

Paying More Attention to Saliency: Image Captioning with Saliency and Context Attention

MARCELLA CORNIA and LORENZO BARALDI, University of Modena and Reggio Emilia
GIUSEPPE SERRA, University of Udine
RITA CUCCHIARA, University of Modena and Reggio Emilia

Image captioning has been recently gaining a lot of attention thanks to the impressive achievements shown by deep captioning architectures, which combine Convolutional Neural Networks to extract image representations and Recurrent Neural Networks to generate the corresponding captions. At the same time, a significant research effort has been dedicated to the development of saliency prediction models, which can predict human eye fixations. Even though saliency information could be useful to condition an image captioning architecture, by providing an indication of what is salient and what is not, research is still struggling to incorporate these two techniques. In this work, we propose an image captioning approach in which a generative recurrent neural network can focus on different parts of the input image during the generation of the caption, by exploiting the conditioning given by a saliency prediction model on which parts of the image are salient and which are contextual. We show, through extensive quantitative and qualitative experiments on large-scale datasets, that our model achieves superior performance with respect to captioning baselines with and without saliency and to different state-of-the-art approaches combining saliency and captioning.

CCS Concepts: • **Computing methodologies** → **Natural language generation**; **Scene understanding**;

Additional Key Words and Phrases: Saliency, visual saliency prediction, image captioning, deep learning

ACM Reference format:

Marcella Cornia, Lorenzo Baraldi, Giuseppe Serra, and Rita Cucchiara. 2018. Paying More Attention to Saliency: Image Captioning with Saliency and Context Attention. *ACM Trans. Multimedia Comput. Commun. Appl.* 14, 2, Article 48 (April 2018), 21 pages.
<https://doi.org/10.1145/3177745>

1 INTRODUCTION

A core problem in computer vision and artificial intelligence is that of building a system that can replicate the human ability of understanding a visual stimulus and describing it in natural language. Indeed, this kind of system would have a great impact on society, opening up new progress

This work is partially supported by “Città educante” (CTN01_00034_393801) of the National Technological Cluster on Smart Communities (cofunded by the Italian Ministry of Education, University and Research - MIUR) and by the project “JUMP - Una piattaforma sensoristica avanzata per rinnovare la pratica e la fruizione dello sport, del benessere, della riabilitazione e del gioco educativo,” funded by the Emilia-Romagna region within the POR-FESR 2014-2020 program. We acknowledge the CINECA award under the ISCRA initiative for the availability of high-performance computing resources and support. We also gratefully acknowledge the support of NVIDIA Corporation with the donation of the GPUs used for this research. Authors’ addresses: M. Cornia, L. Baraldi, and R. Cucchiara, Department of Engineering “Enzo Ferrari,” University of Modena and Reggio Emilia, Via P. Vivarelli 10, 41125 Modena, Italy; emails: {marcella.cornia, lorenzo.baraldi, rita.cucchiara}@unimore.it; G. Serra is with the Department of Computer Science, Mathematics and Physics, University of Udine, Udine, Italy, Via delle Scienze 206, 33100 Udine, Italy; email: giuseppe.serra@uniud.it.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2018 ACM 1551-6857/2018/04-ART48 \$15.00

<https://doi.org/10.1145/3177745>

in human-machine interaction and collaboration. Recent advancements in computer vision and machine translation, together with the availability of large datasets, have made it possible to generate natural sentences describing images. In particular, deep image captioning architectures have shown impressive results in discovering the mapping between visual descriptors and words [24, 55, 56, 59]. They combine Convolutional Neural Networks (CNNs), to extract an image representation, and Recurrent Neural Networks (RNNs), to build the corresponding sentence.

While the progress of these techniques is encouraging, human ability in the construction and formulation of a sentence is still far from being adequately emulated in today's image captioning systems. When humans describe a scene, they look at an object before naming it in a sentence [14], and they do not focus on each region with the same intensity, as selective mechanisms attract their gaze on saliency and relevant parts of the scene [43]. Also, they care about the context using peripheral vision, so that the description of an image alludes not only to the main objects in the scene, and to how they relate to each other, but also to the context in which they are placed in the image.

An intensive research effort has been carried out in the computer vision community to predict where humans look in an image. This task, called saliency prediction, has been tackled in early works by defining hand-crafted features that capture low-level cues such as color and texture or higher-level concepts such as faces, people, and text [4, 19, 23]. Recently, with the advent of deep neural networks and large annotated datasets, saliency prediction techniques have obtained impressive results generating maps that are very close to the ones computed with eye-tracking devices [8, 18, 20].

Despite the encouraging progress in image captioning and visual saliency, and their close connections, these two fields of research have remained almost separate. In fact, only few attempts have been recently presented in this direction [48, 52]. In particular, Sugano et al. [48] presented a gaze-assisted attention mechanism for image caption based on human eye fixations (i.e., the static states of gaze upon a specific location). Although this strategy confirms the importance of using eye fixations, it requires gaze information from a human operator. Therefore, it cannot be applied on general visual data archives, in which this information is missing. To overcome this limit, Tavakoli et al. [52] presented an image captioning method based on saliency maps, which can be automatically predicted from the input image.

In this article, we present an approach that incorporates saliency prediction to effectively enhance the quality of image description. We propose a generative Recurrent Neural Network architecture that can focus on different regions of the input image by means of an attentive mechanism. This attentive behavior, differently from previous works [56], is conditioned by two different attention paths: the former focused on salient spatial regions, predicted by a saliency model, and the latter focused on contextual regions, which are computed as well from saliency maps. Experimental results on five public image captioning datasets (SALICON, COCO, Flickr8k, Flickr30k, and PASCAL-50S) demonstrate that our solution is able to properly exploit saliency cues. Also, we show that this is done without losing the key properties of the generated captions, such as their diversity and the vocabulary size. By visualizing the states of both attentive paths, we finally show that the trained model has learned to attend to both salient and contextual regions during the generation of the caption, and that attention focuses produced by the network effectively correspond, step by step, to generated words.

To sum up, our contributions are as follows. First, we show that saliency can enhance image description, as it provides an indication of what is salient and what is context. Second, we propose a model in which the classic machine attention approach is extended to incorporate two attentive paths, one for salient regions and one for context. These two paths cooperate together during the generation of the caption and generate better captions according to automatic metrics, without

loss of diversity and size of the dictionary. Third, we qualitatively show that the trained model has learned to attend to both salient and contextual regions in an appropriate way.

2 RELATED WORK

In this section, we review the literature related to saliency prediction and image captioning. We also report some recent works that investigate the contribution of saliency for generating natural language descriptions.

2.1 Visual Saliency Prediction

Saliency prediction has been extensively studied by the computer vision community and, in the last few years, has achieved considerable improvement thanks to the large spread of deep neural networks [8, 9, 18, 20, 28, 30, 39]. However, a very large variety of models had been proposed before the advent of deep learning, and almost all of them have been inspired by the seminal work of Itti and Koch [19], in which multiscale low-level features extracted from the input image were linearly combined and then processed by a dynamic neural network with a winner-takes-all strategy. The same idea of properly combining different low-level features was also explored by Harel et al. [15], who defined Markov chains over various image maps and treated the equilibrium distribution over map locations as an activation. In addition to the exploitation of low-level features, several saliency models have also incorporated high-level concepts such as faces, people, and text [4, 23, 61]. In fact, Judd et al. [23] highlighted that, when humans look at images, their gazes are attracted not only by low-level cues typical of the bottom-up attention but also by top-down image semantics. To this end, they proposed a model in which low- and medium-level features were effectively combined and exploited face and people detectors to capture important high-level concepts. Nonetheless, all these techniques have failed to effectively capture the wide variety of causes that contribute to define the visual saliency on images, and, with the advent of deep learning, researchers have developed data-driven architectures capable of overcoming many of the limitations of hand-crafted models.

The first attempts of computing saliency maps through a neural network were lacking because of the absence of sufficiently large training datasets [30, 35, 54]. Vig et al. [54] proposed the first deep architecture for saliency, which was composed by only three convolutional layers. Afterward, Kümmerer et al. [30, 31] based their models on two popular convolutional networks (AlexNet [27] and VGG-19 [46]), obtaining adequate results, despite the fact that the network parameters were not fine-tuned on a saliency dataset. Liu et al. [35] tried to overcome the absence of large-scale datasets by training their model on image patches centered on fixation and nonfixation locations, thus increasing the amount of training data.

With the arrival of the SALICON dataset [21], which is still the largest publicly available dataset for saliency prediction, several deep architectures have moved beyond previous approaches, bringing consistent performance advances. The starting point of all these architectures is a pretrained Convolutional Neural Network (CNN), such as VGG-16 [46], GoogleNet [50], and ResNet [16], to which different saliency-oriented components are added [8, 9], together with different training strategies [9, 18, 20].

In particular, Huang et al. [18] compared three standard CNNs by applying them at two different image scales. In addition, they were the first to train the network using a saliency evaluation metric as a loss function. Jetley et al. [20] introduced a model that formulates a saliency map as generalized Bernoulli distribution. Moreover, they trained their network by using different loss functions that pair the softmax activation function with measures designed to compute distances between probability distributions. Tavakoli et al. [51] investigated interimage similarities to estimate the saliency of a given image using an ensemble of extreme learners, each trained on an

image similar to the input image. Kruthiventi et al. [28], instead, presented a unified framework to predict both eye fixations and salient objects.

Another saliency prediction model was recently presented by Pan et al. [38], who, following the large dissemination of Generative Adversarial Networks, trained their model by using adversarial examples. In particular, their architecture is composed by two agents: a generator, which is responsible for generating the saliency map of a given image, and a discriminator, which performs a binary classification task between generated and real saliency maps. Liu et al. [34], instead, proposed a model to learn long-term spatial interactions and scene contextual modulation to infer image saliency and showed promising results, also thanks to the use of the powerful ResNet-50 architecture [16].

In contrast to all these works, we present two different deep saliency architectures. The first one, called *ML-Net* [8], effectively combines features coming from different levels of a CNN and applies a matrix of learned weights to the predicted saliency map, thus taking into account the center bias present in human eye fixations. The second one, called *SAM* [9], incorporates neural attentive mechanisms that focus on the most salient regions of the input image. The core component of the proposed model is an Attentive Convolutional LSTM that iteratively refines the predicted saliency map. Moreover, to tackle the human center bias, the network is able to learn multiple Gaussian prior maps without predefined information. Since this model achieved state-of-the-art performances, being at the top of different saliency prediction benchmarks, we use it in this work.

2.2 Image Captioning

Recently, the automatic description of images and video has been addressed by computer vision researchers with recurrent neural networks, which, given a vectored description of the visual content, can naturally deal with sequences of words [3, 24, 55]. Before deep learning models, the generation of sentences was mainly tackled by identifying visual concepts, objects, and attributes, which were then combined into sentences using predefined templates [29, 57, 58]. Another strategy was that of posing the image captioning as a retrieval problem, where the closest annotated sentence in the training set was transferred to a test image, or where training captions were split into parts and then reassembled to form new sentences [11, 17, 37, 47]. Obviously, all these approaches limited the variety of possible outputs and could not satisfy the richness of natural language. Recent captioning models, in fact, address the generation of sentences as a machine translation problem in which a visual representation of the image coming from a convolutional network is translated in a language counterpart through a recurrent neural network.

One of the first models based on this idea is that proposed by Karpathy et al. [24], in which sentence snippets are aligned to the visual regions that they describe through a multimodal embedding. After that, these correspondences are treated as training data for a multimodal recurrent neural network that learns to generate the corresponding sentences. Vinyals et al. [55], instead, developed an end-to-end model trained to maximize the likelihood of the target sentence given the input image. Xu et al. [56] introduced an approach to image captioning that incorporates a form of machine attention, by which a generative LSTM can focus on different regions of the image while generating the corresponding caption. They proposed two different versions of their model: the first one, called “Soft Attention,” is trained in a deterministic manner using standard back-propagation techniques, while the second one, called “Hard Attention,” is trained by maximizing a variational lower bound through the reinforcement learning paradigm.

Johnson et al. [22] addressed the task of dense captioning, which jointly localizes and describes in natural language salient image regions. This task consists of generalizing the object detection problem when the descriptions consist of a single word and the image captioning task when one predicted region covers the full image. You et al. [59] proposed a semantic attention model in

which, given an image, a convolutional neural network extracts top-down visual features and at the same time detects visual concepts such as regions, objects, and attributes. The image features and the extracted visual concepts are combined through a recurrent neural network that finally generates the image caption. Differently from previous works that aim at predicting a single caption, Krause et al. [26] introduced the generation of entire paragraphs for describing images. Finally, Shetty et al. [45] employed adversarial training to change the training objective of the caption generator from reproducing ground-truth captions to generating a set of captions that is indistinguishable from human-generated captions.

In this article, we are interested in demonstrating the importance of using saliency along with contextual information during the generation of image descriptions. Our solution falls in the class of neural attentive captioning architectures, and in the experimental section, we compare it against a standard attentive model built upon the Soft Attention approach presented in [56].

2.3 Visual Saliency and Captioning

Only a few other previous works have investigated the contribution of human eye fixations to generate image descriptions. The first work that has explored this idea was that proposed in [48], which presented an extension of a neural attentive captioning architecture. In particular, the proposed model incorporates human fixation points (obtained with eye-tracking devices) instead of computed saliency maps to generate image captions. This kind of strategy mainly suffers from the need of having both eye fixation and caption annotations. Currently, only the SALICON dataset [21], being a subset of the Microsoft COCO dataset [33], is available with both human descriptions and saliency maps.

Ramanishka et al. [41], instead, introduced an encoder-decoder captioning model in which spatiotemporal heatmaps are produced for predicted captions and arbitrary query sentences without explicit attention layers. They refer to these heatmaps as saliency maps, even though they are internal representations of the network, not related to human attention. Experiments showed that the gain in performance with respect to a standard captioning attentive model is not consistent, even though the computational overhead is lower.

A different approach, presented in [52], explores if image descriptions, by humans or models, agree with saliency and if saliency can benefit image captioning. To this end, they proposed a captioning model in which image features are boosted with the corresponding saliency map by exploiting a moving sliding window and mean pooling as aggregation strategies. Comparisons with respect to a no-saliency baseline did not show significant improvements (especially on the Microsoft COCO dataset).

In this article, we instead aim at enhancing image captions by directly incorporating saliency maps in a neural attentive captioning architecture. Differently from previous models that exploit human fixation points, we obtain a more general architecture that can be potentially trained using any image captioning dataset and can predict captions for any input image. In our model, the machine attentive process is split in two different and unrelated paths, one for salient regions and one for context. We demonstrate through extensive experiments that the incorporation of saliency and context can enhance image captioning on different state-of-the-art datasets.

3 WHAT IS HIT BY SALIENCY?

Human gazes are attracted by both low-level cues, such as color, contrast, and texture, and high-level concepts, such as faces and text [6, 23]. Current state-of-the-art saliency prediction methods, thanks to the use of deep networks and large-scale datasets, are able to effectively incorporate all these factors and predict saliency maps that are very close to those obtained from human eye fixations [9]. In this section, we qualitatively investigate which parts of an image are actually hit or




Fig. 1. Ground-truth semantic segmentation and saliency predictions from our model [9] on sample images from Pascal-Context [36] (first row), Cityscapes [7] (second row), and LIP [13] (last row).

ignored by saliency models, by jointly analyzing saliency and semantic segmentation maps. This will motivate the need for using saliency predictions as an additional conditioning for captioning models.

To compute saliency maps, we employ the approach in [9], which has shown good results on popular saliency benchmarks, such as the MIT Saliency [5] and the SALICON dataset [21], and which also won the LSUN Challenge in 2017. It is worthwhile to mention, anyway, that the qualitative conclusions of this section can be applied to any state-of-the-art saliency model.

Since semantic segmentation algorithms are not always completely accurate, we perform the analysis on three semantic segmentation datasets, in which regions have been segmented by human annotators: Pascal-Context [36], Cityscapes [7], and the Look into Person (LIP) [13] dataset. While the first one contains natural images without a specific target, the other two are focused, respectively, on urban streets and human body parts. In particular, the Pascal-Context provides additional annotations for the Pascal VOC 2010 dataset [10], which contains 10,103 training and validation images and 9,637 testing images. It goes beyond the original Pascal semantic segmentation task by providing annotations for the whole scene, and images are annotated by using more than 400 different labels. The Cityscapes dataset, instead, is composed by a set of video sequences recorded in street scenes from 50 different cities. It provides high-quality pixel-level annotations for 5,000 frames and coarse annotations for 20,000 frames. The dataset is annotated with 30 street-specific classes, such as *car*, *road*, *traffic sign*, and so forth. Finally, the LIP dataset is focused on the semantic segmentation of people and provides more than 50,000 images annotated with 19 semantic human part labels. Images contain person instances cropped from the Microsoft COCO dataset [33] and split in training, validation, and testing sets with 30,462, 10,000, and 10,000 images, respectively. For our analyses, we only consider train and validation images for the Pascal-Context and LIP datasets, and the 5,000 pixel-level annotated frames for the Cityscapes dataset. Figure 1 shows, for some sample images, the predicted saliency map and the corresponding semantic segmentation on the three datasets.

We first investigate which are the most and the least salient classes for each dataset. Since there are semantic classes with a low number of occurrences with respect to the total number of images, we only consider relevant semantic classes (i.e., classes with at least N occurrences). Due to the different dataset sizes, we set N to 500 for the Pascal-Context and LIP datasets, and to 200 for the Cityscapes dataset. To collect the number of times that the predicted saliency hits a semantic class, we binarize each map by thresholding the values of its pixels. A low threshold value leads to




Fig. 2. Most salient classes on Pascal-Context, Cityscapes, and LIP datasets.




Fig. 3. Least salient classes on Pascal-Context, Cityscapes, and LIP datasets.

a binarized map with dilated salient regions, while a high threshold creates small salient regions around the fixation points. For this reason, we use two different threshold values to analyze the most and the least salient classes. We choose a threshold near 0 to find the least salient classes for each dataset and a value near 255 to find instead the most salient ones.

Figures 2 and 3 show the most and the least salient classes in terms of the percentage of times that saliency hits a region belonging to a class. As can be seen, there are different distributions depending on the considered dataset. For example, for Pascal-Context, the most salient classes are animals (such as cats, dogs, and birds), people, and vehicles (such as airplanes and cars), while the least salient ones result to be ceiling, floor, and light. As for the Cityscapes dataset, cars are absolutely the most salient class, with a rate 70% of times in which it is hit by saliency. All other classes, instead, do not reach a 40% rate. On the LIP dataset, the most salient classes are all human body parts in the upper body, while the least salient ones are all in the lower body. As expected, people faces are those most hit by saliency, with an absolute number of occurrences near 90%. It can be observed as a general pattern that *the most important or visible objects in the scene are hit by saliency, while objects in the background, and the context itself of the image, are usually ignored*. This leads to the hypothesis that both salient and nonsalient regions are important to generate the description of an image, given that we generally want the context to be included in the caption, and that the distinction between salient regions and context, given by a saliency prediction model, can improve captioning results.

We also investigate the existence of a relation between the size of an object and its saliency values. In Figure 4, we plot the joint distribution of object sizes and saliency values on the three datasets, where the size of an object is simply computed as the number of its pixels normalized by the size of the image. As can be seen, most of the low-saliency instances are small; however,




Fig. 4. Distribution of object sizes and saliency values (best seen in color).

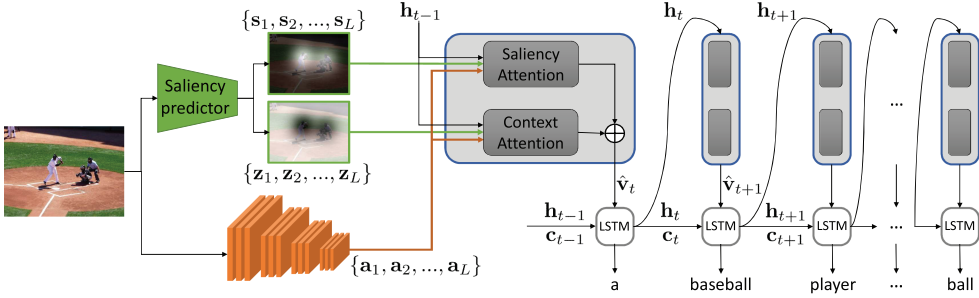


Fig. 5. Overview of the proposed model. Two different attention paths are built for salient regions and contextual regions, to help the model build captions that describe both components (best seen in color).

high-saliency values concentrate on small objects as well as on large ones. In summary, *there is not always a proportionality between the size of an object and its saliency*, so the importance of an object cannot be assessed by simply looking at its size. In the image captioning scenario that we want to tackle, larger objects correspond to larger activations in the last layers of a convolutional architecture, while smaller objects correspond to smaller activations. Since salient and nonsalient regions can have comparable activations, the supervision given by a saliency prediction model on whether a pixel belongs or not to a salient region can be beneficial during the generation of the caption.

4 SALIENCY AND CONTEXT-AWARE ATTENTION

Following the qualitative findings of the previous section, we develop a model in which saliency is exploited to enhance image captioning. Here, a generative recurrent neural network is conditioned, step by step, on salient spatial regions, predicted by a saliency model, and on contextual features, which account for the role of nonsalient regions in the generation of the caption. In the following, we describe the overall model. An overview is presented in Figure 5.

Each input image I is first encoded through a Fully Convolutional Network, which provides a stack of high-level features on a spatial grid $\{a_1, a_2, \dots, a_L\}$, each corresponding to a spatial location of the image. At the same time, we extract a saliency map for the input image using the model in [9] and downscale it to fit the spatial size of the convolutional features to obtain a spatial grid $\{s_1, s_2, \dots, s_L\}$ of salient regions, where $s_i \in [0, 1]$. Correspondingly, we also define a spatial grid of contextual regions, $\{z_1, z_2, \dots, z_L\}$, where $z_i = 1 - s_i$. Under the model, visual features at different locations will be selected or inhibited according to their saliency value.

The generation of the caption is carried out word by word by feeding and sampling words from an LSTM layer, which, at every timestep, is conditioned on features extracted from the input image

and on the saliency map. Formally, the behavior of the generative LSTM is driven by the following equations:

$$\mathbf{i}_t = \sigma(W_{vi}\hat{\mathbf{v}}_t + W_{wi}\mathbf{w}_t + W_{hi}\mathbf{h}_{t-1} + \mathbf{b}_i) \quad (1)$$

$$\mathbf{f}_t = \sigma(W_{vf}\hat{\mathbf{v}}_t + W_{wf}\mathbf{w}_t + W_{hf}\mathbf{h}_{t-1} + \mathbf{b}_f) \quad (2)$$

$$\mathbf{o}_t = \sigma(W_{vo}\hat{\mathbf{v}}_t + W_{wo}\mathbf{w}_t + W_{ho}\mathbf{h}_{t-1} + \mathbf{b}_o) \quad (3)$$

$$\mathbf{g}_t = \phi(W_{vg}\hat{\mathbf{v}}_t + W_{wg}\mathbf{w}_t + W_{hg}\mathbf{h}_{t-1} + \mathbf{b}_g) \quad (4)$$

$$\mathbf{c}_t = \mathbf{f}_t \odot \mathbf{c}_{t-1} + \mathbf{i}_t \odot \mathbf{g}_t \quad (5)$$

$$\mathbf{h}_t = \mathbf{o}_t \odot \phi(\mathbf{c}_t), \quad (6)$$

where, at each timestep, $\hat{\mathbf{v}}_t$ denotes the visual features extracted from I by considering the map of salient regions $\{s_i\}_i$ and those of contextual regions $\{z_i\}_i$. \mathbf{w}_t is the input word, and \mathbf{h} and \mathbf{c} are, respectively, the internal state and the memory cell of the LSTM. \odot denotes the element-wise Hadamard product, σ is the sigmoid function, ϕ is the hyperbolic tangent \tanh , W_* are learned weight matrices, and \mathbf{b}_* are learned biases vectors.

To provide the generative network with visual features, we draw inspiration from the machine attention literature [56] and compute the fixed-length feature vector $\hat{\mathbf{v}}_t$ as a linear combination of spatial features $\{\mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_L\}$ with time-varying weights α_{ti} , normalized over the spatial extent via a *softmax* operator:

$$\hat{\mathbf{v}}_t = \sum_{i=1}^L \alpha_{ti} \mathbf{a}_i, \quad (7)$$

$$\alpha_{ti} = \frac{\exp(e_{ti})}{\sum_{k=1}^L \exp(e_{tk})}. \quad (8)$$

At each timestep, the attention mechanism selects a region of the image, based on the previous LSTM state, and feeds it to the LSTM, so that the generation of a word is conditioned on that specific region, instead of being driven by the entire image.

Ideally, we want weights α_{ti} to be aware of the saliency and contextual value of location \mathbf{a}_i and to be conditioned on the current status of the LSTM, which can be well encoded by its internal state \mathbf{h}_t . In this way, the generative network can focus on different locations of the input image according to their belonging to a salient or contextual region, and to the current generation state. Of course, simply multiplying attention weights with saliency values would result in a loss of context, which is fundamental for caption generation. We instead split attention weights e_{ti} into two contributions, one for saliency and one for context regions, and employ two different fully connected networks to learn the two contributions (Figure 5). Conceptually, this is equivalent to building two separate attention paths, one for salient regions and one for contextual regions, which are merged to produce the final attention. Overall, the model obeys the following equation:

$$e_{ti} = s_i \cdot e_{ti}^{sal} + z_i \cdot e_{ti}^{ctx}, \quad (9)$$

where e_{ti}^{sal} and e_{ti}^{ctx} are, respectively, the attention weights for salient and context regions. Attention weights for saliency and context are computed as follows:

$$e_{ti}^{sal} = \mathbf{v}_{e,sal}^T \cdot \phi(W_{ae,sal} \cdot \mathbf{a}_i + W_{he,sal} \cdot \mathbf{h}_{t-1}) \quad (10)$$

$$e_{ti}^{ctx} = \mathbf{v}_{e,ctx}^T \cdot \phi(W_{ae,ctx} \cdot \mathbf{a}_i + W_{he,ctx} \cdot \mathbf{h}_{t-1}). \quad (11)$$

Notice that our model learns different weights for saliency and contextual regions and combines them into a final attentive map in which the contributions of salient and nonsalient regions are merged together. Similarly to the classical Soft Attention approach [56], the proposed generative LSTM can focus on every region of the image, but the attentive process is aware of the saliency of each location, so that the focus on salient and contextual regions is driven by the output of the saliency predictor.

4.1 Sentence Generation

Words are encoded with one-hot vectors having size equal to that of the vocabulary and are then projected into an embedding space via a learned linear transformation. Because sentences have different lengths, they are also marked with special begin-of-string and end-of-string tokens, to keep the model aware of the beginning and end of a particular sentence.

Given an image and sentence (y_0, y_1, \dots, y_T) , encoded with one-hot vectors, the generative LSTM is conditioned step by step on the first t words of the caption and is trained to produce the next word of the caption. The objective function we optimize is the log-likelihood of correct words over the sequence

$$\max_{\mathbf{w}} \sum_{t=1}^T \log \Pr(y_t | \hat{\mathbf{v}}_t, \mathbf{y}_{t-1}, \mathbf{y}_{t-2}, \dots, \mathbf{y}_0), \quad (12)$$

where \mathbf{w} are all the parameters of the model. The probability of a word is modeled via a softmax layer applied on the output of the LSTM. To reduce the dimensionality, a linear embedding transformation is used to project one-hot word vectors into the input space of the LSTM and, vice versa, to project the output of the LSTM to the dictionary space:

$$\Pr(y_t | \hat{\mathbf{v}}_t, \mathbf{y}_{t-1}, \mathbf{y}_{t-2}, \dots, \mathbf{y}_0) \propto \exp(\mathbf{y}_t^T W_p \mathbf{h}_t), \quad (13)$$

where W_p is a matrix for transforming the LSTM output space to the word space, and \mathbf{h}_t is the output of the LSTM.

At test time, the LSTM is given a begin-of-string tag as input for the first timestep, and then the most probable word according to the predicted distribution is sampled and given as input for the next timestep, until an end-of-string tag is predicted.

5 EXPERIMENTAL EVALUATION

In this section, we perform qualitative and quantitative experiments to validate the effectiveness of the proposed model with respect to different baselines and other saliency-boosted captioning methods. First, we describe datasets and metrics used to evaluate our solution and provide implementation details.

5.1 Datasets and Metrics

To validate the effectiveness of the proposed Saliency and Context-aware Attention, we perform experiments on five popular image captioning datasets: SALICON [21], Microsoft COCO [33], Flickr8k [17], Flickr30k [60], and PASCAL-50S [53].

Microsoft COCO is composed by more than 120,000 images divided into training and validation sets, where each of them is provided with at least five sentences generated by using Amazon Mechanical Turk. SALICON is a subset of this one, created for the visual saliency prediction task. Since its images are from the Microsoft COCO dataset, at least five captions for each image are available. Overall, it contains 10,000 training images, 5,000 validation images, and 5,000 testing images where eye fixations for each image are simulated with mouse movements. In our experiments, we only use train and validation sets for both datasets. The Flickr8k and the Flickr30k datasets are

composed by 8,000 and 30,000 images, respectively. Both of them come with five annotated sentences for each image. In our experiments, we randomly choose 1,000 validation images and 1,000 test images for each of these two datasets. The PASCAL-50S dataset provides additional annotations for the UIUC PASCAL sentences [42]. It is composed of 1,000 images from the PASCAL-VOC dataset, each of them annotated with 50 human-written sentences, instead of five as in the original dataset. Due to the limited number of samples and for a fair comparison with other captioning methods, we first pretrain the model on the Microsoft COCO dataset, and then we test it on the images of this dataset without a specific fine-tuning.

For evaluation, we employ four automatic metrics that are usually employed in image captioning: BLEU [40], ROUGE_L [32], METEOR [2], and CIDEr [53]. BLEU is a modified form of precision between n-grams to compare a candidate translation against multiple reference translations. We evaluate our predictions with BLEU using mono-grams, bi-grams, three-grams, and four-grams. ROUGE_L computes an F-measure considering the longest co-occurring in sequence n-grams. METEOR, instead, is based on the harmonic mean of unigram precision and recall, with recall weighted higher than precision. It also has several features that are not found in other metrics, such as stemming and synonymy matching, along with the standard exact word matching. CIDEr, finally, computes the average cosine similarity between n-grams found in the generated caption and those found in reference sentences, weighting them using TF-IDF. To ensure a fair evaluation, we use the Microsoft COCO evaluation toolkit¹ to compute all scores.

5.2 Implementation Details

Each image is encoded through a convolutional network, which computes a stack of high-level features. We employ the popular ResNet-50 [16], trained over the ImageNet dataset [44], to compute the feature maps over the input image. In particular, the ResNet-50 is composed by 49 convolutional layers, divided into five convolutional blocks, and one fully connected layer. Since we want to maintain the spatial dimensions, we extract the feature maps from the last convolutional layer and ignore the fully connected layer. The output of the ResNet model is a tensor with 2,048 channels. To limit the number of feature maps and the number of learned parameters, we fed this tensor into another convolutional layer with 512 filters and a kernel size of 1, followed by a ReLU activation function. Differently from the weights of the ResNet-50, which are kept fixed, the weights of this last convolutional layer are initialized according to [12] and fine-tuned over the considered datasets. In the LSTM, following the initialization proposed in [1], the weight matrices applied to the inputs are initialized by sampling each element from the Gaussian distribution of 0 mean and 0.01² variance, while the weight matrices applied to the internal states are initialized by using the orthogonal initialization. The vectors $v_{e,sal}$ and $v_{e,ctx}$ as well as all bias vectors \mathbf{b}_* are instead initialized to zero.

To predict the saliency map for each input image, we exploit our Saliency Attentive Model (SAM) [9], which is able to predict accurate saliency maps according to different saliency benchmarks. We note, however, that we do not expect a significant performance variation when using other state-of-the-art saliency methods.

As mentioned, we perform experiments over five different datasets. For the SALICON dataset, since its images all have the same size of 480×640 , we keep the original size of these images, thus obtaining $L = 15 \times 20 = 300$. For all other datasets, which are composed of images with different sizes, we set the input size to 480×480 , obtaining $L = 15 \times 15 = 225$. Since saliency maps are exploited inside the proposed saliency-context attention model, we resize the SALICON saliency maps to have a size of 15×20 , while for all other datasets, we resize them to 15×15 .

¹<https://github.com/tylin/coco-caption>.

All experiments are performed by using the Adam optimizer [25] with Nestorov momentum [49] using an initial learning rate of 0.001 and batch size of 64. The hidden state dimension is set to 1,024 while the embedding size is set to 512. For all datasets, we choose a vocabulary size equal to the number of words that appear at least 5 times in training and validation captions.

5.3 Quantitative Results and Comparisons with Baselines

To assess the performance of our method and to investigate the hypotheses behind it, we first compare with the classic Soft Attention approach, and we then build three baselines in which saliency is used to condition the generative process.

Soft Attention [56]: The visual input to the LSTM is computed via the Soft Attention mechanism to attend to different locations of the image, without considering salient and nonsalient regions. A single feed-forward network is in charge of producing attention values, which can be obtained by replacing Equation (9) with

$$e_{ti} = v_e^T \cdot \phi(W_{ae} \cdot \mathbf{a}_i + W_{he} \cdot \mathbf{h}_{t-1}). \quad (14)$$

This approach is equivalent to the one proposed in [56], although some implementation details are different. In order to achieve a fair evaluation, we use activations from the ResNet-50 model instead of the VGG-19, and we do not include the doubly stochastic regularization trick. For this reason, the numerical results that we report are not directly comparable with those in the original paper (ours are in general higher than the original ones).

Saliency pooling: Visual features from the CNN are multiplied at each location by the corresponding saliency value and then summed, without any attention mechanism. In this case, the visual input of the LSTM is not time dependent, and salient regions are given more focus than nonsalient ones. Comparing with Equation (7), it can be seen as a variation of the Soft Attention in which the network always focuses on salient regions:

$$\hat{\mathbf{v}}_t = \hat{\mathbf{v}} = \sum_{i=1}^L s_i \mathbf{a}_i. \quad (15)$$

Attention on saliency: This is an extension of the Soft Attention approach in which saliency is used to modulate attention values at each location. The attention mechanism, therefore, is conditioned to attend to salient regions with higher probability and to ignore nonsalient regions:

$$e_{ti} = s_i \cdot v_e^T \cdot \phi(W_{ae} \cdot \mathbf{a}_i + W_{he} \cdot \mathbf{h}_{t-1}). \quad (16)$$

Attention on saliency and context (with weight sharing): The attention mechanism is aware of salient and context regions, but weights used to compute the attentive scores of salient and context are shared, excluding the v_e^T vectors. Notice that, if those were shared too, this baseline would be equivalent to the Soft Attention one:

$$e_{ti} = s_i \cdot e_{ti}^{sal} + (1 - s_i) \cdot e_{ti}^{ctx} \quad (17)$$

$$e_{ti}^{sal} = v_{e,sal}^T \cdot \phi(W_{ae} \cdot \mathbf{a}_i + W_{he} \cdot \mathbf{h}_{t-1}) \quad (18)$$

$$e_{ti}^{ctx} = v_{e,ctx}^T \cdot \phi(W_{ae} \cdot \mathbf{a}_i + W_{he} \cdot \mathbf{h}_{t-1}). \quad (19)$$

It is straightforward also to notice that our proposed approach is equivalent to the last baseline, without weight sharing.

In Table 1, we first compare the performance of our method with respect to the Soft Attention approach, to assess the superior performance of the proposal with respect to the published state of the art. We report results on all the datasets, both on validation and test sets, with respect to all the automatic metrics described in Section 5.1. As can be seen, the proposed approach always

Table 1. Image Captioning Results

Dataset	Model	B@1	B@2	B@3	B@4	METEOR	ROUGE _L	CIDEr
SALICON	Soft Attention	69.0	50.9	36.1	25.4	22.5	49.9	70.8
	Saliency+Context Attention	69.2	51.4	37.2	26.9	22.9	50.4	73.3
COCO	Soft Attention	70.6	53.0	38.3	27.5	24.3	51.8	87.9
	Saliency+Context Attention	70.8	53.6	39.1	28.4	24.8	52.1	89.8
Flickr8k (Validation)	Soft Attention	59.9	41.8	27.9	18.2	19.8	45.0	47.7
	Saliency+Context Attention	62.8	44.5	30.2	19.9	20.3	46.5	50.1
Flickr8k (Test)	Soft Attention	61.0	43.2	29.6	20.1	20.8	46.5	53.2
	Saliency+Context Attention	63.5	45.6	31.5	21.2	21.1	47.5	54.1
Flickr30k (Validation)	Soft Attention	61.9	43.3	29.7	20.2	19.9	44.8	43.2
	Saliency+Context Attention	61.3	43.3	30.1	20.9	20.2	45.0	44.5
Flickr30k (Test)	Soft Attention	61.9	43.4	29.9	20.5	19.8	44.5	44.2
	Saliency+Context Attention	61.5	43.8	30.5	21.3	20.0	45.2	46.4
PASCAL-50S	Soft Attention	82.4	70.0	57.0	45.1	32.8	65.9	70.4
	Saliency+Context Attention	82.4	70.2	57.5	45.7	32.9	66.3	70.7

The conditioning of saliency and context (Saliency+Context Attention) enhances the generation of the caption with respect to the traditional machine attention mechanism. Soft Attention here indicates our reimplement of [56], using the same visual features of our model.

overcomes by a significant margin the Soft Attention approach, thus experimentally confirming the benefit of having two separate attention paths, one for salient and one for nonsalient regions, and the role of saliency as a conditioning for captioning. In particular, on the METEOR metric, the relative improvement ranges from $\frac{32.9-32.8}{32.8} = 0.30\%$ on the PASCAL-50S to $\frac{20.3-19.8}{19.8} = 2.53\%$ of the Flickr8k validation set.

In Table 2, instead, we compare our approach with the three baselines that incorporate saliency. First, it can be observed that the Saliency Pooling baseline usually performs worse than the Soft Attention, thus demonstrating that always attending to salient locations is not sufficient to achieve good captioning results. When plugging in attention, as in the Saliency Attention baseline, numerical results are a bit higher, thanks to a time-dependent attention, but still far from the performance achieved by the complete model. It can also be noticed that, even though this baseline does not take into account the context, it sometimes achieves better results than the Soft Attention model (such as in the case of SALICON, with respect to the METEOR metric). Finally, we notice that the baseline with attention on saliency and context, and with weight sharing, is better than Saliency Attention, further confirming the benefit of including the context. Having two completely separated attention paths, such as in our model, is important anyway, as demonstrated by the numerical results of this last baseline with respect to those of our method.

5.4 Comparisons with Other Saliency-Boosted Captioning Models

We also compare to existing captioning models that incorporate saliency during the generation of image descriptions. In particular, we compare to the model proposed in [48], which exploited human fixation points; to the work by Tavakoli et al. [52], which reports experiments on Microsoft COCO and on PASCAL-50S; and to the proposal by Ramanishka et al. [41], which used convolutional activations as a proxy for saliency.

Table 2. Comparison with Image Captioning with Saliency Baselines

Dataset	Model	B@1	B@2	B@3	B@4	METEOR	ROUGE _L	CIDEr
SALICON	Saliency Pooling	66.1	47.8	33.7	24.0	21.1	47.9	62.4
	Attention on Saliency	68.8	51.3	37.0	26.5	22.7	50.1	71.3
	Saliency+Cont. Att. (Weight Sh.)	68.9	51.3	36.8	26.3	22.6	50.2	71.4
	Saliency+Context Attention	69.2	51.4	37.2	26.9	22.9	50.4	73.3
COCO	Saliency Pooling	68.6	50.9	36.3	25.8	23.3	50.2	81.4
	Attention on Saliency	70.4	53.2	38.6	27.6	24.1	51.6	86.6
	Saliency+Cont. Att. (Weight Sh.)	70.4	53.1	38.8	28.2	24.7	52.1	89.4
	Saliency+Context Attention	70.8	53.6	39.1	28.4	24.8	52.1	89.8
Flickr8k (Validation)	Saliency Pooling	56.1	37.7	24.3	15.6	18.3	42.8	37.0
	Attention on Saliency	58.7	40.4	26.8	17.6	19.7	45.1	44.7
	Saliency+Cont. Att. (Weight Sh.)	62.0	43.9	29.6	19.8	20.2	45.7	50.2
	Saliency+Context Attention	62.8	44.5	30.2	19.9	20.3	46.5	50.1
Flickr8k (Test)	Saliency Pooling	56.5	37.8	24.6	16.2	18.5	42.9	37.7
	Attention on Saliency	59.6	42.2	28.7	19.5	20.7	46.1	50.1
	Saliency+Cont. Att. (Weight Sh.)	62.4	44.2	29.9	19.7	21.1	46.7	51.7
	Saliency+Context Attention	63.5	45.6	31.5	21.2	21.1	47.5	54.1
Flickr30k (Validation)	Saliency Pooling	58.7	40.5	27.1	18.4	18.3	43.0	34.2
	Attention on Saliency	63.0	44.5	30.8	21.3	19.4	44.7	43.5
	Saliency+Cont. Att. (Weight Sh.)	62.0	43.8	30.0	20.5	19.7	44.6	43.3
	Saliency+Context Attention	61.3	43.3	30.1	20.9	20.2	45.0	44.5
Flickr30k (Test)	Saliency Pooling	58.3	40.6	27.5	18.6	18.7	43.0	36.2
	Attention on Saliency	62.5	44.2	30.5	21.0	19.6	44.9	45.0
	Saliency+Cont. Att. (Weight Sh.)	61.7	43.7	30.0	20.4	19.6	44.2	43.1
	Saliency+Context Attention	61.5	43.8	30.5	21.3	20.0	45.2	46.4
PASCAL-50S	Saliency Pooling	79.9	67.1	53.6	41.8	31.4	64.1	65.3
	Attention on Saliency	82.4	70.3	57.4	45.5	32.7	66.3	70.2
	Saliency+Cont. Att. (Weight Sh.)	82.0	69.7	56.4	44.2	32.7	65.2	70.0
	Saliency+Context Attention	82.4	70.2	57.5	45.7	32.9	66.3	70.7

While the use of machine attention strategies is beneficial (see Saliency Pooling vs. Attention on Saliency), saliency and context are both important for captioning. The use of different attention paths for saliency and context also enhances the performance (see Saliency+Context Attention (With Weight Sharing) vs. Saliency+Context Attention).

Table 3 shows the results on the three considered datasets in terms of BLEU@4, METEOR, ROUGE_L, and CIDEr. We compare our solutions to both versions of the model presented in [52]. The GBVS version exploits saliency maps calculated by using a traditional bottom-up model [15], while the other one includes saliency maps extracted from a deep convolutional network [51].

Overall, results show that the proposed Saliency and Context Attention model can overcome the other methods on different metrics, thus confirming the strategy of including two attention

Table 3. Comparison with Existing Saliency-Boosted Captioning Models

Dataset	Model	B@4	METEOR	ROUGE _L	CIDE _r
SALICON	Sugano et al. [48]	24.5	21.9	52.4	63.8
	Ours	26.9	22.9	50.4	73.3
COCO	Tavakoli et al. [52] (GBVS)	28.7	23.5	51.2	84.1
	Tavakoli et al. [52] (iSEEL)	28.3	23.5	50.8	83.6
	Ours	28.4	24.8	52.1	89.8
Flickr30k (Test)	Ramanishka et al. [41]	—	18.3	—	—
	Ours	21.3	20.0	45.2	46.4
PASCAL-50S	Tavakoli et al. [52] (GBVS)	40.0	30.2	63.5	61.5
	Tavakoli et al. [52] (iSEEL)	39.6	30.2	63.2	61.4
	Ours	45.7	32.9	66.3	70.7

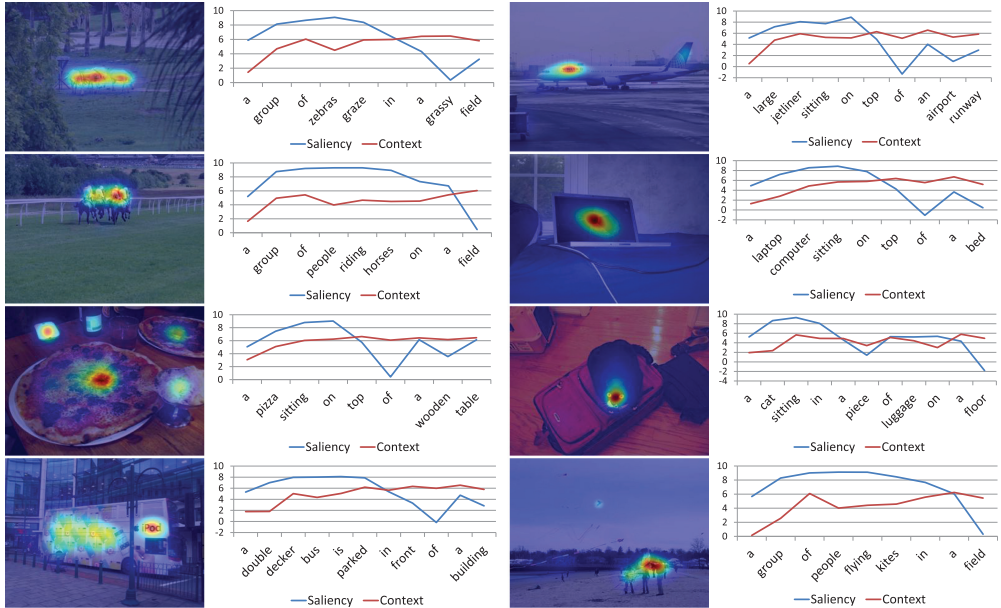


Fig. 6. Examples of attention weight changes between saliency and context along with the generation of captions (best seen in color). Images are from the Microsoft COCO dataset [33].

paths. In particular, on the METEOR metric, we obtain a relative improvement of 4.57% on the SALICON dataset, 5.53% on the Microsoft COCO, and 8.94% on the PASCAL-50S.

5.5 Analysis of Generated Captions

We further collect statistics on captions generated by our method and by the Soft Attention model, to quantitatively assess the quality of generated captions. First, we define three metrics that evaluate the vocabulary size and the difference between the corpus of captions generated by the two models and the ground truth:

Table 4. Statistics on Vocabulary Size and Diversity of the Generated Captions

Dataset	Model	Div-1	Div-2	Vocabulary	% Novel Sent.	% Different Sent.
SALICON	Soft Attention	0.0136	0.0498	658	95.22%	95.34%
	Saliency+Context Attention	0.0125	0.0549	614	93.12%	
COCO	Soft Attention	0.0038	0.0187	1490	81.81%	93.80%
	Saliency+Context Attention	0.0037	0.0182	1444	78.02%	
Flickr8k (Validation)	Soft Attention	0.0367	0.1026	389	98.30%	97.90%
	Saliency+Context Attention	0.0400	0.1094	411	99.30%	
Flickr8k (Test)	Soft Attention	0.0385	0.1041	404	98.50%	97.60%
	Saliency+Context Attention	0.0419	0.1119	423	99.60%	
Flickr30k (Validation)	Soft Attention	0.0577	0.1445	699	99.90%	98.62%
	Saliency+Context Attention	0.0565	0.1439	715	99.61%	
Flickr30k (Test)	Soft Attention	0.0580	0.1508	682	99.90%	98.20%
	Saliency+Context Attention	0.0585	0.1549	711	99.70%	
PASCAL-50S	Soft Attention	0.0475	0.1379	465	97.10%	94.80%
	Saliency+Context Attention	0.0468	0.1359	456	96.40%	

Including saliency and context in two different machine attention paths (Saliency+Context Attention) produces different captions with respect to the traditional machine attention approach (Soft Attention), while preserving almost the same diversity statistics.

- *Vocabulary size*: number of unique words generated in all captions
- *Percentage of novel sentences*: percentage of generated sentences that are not seen in the training set
- *Percentage of different sentences*: percentage of images that are described differently by the two models

Then, we measure the diversity of the set of captions generated by each of the two models, via the following two metrics [45]:

- *Div-1*: ratio of number of unique unigrams in a set of captions to the number of words in the same set. Higher is more diverse.
- *Div-2*: ratio of number of unique bigrams in a set of captions to the number of words in the same set. Higher is more diverse.

In Table 4, we compare the set of captions generated by our model with that generated by the Soft Attention baseline. Although our model features a slight reduction of the vocabulary size on SALICON, COCO, and PASCAL-50S, captions generated by the two models are very often different, thus confirming that the two approaches have learned different captioning models. Moreover, the diversity and the number of novel sentences of the Soft Attention approach are entirely preserved.

5.6 Analysis of Attentive States

The selection of a location in our model is based on the competition between the saliency attentive path and the context attentive path (see Equation (9)). To investigate how the two paths interact and contribute to the generation of a word, in Figure 6 we report, for several images from the Microsoft COCO dataset, the changes in attention weights between the two paths. Specifically, for

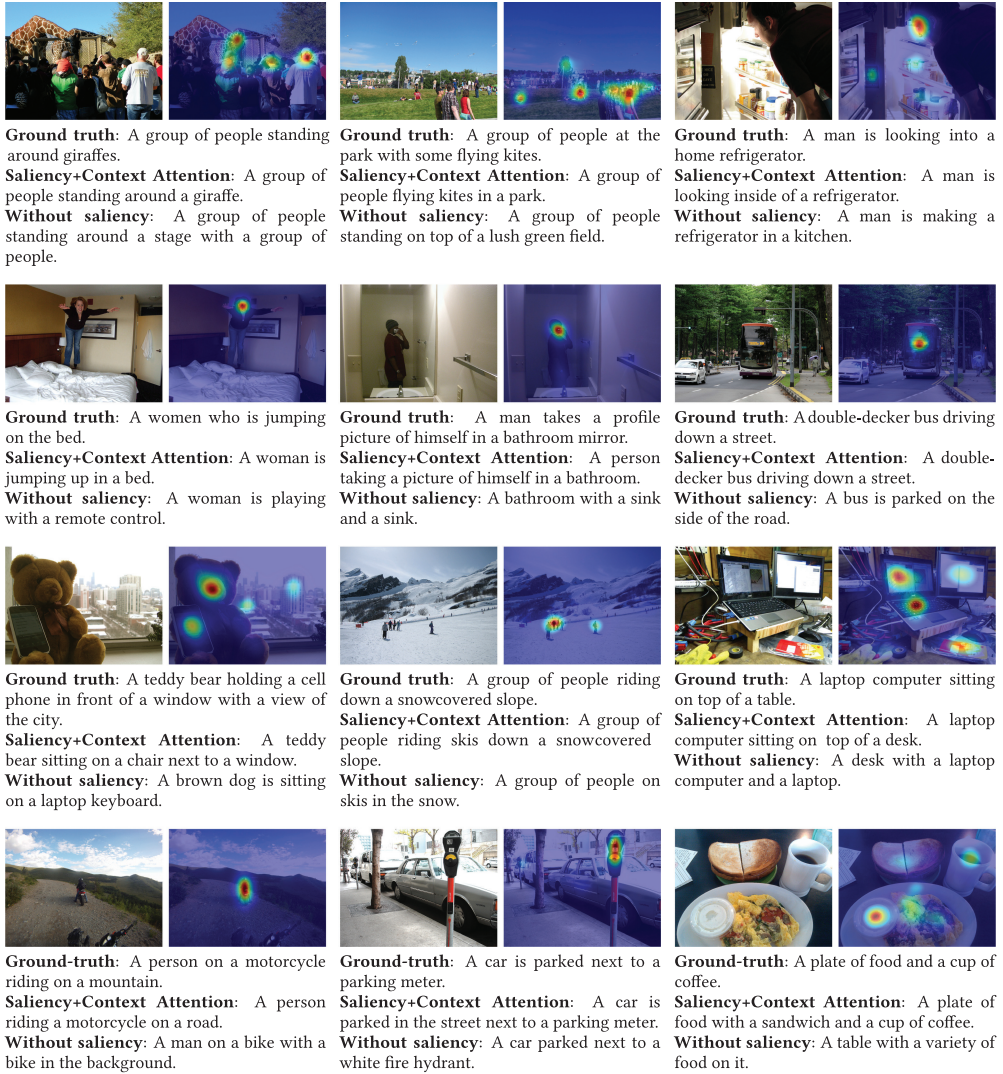


Fig. 7. Example results on the Microsoft COCO dataset [33].

each image, we report the average of e_{ti}^{sal} and e_{ti}^{ctx} values at each timestep, along with a visualization of its saliency map. It is interesting to see how the model was able to correctly exploit the two attention paths for generating different parts of the caption, and how generated words correspond in most cases to the attended regions. For example, in the case of the first image (“a group of zebras graze in a grassy field”), the saliency attentive path is more active than the context path during the generation of words corresponding to the “group of zebras,” which is captured by saliency. Instead, when the model has to describe the context (“in a grassy field”), the saliency attentive path has lower weights with respect to the context attentive path. The same can be observed for all the reported images; it can also be noticed that generated captions tend to describe both salient objects and the context, and that usually the salient part, which is also the most important, is described before the context.

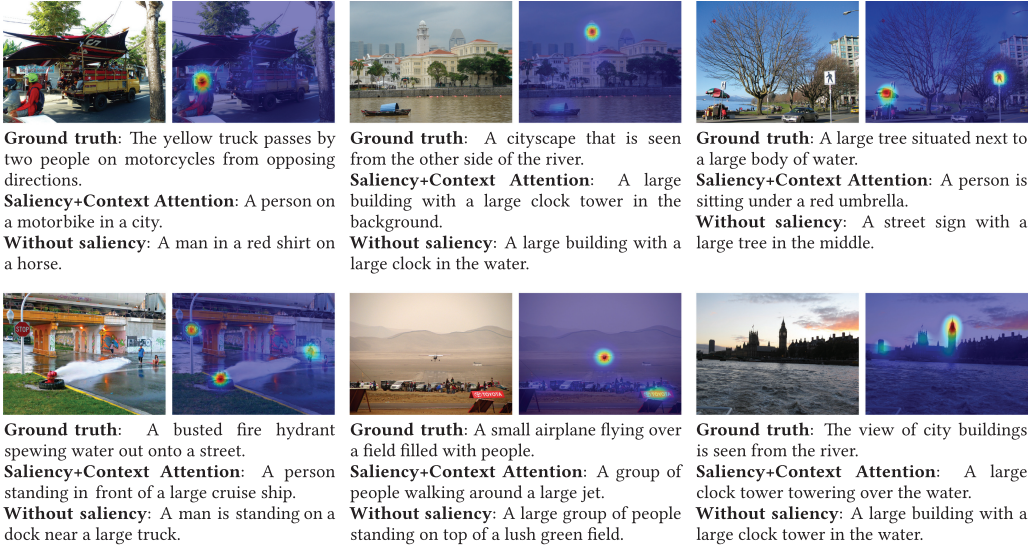


Fig. 8. Failure cases on sample images of the Microsoft COCO dataset [33].

5.7 Qualitative Results

Finally, in Figure 7, we report some sample results on images taken from the Microsoft COCO dataset. For each image, we report the corresponding saliency map and captions generated by our model and by the Soft Attention baseline compared to the ground truth. It can be seen that, on average, captions generated by our model are more consistent with the corresponding image and the human-generated caption, and that, as also observed in the previous section, salient parts are described as well as the context. The incorporation of saliency and context also help the model to avoid failures due to hallucination, such as in the case of the fourth image, in which the Soft Attention model predicts a remote control that is not depicted in the image. Other failure cases, which are avoided by our model, include the repetition of words (as in the fifth image) and the failure to describe the context (first image). We speculate that the presence of two separate attention paths, which the model has learned to attend to during the generation of the caption, helps to avoid such failures more effectively than the classic machine attention approach.

For completeness, some failure cases of the proposed model are reported in Figure 8. The majority of failures occur when the salient regions of the image are not described in the corresponding ground-truth caption (as, for example, in the first row), thus causing a performance loss. Some problems arise also in the presence of complex scenes (such as in the fourth image). However, we observe that the Soft Attention baseline fails as well to predict correct and complete captions in these cases.

6 CONCLUSION

We proposed a novel image captioning architecture that extends the machine attention paradigm by creating two attentive paths conditioned on the output of a saliency prediction model. The first one is focused on salient regions, and the second on contextual regions: the overall model exploits the two paths during the generation of the caption by giving more importance to salient or contextual regions as needed. The role of saliency with respect to context has been investigated by collecting statistics on semantic segmentation datasets, while the captioning model has been

evaluated on large-scale captioning datasets using standard automatic metrics and by evaluating the diversity and the dictionary size of the generated corpora. Finally, the activations of the two attentive paths have been investigated, and we have shown that they correspond, word by word, to a focus on salient objects or on the context in the generated caption; moreover, we qualitatively assessed the superiority of the captions generated by our method with respect to those generated by the Soft Attention approach. Although our focus has been that of demonstrating the effectiveness of saliency on captioning, rather than that of beating captioning approaches that rely on different cues, we point out that our method can be easily incorporated into those architectures.

REFERENCES

- [1] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *International Conference on Learning Representations*.
- [2] Satyanjeev Banerjee and Alon Lavie. 2005. METEOR: An automatic metric for MT evaluation with improved correlation with human judgments. In *ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization*.
- [3] Lorenzo Baraldi, Costantino Grana, and Rita Cucchiara. 2017. Hierarchical boundary-aware neural encoder for video captioning. In *IEEE International Conference on Computer Vision and Pattern Recognition*.
- [4] Ali Borji. 2012. Boosting bottom-up and top-down visual features for saliency estimation. In *IEEE International Conference on Computer Vision and Pattern Recognition*.
- [5] Zoya Bylinskii, Tilke Judd, Ali Borji, Laurent Itti, Frédo Durand, Aude Oliva, and Antonio Torralba. 2017. MIT Saliency Benchmark. Retrieved from <http://saliency.mit.edu/>.
- [6] Zoya Bylinskii, Adrià Recasens, Ali Borji, Aude Oliva, Antonio Torralba, and Frédo Durand. 2016. Where should saliency models look next? In *European Conference on Computer Vision*.
- [7] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. 2016. The cityscapes dataset for semantic urban scene understanding. In *IEEE International Conference on Computer Vision and Pattern Recognition*.
- [8] Marcella Cornia, Lorenzo Baraldi, Giuseppe Serra, and Rita Cucchiara. 2016. A deep multi-level network for saliency prediction. In *International Conference on Pattern Recognition*.
- [9] Marcella Cornia, Lorenzo Baraldi, Giuseppe Serra, and Rita Cucchiara. 2017. Predicting human eye fixations via an LSTM-based saliency attentive model. *arXiv preprint arXiv:1611.09571* (2017).
- [10] Mark Everingham, Luc Van Gool, Christopher K. I. Williams, John Winn, and Andrew Zisserman. 2010. The pascal visual object classes (VOC) challenge. *International Journal of Computer Vision* 88, 2 (2010), 303–338.
- [11] Ali Farhadi, Mohsen Hejrati, Mohammad Amin Sadeghi, Peter Young, Cyrus Rashtchian, Julia Hockenmaier, and David Forsyth. 2010. Every picture tells a story: Generating sentences from images. In *European Conference on Computer Vision*.
- [12] Xavier Glorot and Yoshua Bengio. 2010. Understanding the difficulty of training deep feedforward neural networks. In *International Conference on Artificial Intelligence and Statistics*.
- [13] Ke Gong, Xiaodan Liang, Xiaohui Shen, and Liang Lin. 2017. Look into person: Self-supervised structure-sensitive learning and a new benchmark for human parsing. In *IEEE International Conference on Computer Vision and Pattern Recognition*.
- [14] Zenzi M. Griffin and Kathryn Bock. 2000. What the eyes say about speaking. *Psychological Science* 11, 4 (2000), 274–279.
- [15] Jonathan Harel, Christof Koch, and Pietro Perona. 2006. Graph-based visual saliency. In *Advances in Neural Information Processing Systems*.
- [16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *IEEE International Conference on Computer Vision and Pattern Recognition*.
- [17] Micah Hodosh, Peter Young, and Julia Hockenmaier. 2013. Framing image description as a ranking task: Data, models and evaluation metrics. *Journal of Artificial Intelligence Research* 47 (2013), 853–899.
- [18] Xun Huang, Chengyao Shen, Xavier Boix, and Qi Zhao. 2015. SALICON: Reducing the semantic gap in saliency prediction by adapting deep neural networks. In *IEEE International Conference on Computer Vision*.
- [19] Laurent Itti, Christof Koch, and Ernst Niebur. 1998. A model of saliency-based visual attention for rapid scene analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 20, 11 (1998), 1254–1259.
- [20] Saumya Jetley, Naila Murray, and Eleonora Vig. 2016. End-to-end saliency mapping via probability distribution prediction. In *IEEE International Conference on Computer Vision and Pattern Recognition*.

- [21] Ming Jiang, Shengsheng Huang, Juanyong Duan, and Qi Zhao. 2015. SALICON: Saliency in context. In *IEEE International Conference on Computer Vision and Pattern Recognition*.
- [22] Justin Johnson, Andrej Karpathy, and Li Fei-Fei. 2016. Densecap: Fully convolutional localization networks for dense captioning. In *IEEE International Conference on Computer Vision and Pattern Recognition*.
- [23] Tilke Judd, Krista Ehinger, Frédo Durand, and Antonio Torralba. 2009. Learning to predict where humans look. In *IEEE International Conference on Computer Vision*.
- [24] Andrej Karpathy and Li Fei-Fei. 2015. Deep visual-semantic alignments for generating image descriptions. In *IEEE International Conference on Computer Vision and Pattern Recognition*.
- [25] Diederik Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [26] Jonathan Krause, Justin Johnson, Ranjay Krishna, and Li Fei-Fei. 2017. A hierarchical approach for generating descriptive image paragraphs. In *IEEE International Conference on Computer Vision and Pattern Recognition*.
- [27] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. 2012. ImageNet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems*.
- [28] Srinivas S. S. Kruthiventi, Vennela Gudisa, Jaley H. Dholakiya, and R. Venkatesh Babu. 2016. Saliency unified: A deep architecture for simultaneous eye fixation prediction and salient object segmentation. In *IEEE International Conference on Computer Vision and Pattern Recognition*.
- [29] Girish Kulkarni, Visruth Premraj, Vicente Ordonez, Sagnik Dhar, Siming Li, Yejin Choi, Alexander C. Berg, and Tamara L. Berg. 2013. Babytalk: Understanding and generating simple image descriptions. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 35, 12 (2013), 2891–2903.
- [30] Matthias Kümmerer, Lucas Theis, and Matthias Bethge. 2015. DeepGaze I: Boosting saliency prediction with feature maps trained on ImageNet. In *International Conference on Learning Representations Workshops*.
- [31] Matthias Kümmerer, Thomas S. A. Wallis, and Matthias Bethge. 2016. DeepGaze II: Reading fixations from deep features trained on object recognition. *arXiv preprint arXiv:1610.01563* (2016).
- [32] Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *ACL Workshop on Text Summarization Branches Out*.
- [33] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. 2014. Microsoft COCO: Common objects in context. In *European Conference on Computer Vision*.
- [34] Nian Liu and Junwei Han. 2016. A deep spatial contextual long-term recurrent convolutional network for saliency detection. *arXiv preprint arXiv:1610.01708* (2016).
- [35] Nian Liu, Junwei Han, Dingwen Zhang, Shifeng Wen, and Tianming Liu. 2015. Predicting eye fixations using convolutional neural networks. In *IEEE International Conference on Computer Vision and Pattern Recognition*.
- [36] Roozbeh Mottaghi, Xianjie Chen, Xiaobai Liu, Nam-Gyu Cho, Seong-Whan Lee, Raquel Urtasun, and Alan Yuille. 2014. The role of context for object detection and semantic segmentation in the wild. In *IEEE International Conference on Computer Vision and Pattern Recognition*.
- [37] Vicente Ordonez, Girish Kulkarni, and Tamara L. Berg. 2011. Im2text: Describing images using 1 million captioned photographs. In *Advances in Neural Information Processing Systems*. 1143–1151.
- [38] Junting Pan, Cristian Canton, Kevin McGuinness, Noel E. O'Connor, Jordi Torres, Elisa Sayrol, and Xavier Giro-i Nieto. 2017. SalGAN: Visual saliency prediction with generative adversarial networks. In *IEEE International Conference on Computer Vision and Pattern Recognition Workshops*.
- [39] Junting Pan, Kevin McGuinness, Elisa Sayrol, Noel O'Connor, and Xavier Giró-i Nieto. 2016. Shallow and deep convolutional networks for saliency prediction. In *IEEE International Conference on Computer Vision and Pattern Recognition*.
- [40] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. BLEU: A method for automatic evaluation of machine translation. In *40th Annual Meeting on Association for Computational Linguistics*.
- [41] Vasili Ramanishka, Abir Das, Jianming Zhang, and Kate Saenko. 2017. Top-down visual saliency guided by captions. In *IEEE International Conference on Computer Vision and Pattern Recognition*.
- [42] Cyrus Rashtchian, Peter Young, Micah Hodosh, and Julia Hockenmaier. 2010. Collecting image annotations using Amazon's mechanical turk. In *NAACL HLT Workshops*.
- [43] Ronald A. Rensink. 2000. The dynamic representation of scenes. *Visual Cognition* 7, 1–3 (2000), 17–42.
- [44] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. 2015. ImageNet large scale visual recognition challenge. *International Journal of Computer Vision* 115, 3 (2015), 211–252.
- [45] Rakshith Shetty, Marcus Rohrbach, Lisa Anne Hendricks, Mario Fritz, and Bernt Schiele. 2017. Speaking the same language: Matching machine to human captions by adversarial training. *arXiv preprint arXiv:1703.10476* (2017).
- [46] Karen Simonyan and Andrew Zisserman. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556* (2014).

- [47] Richard Socher, Andrej Karpathy, Quoc V. Le, Christopher D. Manning, and Andrew Y. Ng. 2014. Grounded compositional semantics for finding and describing images with sentences. *Transactions of the Association for Computational Linguistics* 2 (2014), 207–218.
- [48] Yusuke Sugano and Andreas Bulling. 2016. Seeing with humans: Gaze-assisted neural image captioning. *arXiv preprint arXiv:1608.05203* (2016).
- [49] Ilya Sutskever, James Martens, George Dahl, and Geoffrey Hinton. 2013. On the importance of initialization and momentum in deep learning. In *International Conference on Machine Learning*.
- [50] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. 2015. Going deeper with convolutions. In *IEEE International Conference on Computer Vision and Pattern Recognition*.
- [51] Hamed R. Tavakoli, Ali Borji, Jorma Laaksonen, and Esa Rahtu. 2017. Exploiting inter-image similarity and ensemble of extreme learners for fixation prediction using deep features. *Neurocomputing* 244 (2017), 10–18.
- [52] Hamed R. Tavakoli, Rakshith Shetty, Ali Borji, and Jorma Laaksonen. 2017. Paying attention to descriptions generated by image captioning models. In *IEEE International Conference on Computer Vision*.
- [53] Ramakrishna Vedantam, C. Lawrence Zitnick, and Devi Parikh. 2015. Cider: Consensus-based image description evaluation. In *IEEE International Conference on Computer Vision and Pattern Recognition*.
- [54] Eleonora Vig, Michael Dorr, and David Cox. 2014. Large-scale optimization of hierarchical features for saliency prediction in natural images. In *IEEE International Conference on Computer Vision and Pattern Recognition*.
- [55] Oriol Vinyals, Alexander Toshev, Samy Bengio, and Dumitru Erhan. 2015. Show and tell: A neural image caption generator. In *IEEE International Conference on Computer Vision and Pattern Recognition*.
- [56] Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhudinov, Rich Zemel, and Yoshua Bengio. 2015. Show, attend and tell: Neural image caption generation with visual attention. In *International Conference on Machine Learning*.
- [57] Yezhou Yang, Ching Lik Teo, Hal Daumé III, and Yiannis Aloimonos. 2011. Corpus-guided sentence generation of natural images. In *Conference on Empirical Methods in Natural Language Processing*.
- [58] Benjamin Z. Yao, Xiong Yang, Liang Lin, Mun Wai Lee, and Song-Chun Zhu. 2010. I2t: Image parsing to text description. *Proceedings of the IEEE* 98, 8 (2010), 1485–1508.
- [59] Quanzeng You, Hailin Jin, Zhaowen Wang, Chen Fang, and Jiebo Luo. 2016. Image captioning with semantic attention. In *IEEE International Conference on Computer Vision and Pattern Recognition*.
- [60] Peter Young, Alice Lai, Micah Hodosh, and Julia Hockenmaier. 2014. From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions. *Transactions of the Association for Computational Linguistics* 2 (2014), 67–78.
- [61] Jianming Zhang and Stan Sclaroff. 2013. Saliency detection: A boolean map approach. In *IEEE International Conference on Computer Vision*.

Received November 2017; accepted December 2017