.6982	<b>0.7084</b>	$\text{Grads}$
Average	0.6075	0.6010	0.6017	0.6049	<b>0.6091</b>		0.5435	0.5289	0.5349	0.5402	<b>0.5438</b>	

we use three subsets of images,  $I \supset I_2 \supset I_1$ . The set  $I_2$  is a subset of 1000 images, while  $I_1$  contains 100 images. We apply SyCAM using two metrics: the  $m_{GT}$  and the SCH metrics defined in Sect. 2.3. Because these metrics give importance to the ground truth COVID-19 masks, the SyCAM expressions incorporate expert knowledge into the saliency maps generated. Since we are only considering a binary classification, we employ the enumerative approach. For both metrics defined above, we let the experiments run for 24 h. In this experiment, evaluating  $m_{GT}$  or SCH takes around 0.05s per image.

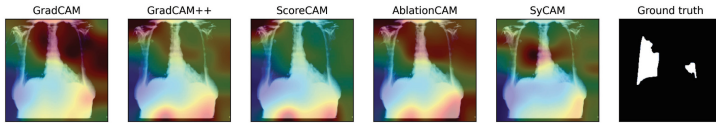
The results are presented in Table 3. We can observe that, for both metrics, SyCAM synthesizes expressions whose basis is of the form  $\exp^n$ , with  $\exp$  a

terminal and  $n \in \mathbb{N}$ :  $CICScores^4$  for the  $m_{GT}$  metric, and  $Grads^5$  for the SCH metric. This may indicate that dominant expressions have this form, which highly reduces the values lower than 1 and gives more importance to the high values.

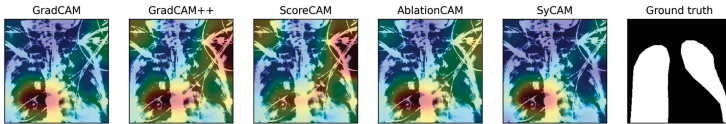
Examples of saliency maps generated by SYCAM are included in Figs. 1, 5a and 5b. We can observe in the X-ray images of Fig. 1 that the saliency map generated by SYCAM is more concentrated on the ground-truth mask, while the rest of the CAM methods fail to do so. In fact, we can see that GradCAM++ and ScoreCAM give high importance to a region outside of the body in the bottom right. Something similar happens in Fig. 5a. In Fig. 5b, because the SYCAM expression is  $2 \cdot Grads^5$ , the heatmap obtained is a more compact version of the GradCAM heatmap. This reduces the importance of areas outside of the ground truth, achieving a higher SCH score.

**Table 3.** Enumerative approach for the fine-tuned model ResNet50 over the COVID-19 benchmark. The metrics used are the  $m_{GT}$  and the SCH metrics (higher is better).

Metric	GradCAM	GradCAM++	ScoreCAM	AblationCAM	SyCAM	SyCAM Exp.
$m_{GT}$	0.1861	0.1712	0.1512	0.1876	<b>0.1972</b>	$Grads * top_{20} * CIC^4$
SCH	0.1567	0.1491	0.1458	0.1524	<b>0.1610</b>	$2 \cdot Grads^5$



(a) Heatmaps generated for the  $m_{GT}$  metric (higher is better). The score for each method is 0.0 for every method except for SYCAM, for which it is 0.3389. The SYCAM heatmap is more concentrated on the ground truth mask.



(b) Heatmaps generated for the SCH metric (higher is better). The scores for each method are 0.3474, 0.2874, 0.2650, 0.3377, and 0.3803, respectively. The SYCAM heatmap is a more compact version of GradCAM.

**Fig. 5.** Saliency maps generated by GradCAM, GradCAM++, ScoreCAM, AblationCAM, and the SyCAM expression synthesized for ResNet50.

## 6 Discussion and Future Work

We presented SyCAM, a metric-based synthesis framework for automatically generating CAM expressions. SyCAM offers advantages in tailoring CAM expressions to specific syntactic restrictions, datasets, and evaluation metrics. However, our approach still has limitations that we plan to address in future work. One key limitation is the high computation time required to generate expressions, which stems from the complexity of the synthesis algorithms. Improvements in synthesis methods will directly enhance our approach. However, we emphasize that generating a SyCAM expression is a one-time process, and the computation time for generating saliency maps remains reasonable. Additionally, the SyCAM framework relies on enumerating expressions and verifying their validity. Future work will explore ways to guide this search process more efficiently, thereby reducing computation time. Lastly, we also aim to extend our study to multi-objective settings [35], incorporating multiple evaluation metrics simultaneously and identifying Pareto-optimal expressions as solutions.

## References

1. Albarghouthi, A., Gulwani, S., Kincaid, Z.: Recursive program synthesis. In: Sharygina, N., Veith, H. (eds.) CAV 2013. LNCS, vol. 8044, pp. 934–950. Springer, Heidelberg (2013). [https://doi.org/10.1007/978-3-642-39799-8\\_67](https://doi.org/10.1007/978-3-642-39799-8_67)
2. Alur, R., et al.: Syntax-guided synthesis. In: Formal Methods in Computer-Aided Design, FMCAD 2013, Portland, OR, USA, 20–23 October 2013, pp. 1–8. IEEE (2013), <https://ieeexplore.ieee.org/document/6679385/>
3. Arras, L., Osman, A., Samek, W.: CLEVR-XAI: a benchmark dataset for the ground truth evaluation of neural network explanations. *Inf. Fusion* **81**, 14–40 (2022). <https://doi.org/10.1016/J.INFFUS.2021.11.008>
4. Chattopadhyay, A., Sarkar, A., Howlader, P., Balasubramanian, V.N.: Grad-CAM++: generalized gradient-based visual explanations for deep convolutional networks. In: 2018 IEEE Winter Conference on Applications of Computer Vision, WACV 2018, Lake Tahoe, NV, USA, 12–15 March 2018, pp. 839–847. IEEE Computer Society (2018). <https://doi.org/10.1109/WACV.2018.00097>
5. Deng, J., Dong, W., Socher, R., Li, L., Li, K., Fei-Fei, L.: Imagenet: a large-scale hierarchical image database. In: 2009 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR 2009), 20–25 June 2009, Miami, Florida, USA, pp. 248–255. IEEE Computer Society (2009). <https://doi.org/10.1109/CVPR.2009.5206848>
6. Desai, S., Ramaswamy, H.G.: Ablation-CAM: visual explanations for deep convolutional network via gradient-free localization. In: IEEE Winter Conference on Applications of Computer Vision, WACV 2020, Snowmass Village, CO, USA, 1–5 March 2020, pp. 972–980. IEEE (2020). <https://doi.org/10.1109/WACV45572.2020.9093360>



7. Dhurandhar, A., Chen, P., Luss, R., Tu, C., Ting, P., Shanmugam, K., Das, P.: Explanations based on the missing: towards contrastive explanations with pertinent negatives. In: Bengio, S., Wallach, H.M., Larochelle, H., Grauman, K., Cesa-Bianchi, N., Garnett, R. (eds.) *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, 3–8 December 2018, Montréal, Canada*, pp. 590–601 (2018), <https://proceedings.neurips.cc/paper/2018/hash/c5ff2543b53f4cc0ad3819a36752467b-Abstract.html>
8. Everingham, M., Van Gool, L., Williams, C.K.I., Winn, J., Zisserman, A.: The PASCAL visual object classes challenge 2007 (VOC2007) Results. <http://www.pascal-network.org/challenges/VOC/voc2007/workshop/index.html>
9. Fletcher, L., van der Klis, R., Sedláček, M., Vasilev, S., Athanasiadis, C.: Reproducibility study of “LICO: Explainable models with language-image consistency” (2024), <https://doi.org/10.48550/arXiv.2410.13989>
10. Fu, R., Hu, Q., Dong, X., Guo, Y., Gao, Y., Li, B.: Axiom-based grad-CAM: towards accurate visualization and explanation of CNNs. In: *31st British Machine Vision Conference 2020, BMVC 2020, Virtual Event, UK, 7–10 September 2020*. BMVA Press (2020), <https://www.bmvc2020-conference.com/assets/papers/0631.pdf>
11. Girshick, R., Donahue, J., Darrell, T., Malik, J.: Rich feature hierarchies for accurate object detection and semantic segmentation. In: *2014 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 580–587. IEEE Computer Society, Los Alamitos, CA, USA, June 2014. <https://doi.org/10.1109/CVPR.2014.81>
12. Hooker, S., Erhan, D., Kindermans, P., Kim, B.: A benchmark for interpretability methods in deep neural networks. In: Wallach, H.M., Larochelle, H., Beygelzimer, A., d’Alché-Buc, F., Fox, E.B., Garnett, R. (eds.) *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, 8–14 December 2019, Vancouver, BC, Canada*, pp. 9734–9745 (2019), <https://proceedings.neurips.cc/paper/2019/hash/fe4b8556000d0f0cae99daa5c5c5a410-Abstract.html>
13. Howard, J.: Imagenette, <https://github.com/fastai/imagenette/>
14. Ibrahim, R., Shafiq, M.O.: Explainable convolutional neural networks: a taxonomy, review, and future directions. *ACM Comput. Surv.* **55**(10), 206:1–206:37 (2023). <https://doi.org/10.1145/3563691>
15. Jha, S., Seshia, S.A.: A theory of formal synthesis via inductive learning. *Acta Informatica* **54**(7), 693–726 (2017). <https://doi.org/10.1007/S00236-017-0294-5>
16. Kadir, M.A., Mosavi, A., Sonntag, D.: Evaluation metrics for XAI: a review, taxonomy, and practical applications. In: *2023 IEEE 27th International Conference on Intelligent Engineering Systems (INES)*, pp. 000111–000124 (2023). <https://doi.org/10.1109/INES59282.2023.10297629>
17. Karpathy, A., Toderici, G., Shetty, S., Leung, T., Sukthankar, R., Fei-Fei, L.: Large-scale video classification with convolutional neural networks. In: *2014 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1725–1732 (2014). <https://doi.org/10.1109/CVPR.2014.223>
18. Khalil, H., et al.: Classification of diabetic retinopathy types based on convolution neural network (CNN). *Menoufia J. Electron. Eng. Res.* **28**(ICEEM2019-Special Issue), 126–153 (2019). <https://doi.org/10.21608/mjeer.2019.76962>
19. Lawrence, S., Giles, C., Tsoi, A.C., Back, A.: Face recognition: a convolutional neural-network approach. *IEEE Trans. Neural Networks* **8**(1), 98–113 (1997). <https://doi.org/10.1109/72.554195>



20. Leem, S., Seo, H.: Attention guided CAM: visual explanations of vision transformer guided by self-attention. In: Wooldridge, M.J., Dy, J.G., Natarajan, S. (eds.) *Thirty-Eighth AAAI Conference on Artificial Intelligence, AAAI 2024, Thirty-Sixth Conference on Innovative Applications of Artificial Intelligence, IAAI 2024, Fourteenth Symposium on Educational Advances in Artificial Intelligence, EAAI 2024*, 20–27 February 2024, Vancouver, Canada, pp. 2956–2964. AAAI Press (2024). <https://doi.org/10.1609/AAAI.V38I4.28077>
21. Lerma, M., Lucas, M.: Grad-CAM++ is equivalent to Grad-CAM with positive gradients (2022). <https://doi.org/10.48550/ARXIV.2205.10838>
22. Marcel, S., Rodriguez, Y.: Torchvision the machine-vision package of torch. In: Bimbo, A.D., Chang, S., Smeulders, A.W.M. (eds.) *Proceedings of the 18th International Conference on Multimedia 2010*, Firenze, Italy, 25–29 October 2010, pp. 1485–1488. ACM (2010). <https://doi.org/10.1145/1873951.1874254>
23. Miró-Nicolau, M., i Capó, A.J., Moyà-Alcove, G.: Assessing fidelity in xai post-hoc techniques: a comparative study with ground truth explanations datasets. *Artif. Intell.* **335**, 104179 (2024). <https://doi.org/10.1016/j.artint.2024.104179>
24. Morbidelli, P., Carrera, D., Rossi, B., Fragneto, P., Boracchi, G.: Augmented Grad-CAM: heat-maps super resolution through augmentation. In: *2020 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2020*, Barcelona, Spain, 4–8 May 2020, pp. 4067–4071. IEEE (2020). <https://doi.org/10.1109/ICASSP40776.2020.9054416>
25. Muhammad, M.B., Yeasin, M.: Eigen-CAM: visual explanations for deep convolutional neural networks. *SN Comput. Sci.* **2**(1), 47 (2021). <https://doi.org/10.1007/S42979-021-00449-3>
26. Omeiza, D., Speakman, S., Cintas, C., Weldemariam, K.: Smooth Grad-CAM++: an enhanced inference level visualization technique for deep convolutional neural network models (2019), <http://arxiv.org/abs/1908.01224>
27. Petsiuk, V., Das, A., Saenko, K.: RISE: randomized input sampling for explanation of black-box models. In: *British Machine Vision Conference 2018, BMVC 2018*, Newcastle, UK, 3–6 September 2018, p. 151. BMVA Press (2018), <http://bmvc2018.org/contents/papers/1064.pdf>
28. Rafferty, A., Ramaesh, R., Rajan, A.: Transparent and clinically interpretable AI for lung cancer detection in chest X-rays (2024), <https://arxiv.org/abs/2403.19444>
29. Rieger, L., Hansen, L.K.: IROF: a low resource evaluation metric for explanation methods (2020), <https://arxiv.org/abs/2003.08747>
30. Rong, Y., Leemann, T., Borisov, V., Kasneci, G., Kasneci, E.: A consistent and efficient evaluation strategy for attribution methods. In: Chaudhuri, K., Jegelka, S., Song, L., Szepesvári, C., Niu, G., Sabato, S. (eds.) *International Conference on Machine Learning, ICML 2022*, 17–23 July 2022, Baltimore, Maryland, USA. *Proceedings of Machine Learning Research*, vol. 162, pp. 18770–18795. PMLR (2022), <https://proceedings.mlr.press/v162/rong22a.html>
31. Samek, W., Binder, A., Montavon, G., Lapuschkin, S., Müller, K.: Evaluating the visualization of what a deep neural network has learned. *IEEE Trans. Neural Networks Learn. Syst.* **28**(11), 2660–2673 (2017). <https://doi.org/10.1109/TNNLS.2016.2599820>
32. Selvaraju, R.R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., Batra, D.: Grad-cam: visual explanations from deep networks via gradient-based localization. In: *2017 IEEE International Conference on Computer Vision (ICCV)*, pp. 618–626 (2017). <https://doi.org/10.1109/ICCV.2017.74>

33. Szczepankiewicz, K., et al.: Ground truth based comparison of saliency maps algorithms. *Sci. Rep.* **13**(1), 16887 (2023). <https://doi.org/10.1038/s41598-023-42946-w>
34. Tahir, A.M., et al.: Covid-qu-ex. Kaggle (2021). <https://doi.org/10.34740/3122958>
35. Torfah, H., Shah, S., Chakraborty, S., Akshay, S., Seshia, S.A.: Synthesizing pareto-optimal interpretations for black-box models. In: *Formal Methods in Computer Aided Design, FMCAD 2021, New Haven, CT, USA, 19–22 October 2021*, IEEE (2021), [https://doi.org/10.34727/2021/isbn.978-3-85448-046-4\\_24](https://doi.org/10.34727/2021/isbn.978-3-85448-046-4_24)
36. Wang, H., et al.: Score-CAM: score-weighted visual explanations for convolutional neural networks. In: *2020 Conference on Computer Vision and Pattern Recognition, CVPR Workshops 2020, Seattle, WA, USA, 14–19 June 2020*, pp. 111–119. Computer Vision Foundation/IEEE (2020). <https://doi.org/10.1109/CVPRW50498.2020.00020>
37. Zhou, B., Khosla, A., Lapedriza, À., Oliva, A., Torralba, A.: Learning deep features for discriminative localization. In: *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, 27–30 June 2016*, pp. 2921–2929. IEEE Computer Society (2016). <https://doi.org/10.1109/CVPR.2016.319>

**Open Access** This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

