Seeing with Humans: Gaze-Assisted Neural Image Captioning

Yusuke Sugano and Andreas Bulling

Abstract—Gaze reflects how humans process visual scenes and is therefore increasingly used in computer vision systems. Previous works demonstrated the potential of gaze for object-centric tasks, such as object localization and recognition, but it remains unclear if gaze can also be beneficial for scene-centric tasks, such as image captioning. We present a new perspective on gaze-assisted image captioning by studying the interplay between human gaze and the attention mechanism of deep neural networks. Using a public large-scale gaze dataset, we first assess the relationship between state-of-the-art object and scene recognition models, bottom-up visual saliency, and human gaze. We then propose a novel split attention model for image captioning. Our model integrates human gaze information into an attention-based long short-term memory architecture, and allows the algorithm to allocate attention selectively to both fixated and non-fixated image regions. Through evaluation on the COCO/SALICON datasets we show that our method improves image captioning performance and that gaze can complement machine attention for semantic scene understanding tasks.

Index Terms—Eye tracking, Attention, Image Captioning

✦

1 INTRODUCTION

Human gaze reflects processes of cognition and perception and therefore represents a rich source of information about the observer of a visual scene. Consequently, gaze has successfully been used for tasks such as eye-based user modeling [1], [2], [3], [4], [5] and opened up new opportunities to further advance human-machine collaboration – collaborative human-machine vision systems in which part of the processing is carried out by the machine while another part is performed by a human and conveyed to the machine via gaze information.

As eye tracking techniques mature and become integrated into daily-life devices such as smart glasses, it is becoming more realistic to assume the availability of additional gaze information. Recent advances in crowd-sourcing techniques [6], [7], [8] and appearance-based estimation methods [9], [10], [11] are also further paving the way for low-cost and large-scale gaze data collection. For these reasons, gaze-enabled computer vision methods have been attracting increasing interest in recent years. Prior work has typically focused on object-centric tasks, such as object localization [5], [12] or recognition [3], [13]. Although gaze information is potentially even richer for tasks that require holistic scene understanding [14], [15], the integration of gaze information for scene-centric computer vision algorithms, such as for image captioning, has not yet been explored.

At the same time, attention mechanisms – which are inspired by how humans selectively attend to visual input – have recently become popular in machine learning literature [16], [17]. Models that include attention mechanisms

- Y. Sugano and A. Bulling are with the Perceptual User Interfaces Group, Max Planck Institute for Informatics, Saarbrücken, Germany.
  E-mail: {sugano, bulling}@mpi-inf.mpg.de
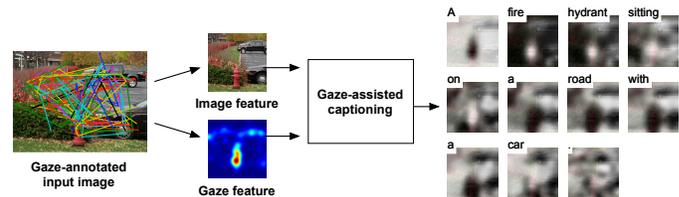
Fig. 1. Our method takes gaze-annotated images as input, and uses both human gaze and bottom-up visual features for attention-based captioning.

were shown to improve performance and efficiency for a variety of computer vision tasks, such as facial expression recognition [18] or image captioning [19]. However, given that most current attention models are trained without any supervision from human attention, it remains unclear whether gaze information can further improve model performance.

This gap between human gaze and the machine learning algorithms' own feature localization capability is a fundamental question for the role of human gaze in computer vision tasks. There has been remarkable progress in learning-based visual saliency models that demonstrated good performance for predicting where humans look in a purely bottom-up manner [20]. Also, state-of-the-art convolutional neural nets (CNNs) were reported to be able to efficiently localize key points in visual scenes [21], indicating that machines can be better at task-specific feature localization than humans in some scenarios. It remains unclear, however, whether human gaze can complement bottom-up visual information for feature localization.

The goal of this work is to shed some light on these questions and to explore how the performance of image captioning models can be improved by incorporating gaze information. Using the SALICON dataset [6], we first compare localization capability of state-of-the-art object and scene recognition models with human gaze. Inspired by this analysis, we then propose a novel image captioning model (see Figure 1) that integrates gaze information into a state-of-the-art long short-term memory (LSTM) architecture with an attention mechanism [19]. Human gaze is represented as a set of fixations, i.e., the static states of gaze upon a specific location, and the model localizes its machine attention selectively to both fixated and non-fixated regions.

The main contributions of this work are twofold. First, we present an analysis of the relationship between object and scene recognition models and human gaze. We take state-of-the-art CNN models as examples, and discuss how human gaze can influence their performance. Second, we present a novel gaze-assisted attention mechanism for image captioning that is based on a split attention model. We show that the proposed model improves captioning performance of the baseline attention model [19]. To the best of our knowledge, this is the first work to 1) propose an actual model for gaze-assisted image captioning, and 2) relate human gaze input to deep neural network models of object recognition, scene recognition, and image captioning. In this manner, this work provides the first unified overview on gaze, machine attention, image captioning and deep object/scene recognition models.

2 RELATED WORK

Our gaze-assisted neural image captioning method is related to previous works in the emerging domain of collaborative human-machine vision as well as to attention mechanisms in current deep learning methods.

Gaze for Object Localization and Recognition

The most well-studied approach is to use gaze information as a cue to infer object locations. These efforts are motivated by findings from vision research that showed that humans tend to fixate on important or salient objects in the visual scene [22], [23]. The correlation between fixations and objects is also becoming important in the context of saliency prediction [24], [25]. There have been several methods proposed for robust gaze-guided object segmentation [12], [26], [27], and based on a similar assumption fixation information is also used for localizing important objects in videos [5], [28], [29]. While these works mainly focused on object localization and the recognition was assumed to be purely vision-based, there have been several attempts to further use gaze information for object recognition. Gaze-related features were used for object recognition [13], or incorporated into the image-based object recognition pipeline [3], [30], [31]. However, the tasks were still object-centric and recognition tasks beyond object bounding boxes were not explored.

Gaze for Scene Description

The usefulness of gaze information for image recognition is not limited to object-centric tasks. For example, fixation locations were shown to provide information about the visual scene that can be used for both first-person and third-person activity recognition [2], [32]. However, only a few previous works explored the link between gaze and scene descriptions in the context of computer vision. Subramanian et al. identified correlations between gaze and scene semantics and showed that saccadic eye movements between different object regions can be used to discover object-object relations [33]. Yun et al. showed that such correlations even exist in free-viewing conditions, i.e. without any specific tasks, and between image descriptions provided by different observers [14], [30]. Coco and Keller further provided a detailed analysis on what kinds of visual guidance mechanisms cause such correlations [34], and they demonstrated that annotators' gaze information can be used to predict the given captions [35]. However, all of these works provided only basic analyses of the link between gaze, visual feature, and scene descriptions and did not integrate gaze information into image captioning algorithms.

Attention Mechanisms in Deep Learning

An increasing number of works investigate the potential of attention mechanisms for deep neural networks for computer vision tasks. Attention mechanisms were shown to reduce computation and to make the networks more robust to changes in input resolution [16], [17], [36]. Even more importantly, attention mechanisms add localization capabilities to neural networks, resulting in improved performance for target localization and recognition tasks [18], [37], [38]. Information localization capabilities are also important for tasks such as machine translation [39] or semantic description [40].

With image captioning having recently emerged as a core task for deep neural networks [41], [42], [43], [44], [45], [46], [47], [48], this is where a promising link with human attention arises. The benefit of the attention mechanism for image captioning has been demonstrated by Xu et al. [19], and You et al. proposed to use the attention mechanism in the semantic space [49]. Although the visual attention mechanism of humans is often referred to as the inspirational source for these methods, none of them related their attention mechanism to actual human gaze data. Instead, the attention mechanisms were treated as pure machine optimization tasks and were neither designed nor evaluated to resemble human attentive behavior.

3 GAZE FOR OBJECT AND SCENE RECOGNITION

We first conduct fundamental analyses on the relationship between gaze and image recognition models on the SAL-ICON dataset [6]. Object/scene category classification is key for holistic image understanding and our goal is to quantify whether and how human gaze can help state-of-the-art classification models. In addition, we also compare human gaze and bottom-up visual saliency using the boolean map saliency (BMS) algorithm [50], which is one of the best-performing saliency models with a publicly available implementation. Extending on prior work [30], [33], we evaluate the localization capabilities of state-of-the-art object recognition models in comparison with human gaze and bottom-up saliency, and provide the first analysis of scene recognition models. We take the 16-layer VGGNet [51] architecture as an example and discuss the correlation with human gaze for both object and scene recognition. For object recognition, we use Simonyan et al.'s pre-trained model [51] on the ILSVRC-2012 dataset [52]. Similarly, Wang et al.'s pre-trained model [53] on the Places205 dataset [54] is used for scene recognition.

3.1 Dataset

The SALICON dataset provides gaze data collected through a crowd-sourced eye tracking experiment based on a mouse-contingent paradigm. The authors quantitatively showed that such mouse tracking data can resemble real eye tracking data under free viewing conditions. The currently public part of the SALICON dataset contains 10,000 and 5,000 images from the training and validation sets of the COCO dataset, respectively.

Given that the SALICON dataset is a subset of the Microsoft COCO dataset [55], ground-truth gaze annotations can be obtained from the original annotations. In addition, we assigned the original image annotations (captions, object segmentations, and categories) from the COCO dataset according to their image file names. To associate the annotated 80 object categories to the ILSVRC object categories, we used the WordNet synset IDs corresponding to the COCO object categories given by Chao et al. [56]. We associated *all* object categories in the ILSVRC dataset that are children of the COCO categories in the WordNet hierarchy.

To study the relationship between image semantics and fixation locations in more detail, we further divided the object categories into two disjoint sets depending on whether

(a) Object recognition

(b) Scene recognition
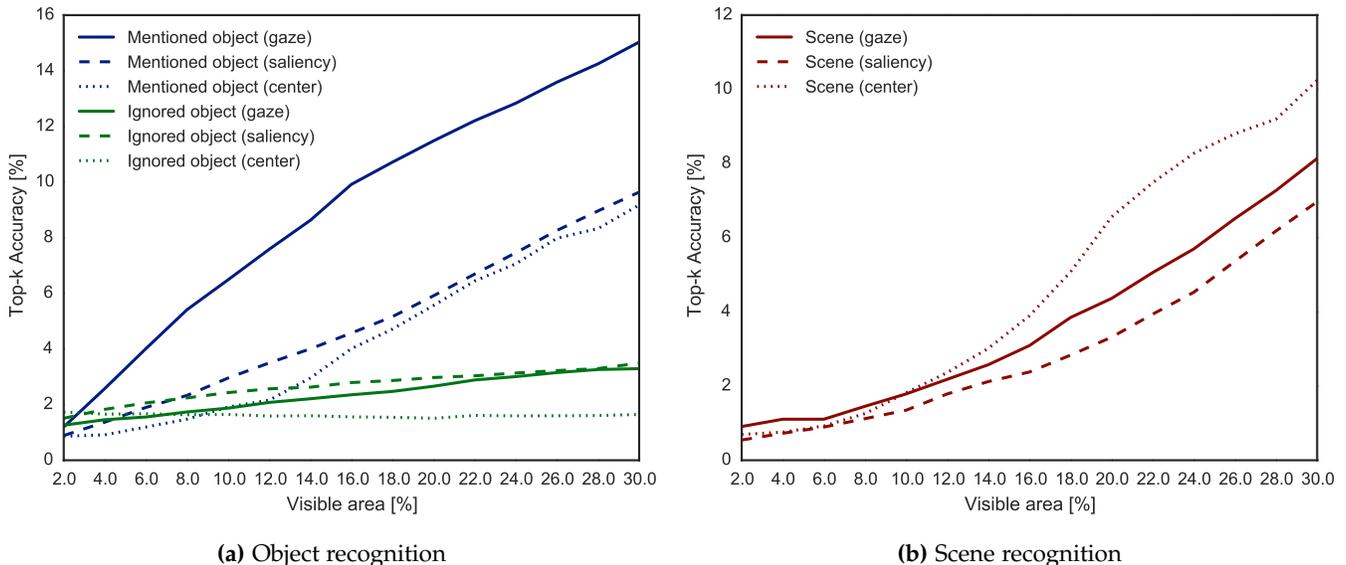
Fig. 2. Top-k accuracy for object and scene recognition. The horizontal axis indicates the ratio of the visible area given by thresholding the fixation, saliency and center maps. k is set to the number of labels associated with each image.

they were mentioned in the captions. We used the NLTK library [57] to extract all nouns from the captions. To reduce the effect of errors in the natural language processing, we only took nouns appearing more than twice in different captions. On average, this process yielded 4.02 nouns per image. Then, all possible WordNet synsets corresponding to these nouns were extracted and associated with all child ILSVRC object categories in the same manner. The object categories obtained from the COCO annotations were divided into two subsets, referred to as *mentioned* and *ignored* below, depending on whether the same synset was obtained from the captions. While the *mentioned* subset corresponds to the ILSVRC object IDs found both in annotations and captions, the *ignored* subset contains IDs whose parent COCO ID had no correspondence in the caption ID set. Given that the COCO dataset does not provide scene labels, we assigned scene categories according to the set of nouns extracted from the captions. For some scene category names whose correspondences to WordNet synsets were ambiguous, we manually assigned IDs. Finally, we extracted scene categories which have exactly the same synset in the set of nouns.

3.2 Classification Performance with Fixation Masks

Figure 2 shows the mean top-k classification accuracy of object (both mentioned and ignored) and scene categories, when fixation/saliency maps were used. Since in our setting each image can have multiple ground-truth object/scene labels, we set k to the number of labels associated with each image. The horizontal axis indicates the ratio of the visible area made by thresholding saliency/fixation maps, and the vertical axis shows classification scores where the model can only observe fixated or salient regions. As a baseline, Figure 2 also shows an importance value plot with the same visible area ratio at the center of the input image. We show performance values for both mentioned and ignored object categories independently.

In general, fixation and saliency information is more closely related to the object classification performance for the mentioned rather than ignored objects in Figure 2. It can also be seen that fixation maps achieve significantly better classification performance for mentioned objects than saliency maps. This indicates that the performance gap between gaze and saliency is still significant in terms of recall. It is in general more difficult for saliency models to suppress false positives, and the prediction result is not selective enough to support semantic image recognition. If the ground-truth mentioned object segmentation is used for masking, the classification scores of mentioned and ignored objects become 24.50% and 3.54%, respectively. Therefore, although the fixation mask cannot cover all of the important locations to discover all important objects, it can have roughly the same discrimination capability for non-important objects as ground-truth masks. In contrast, for scene classification, the score is significantly lower at the beginning but becomes higher if the visible area is expanded. While the score for fixation maps is consistently higher than for saliency maps, the center map baseline shows better performance than both fixation and saliency maps when the visible area is expanded to around 15%.

In addition, Figure 3 shows some examples of the feature importance maps for object and scene recognition models. Following the approach of Zeiler et al. [58], we evaluated the importance of image regions by measuring the decrease in recognition performance when the region is masked. If the masked region contains important visual information for recognizing the target label, the recognition score should decrease substantially. Hence, the feature importance can be measured via a negative value of the score decrement.

The first column shows the ground-truth segmentations of mentioned objects obtained from the original COCO dataset annotation. The second and third column show importance maps of the object and scene recognition models, respectively. The fourth column shows human fixation maps

Object (GT)	Object	Scene	Fixation	Saliency
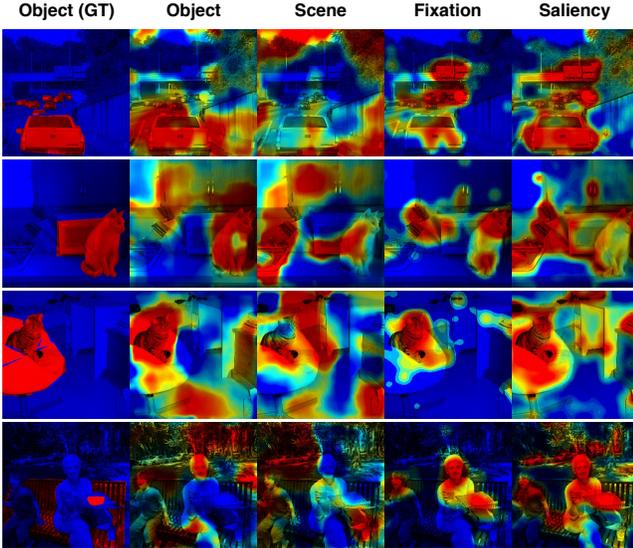
Fig. 3. Comparison of feature importance maps. Mean maps over all corresponding labels are overlaid onto the image with a color coding from blue (lowest importance) to red (highest importance). For better visual comparison, all maps were histogram equalized.

taken from the SALICON dataset, and the fifth column shows the corresponding purely bottom-up saliency prediction results. While there is generally a larger similarity between fixation and object recognition maps than scene recognition maps, there are some object categories that do not attract human attention, such as the bench in the last image, even if the object is semantically important.

From these results, we can make several important observations. First, fixation positions are indeed related to important locations for object recognition models to find semantically important objects. By focusing on fixated regions, object recognition models can potentially discriminate between mentioned and ignored objects. The gap between human fixation and bottom-up saliency prediction is also related to this point, and human fixation gives better localization of important features. Second, fixation positions are not significantly related to important locations for scene recognition models, and the area of focus has to be extended to find relevant information.

4 GAZE-ASSISTED IMAGE CAPTIONING

The previous analysis provides valuable insights into how gaze information can be exploited for semantic scene understanding and image captioning tasks. While human gaze can help object recognition models to discriminate semantically important objects from non-important objects, the model also needs to pay attention to objects that do not attract human gaze. Scene categories also cannot be fully recovered only from fixated regions even if they are important for semantic description. Therefore, unlike humans, machines also need to obtain visual features from background regions to recognize scene categories. Based on these observations, we propose to use gaze information to guide the attention-based captioning architecture [19] so that the model can allocate attention selectively according to fixation distribution.

We first briefly summarize Xu et al.'s model based on the standard LSTM architecture [59] and the soft attention mechanism [39]. The input image is encoded as a set of L feature vectors $a = \{\boldsymbol{a}_1, \dots, \boldsymbol{a}_L\}$ extracted from the last convolutional layer of an object recognition CNN. Each $\boldsymbol{a}_i$ represents a D-dimensional feature vector corresponding to one part of the L image regions. The task is to output a caption $y = \{\boldsymbol{y}_1, \dots, \boldsymbol{y}_C\}$ encoded as a sequence of words. The input to the LSTM cell is time step-dependent context vectors $\hat{\boldsymbol{z}}_t$ representing the specific part of the input image

$$\hat{\boldsymbol{z}}_t = \sum_{i=1}^{L} \alpha_{t,i} \boldsymbol{a}_i, \tag{1}$$

which is defined as a weighted sum of the feature vectors. The weight $\alpha_{t,i}$ represents the current state of the machine attention, and is defined as a function of the original image feature $\boldsymbol{a}_i$ and the previous hidden state $\boldsymbol{h}_{t-1}$ of the LSTM:

$$e_{t,i} = f_{\text{att}}(\boldsymbol{a}_i, \boldsymbol{h}_{t-1}), \tag{2}$$

$$\alpha_{t,i} = \frac{\exp(e_{t,i})}{\sum_{i=1}^{L} \exp(e_{t,i})}. \tag{3}$$

Then, the caption sequence is estimated using the deep output layer as a function of the context vector, current hidden state, and the previous word [60]: $p(\boldsymbol{y}_t | \boldsymbol{a}, \boldsymbol{y}_1^{t-1}) = f_{\text{out}}(\boldsymbol{y}_{t-1}, \boldsymbol{h}_t, \hat{\boldsymbol{z}}_t)$. They also introduced a doubly stochastic regularization $\sum_t \alpha_{t,i} \approx 1$ to the final cost function, i.e., the model is regularized so that attention is paid uniformly over the whole image. Our interest in this work is to integrate human gaze feature into the attention function in Eq (2).

4.1 Integration of Gaze Information

As illustrated in Figure 4, our method takes both image and gaze features as input. The gaze feature is represented as a normalized fixation histogram $g = \{g_1, \dots, g_L\}$. The fixation histogram obtained from gaze recording is cropped and resized to the same L grid, and the g_i is taken from the same grid as $\boldsymbol{a}_i$. The gaze feature is given to the attention function together with the image feature a. In the original model, f_{att} is defined as a linear function

$$e_{t,i} = \boldsymbol{w}_{\text{att}} \boldsymbol{p}_{t,i} + c_{\text{att}}, \tag{4}$$

where $\boldsymbol{p}_{t,i}$ is a nonlinear projection of $\boldsymbol{a}_i$ and $\boldsymbol{h}_{t-1}$. According to the previous analysis, the gaze-assisted attention function requires the flexibility to allocate attention to non-fixated regions. Hence, we propose to split the machine attention according to human fixation as

$$e_{t,i} = g_i \boldsymbol{w}_{\text{pos}} \boldsymbol{p}_{t,i} + (1 - g_i) \boldsymbol{w}_{\text{neg}} \boldsymbol{p}_{t,i} + c_{\text{att}}. \tag{5}$$

This model can learn different weights for fixated ($\boldsymbol{w}_{\text{pos}}$) and non-fixated ($\boldsymbol{w}_{\text{neg}}$) regions, and can efficiently utilize the gaze feature without losing too much information from non-fixated image regions. We assess the advantage of this split attention model through experiments.
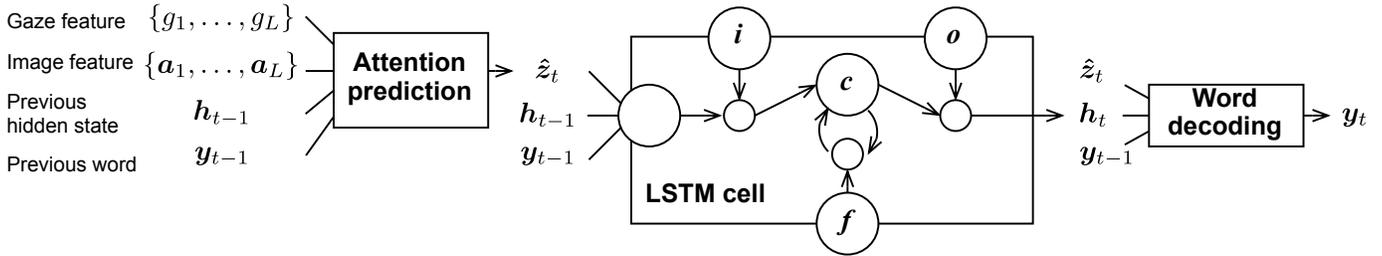
Fig. 4. Pipeline of the gaze-assisted image captioning. The attention function takes both image and gaze features as input, and the context vector weighted with the attention is given to the LSTM cell for word-by-word captioning.

4.2 Implementation Details

Following the original implementation, the image features were extracted from the last convolutional layer before max pooling using the 19-layer VGGNet [51] without fine tuning. This results in image features with $L = 14 \times 14 = 196$ and $D = 512$. Input images were resized so that the shortest side had a length of 256 pixels, and the center cropped 224×224 image was given to the network. Most of the other details were kept the same as in the original model implementation, while the dimensionality of the hidden state h was set to a lower value of 1,400 to account for the smaller amount of training data available. In the experiments, all models were trained using the Adam algorithm [61].

5 EXPERIMENTS

In this section, we report experimental results of our gaze-assisted image captioning models on the SALICON dataset. Since the test set in the SALICON dataset does not provide gaze data, we randomly split the 5,000 validation images into two 2,500 images for validation and test. The validation set is used for early stopping, and we show the performance of the best model on the test set. As suggested by Xu et al., we used the BLEU score [62] on the validation set for model selection. In order to give a fair comparison given the limited amount of training data, we fixed the random seed value for weight initialization and mini-batch selection between different models. Hence, all models were trained under exactly the same conditions.

5.1 Captioning Performance

We first show quantitative comparison of captioning performance. As evaluation metrics we used the implementations of commonly used metrics (BLEU 1-4, METEOR, ROUGE$_L$ and CIDEr scores) provided with the COCO dataset [63]. BLEU is a standard machine translation score measuring co-occurrences of n-grams [62], and ROUGE$_L$ is a text summarization metric based on the longest common subsequence [64]. METEOR score is based on the word alignment [65], and CIDEr measures consensus in captions [66]. Although the limitations of these computational metrics have often been pointed out and it is fundamentally difficult to evaluate goodness of natural language captions, they give us some insights into how the additional gaze information changes the captioning performance.

In Table 1, we report BLEU 1-4, METEOR, ROUGE$_L$ and CIDEr scores of different models. In addition to the original

machine attention model (*Machine*) [19] as well as our gaze-assisted split attention model (*Split attention*), we also show two additional baseline models. In the second row (*Gaze-only*), we show the case where the weight for non-fixated region in Eq. (5) is not used. This model represents a more straightforward design of the gaze-assisted attention model which heavily relies on fixated regions. In addition, we show results when the BMS saliency maps are used instead of fixation maps with the same architecture as the proposed model in the third row (*Saliency*).

As can be seen from the table, the proposed split attention model performs consistently better than the baseline models for all metrics even with the same underlying architecture. The difference between *Gaze-only* and *Split attention* results clearly illustrates the significance of the proposed split attention architecture. The performance gap between human gaze and bottom-up saliency also indicates the fundamental importance of the gaze information. Together with the previous analysis, our results indicate the importance of human gaze information for semantic image understanding.

5.2 Attention Allocation Examples

As discussed above, computational metrics do not fully explain how the captioning model is improved by the proposed attention model, and it is more important to discuss concrete examples of generated captions and their corresponding attention allocation results. Figure 5 visualizes some examples of the attention allocation results of the baseline and gaze-assisted models. Each row shows the input image, attention map examples at each step t, and their corresponding output words. Gaze information typically helps the model to find small but important objects in the scene, like the kites in the first example in Figure 5. Since gaze information explicitly provides object locations, it is also beneficial for avoiding repetition of object discovery in cluttered scenes. While the baseline model describes the laptop twice in the second example, the proposed model properly allocates the attention to the object region and generates the word only once. It is also noteworthy that, although there are some objects which do not attract human fixations, the proposed split attention model has the flexibility to allocate attention to such objects. In the third example, the *Gaze-only* model fails to describe the snowboard; the proposed split attention model successfully describes it.

Many prior works have reported the strong relationship between eye movements and viewing task [67], and how humans look at images heavily depends on the task the viewer is performing. From these examples it can be

TABLE 1. Image captioning performance. Columns show BLEU 1-4, METEOR, ROUGE$_L$ and CIDEr scores of different models. Columns show the performance of the original model [19], the model without the attention term for non-fixated regions, the model using saliency maps instead of fixation maps, and the proposed split attention model.

Model	BLEU				METEOR	ROUGE	CIDEr
	1	2	3	4			
Machine [19]	0.706	0.495	0.342	0.237	0.218	0.520	0.626
Gaze-only	0.704	0.492	0.340	0.236	0.215	0.519	0.613
Saliency	0.708	0.496	0.342	0.236	0.217	0.519	0.623
Split attention	**0.714**	**0.505**	**0.352**	**0.245**	**0.219**	**0.524**	**0.638**

Fig. 5. Sample images, the machine attention map at each step as well as the corresponding output words for the baseline and gaze-assisted models. The first example illustrates the case where the proposed model finds small but important objects (*kite*) in the scene. It also helps to suppress the repetition of object description in cluttered scenes (*laptop* in the second example). The proposed split attention model can also describe objects without strong fixation, such as *snowboard* in the third example. See the supplementary material for more examples.

seen that the behavior of the gaze-assisted model better resembles how humans see images, and they indicate the potential of utilizing different types of gaze behavior for task-oriented captioning. In addition, the performance gap between human gaze and visual saliency poses another important research question; whether computational saliency models can achieve similar performance, especially with recently proposed deep models [68], [69].

5.3 Word Prediction Performance

To better understand which words are correctly discovered by the proposed gaze-assisted model, we analyzed individual word prediction scores. In Figure 6a, we plot precision-recall points for words whose F-scores are improved by more than a threshold (0.05) by the fixation model. To improve

intelligibility, we show only words appearing more than 10 times in the ground-truth captions. We used the same weighted precision/recall measures as Chen et al. [63]; while negative images have a weight of 1, the number of captions containing the word is used as a weight for positive images. Words in both predicted and ground-truth captions are lemmatized using the NLTK library. Figure 6b shows the same precision-recall plot for words whose recall scores become worse with the proposed model.

As discussed above, the proposed model improves word discovery scores for small important objects such as *kite*, *knife*, *umbrella*, *(fire) hydrant*, *traffic (sign)*. The improvement of the words like *top*, *front*, and *about* could be related to the above-mentioned switching of the description subject. Although the proposed attention function can attend to background

(a)

(b)
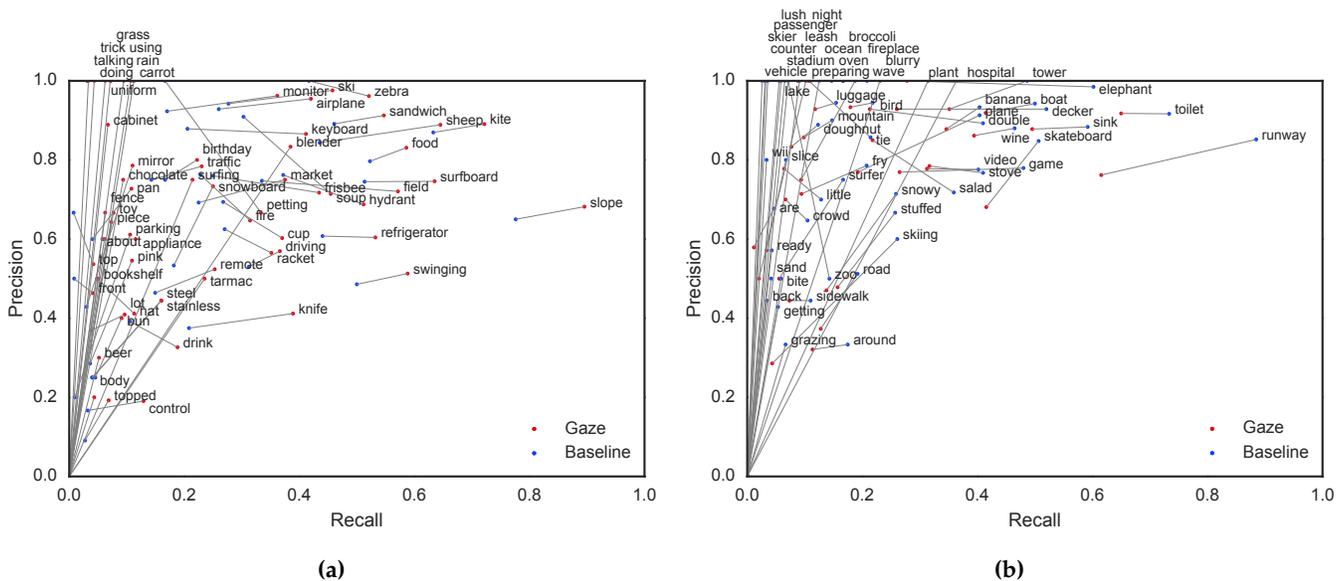
Fig. 6. Precision and recall of individual words for which the F-scores (a) improved and (b) degraded by using the proposed gaze-assisted image captioning model.

regions, it still loses some performance on words related to background scene categories, such as *tower*, *table*, *runway*, and *fireplace*, as can be seen in Figure 6b. It can also be seen that some words related to activity or context, like *game*, *blurry*, *night*, are also losing performance. However, as illustrated in Figure 5, the proposed split attention model also helps the model to find objects which do not attract human attention, such as *snowboard* and *ski*. Since the word discovery performance should also depend on the amount of training data, investigation of larger-scale eye tracking experiments is one of the most important future directions.

6 CONCLUSION

In this paper we presented a detailed study on how human gaze information can help holistic image understanding and captioning tasks. We first analyzed the relationship between gaze and state-of-the-art recognition models for both object and scene categories. We showed that human gaze is more correlated with important locations for object recognition models and can help to find more semantically important objects than bottom-up saliency models. We further presented the first gaze-assisted image captioning model and quantified its performance. With the previous findings in mind, we proposed a split attention model where the machine attention can be allocated selectively to fixated and non-fixated image regions. Our model improves the captioning performance of a baseline model on the challenging COCO/SALICON dataset, and achieved a similar performance improvement compared to state-of-the-art bottom-up saliency models.

These results underline the potential of gaze-assisted image captioning, particularly for cluttered images without a clearly depicted central object. The approach is similarly appealing for gaze-assisted captioning of unorganized image streams, for example those recorded using life-logging devices or other egocentric cameras. Since the SALICON dataset was collected using a pseudo-eye tracking setup, investigating gaze-assisted image understanding in an egocentric

setting with real gaze data is one of the most interesting directions for future work.

REFERENCES

[1] A. Bulling, J. A. Ward, H. Gellersen, and G. Tröster, "Eye movement analysis for activity recognition using electrooculography," *IEEE TPAMI*, vol. 33, no. 4, pp. 741–753, 2011.

[2] A. Fathi, Y. Li, and J. M. Rehg, "Learning to recognize daily actions using gaze," in *Proc. ECCV*, 2012, pp. 314–327.

[3] D. P. Papadopoulos, A. D. Clarke, F. Keller, and V. Ferrari, "Training object class detectors from eye tracking data," in *Proc. ECCV*, 2014, pp. 361–376.

[4] H. Sattar, S. Müller, M. Fritz, and A. Bulling, "Prediction of search targets from fixations in open-world settings," in *Proc. CVPR*, 2015, pp. 981–990.

[5] S. Karthikeyan, T. Ngo, M. Eckstein, and B. Manjunath, "Eye tracking assisted extraction of attentionally important objects from videos," in *Proc. CVPR*, 2015, pp. 3241–3250.

[6] M. Jiang, S. Huang, J. Duan, and Q. Zhao, "SALICON: Saliency in context," in *Proc. CVPR*, 2015, pp. 1072–1080.

[7] D. Rudoy, D. B. Goldman, E. Shechtman, and L. Zelnik-Manor, "Crowdsourcing gaze data collection," in *Proc. Collective Intelligence*, 2012.

[8] P. Xu, K. A. Ehinger, Y. Zhang, A. Finkelstein, S. R. Kulkarni, and J. Xiao, "Turkergaze: Crowdsourcing saliency with webcam based eye tracking," *arXiv preprint arXiv:1504.06755*, 2015.

[9] K. A. Funes Mora and J.-M. Odobez, "Geometric generative gaze estimation (G3E) for remote RGB-D cameras," in *Proc. CVPR*, 2014, pp. 1773–1780.

[10] Y. Sugano, Y. Matsushita, and Y. Sato, "Learning-by-synthesis for appearance-based 3d gaze estimation," in *Proc. CVPR*. IEEE, 2014, pp. 1821–1828.

[11] X. Zhang, Y. Sugano, M. Fritz, and A. Bulling, "Appearance-based gaze estimation in the wild," in *Proc. CVPR*, 2015, pp. 4511–4520.

[12] A. K. Mishra, Y. Aloimonos, L. F. Cheong, and A. Kassim, "Active visual segmentation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 34, no. 4, pp. 639–653, April 2012.

[13] S. Karthikeyan, V. Jagadeesh, R. Shenoy, M. Eckstein, and B. Manjunath, "From where and how to what we see," in *Proc. ICCV*, 2013, pp. 625–632.

[14] K. Yun, Y. Peng, D. Samaras, G. J. Zelinsky, and T. L. Berg, "Exploring the role of gaze behavior and object detection in scene understanding," *Frontiers in psychology*, vol. 4, 2013.

[15] G. J. Zelinsky, "Understanding scene understanding," *Frontiers in psychology*, vol. 4, 2013.

[16] H. Larochelle and G. E. Hinton, "Learning to combine foveal glimpses with a third-order boltzmann machine," in *Proc. NIPS*, 2010, pp. 1243–1251.

[17] V. Mnih, N. Heess, A. Graves *et al.*, "Recurrent models of visual attention," in *Proc. NIPS*, 2014, pp. 2204–2212.

[18] Y. Zheng, R. S. Zemel, Y.-J. Zhang, and H. Larochelle, "A neural autoregressive approach to attention-based recognition," *IJCV*, vol. 113, no. 1, pp. 67–79, 2014.

[19] K. Xu, J. Ba, R. Kiros, A. Courville, R. Salakhutdinov, R. Zemel, and Y. Bengio, "Show, attend and tell: Neural image caption generation with visual attention," in *Proc. ICML*, 2015.

[20] A. Borji and L. Itti, "State-of-the-art in visual attention modeling," *IEEE TPAMI*, vol. 35, no. 1, pp. 185–207, 2013.

[21] J. L. Long, N. Zhang, and T. Darrell, "Do convnets learn correspondence?" in *Proc. NIPS*, 2014, pp. 1601–1609.

[22] W. Einhäuser, M. Spain, and P. Perona, "Objects predict fixations better than early saliency," *Journal of Vision*, vol. 8, no. 14, p. 18, 2008.

[23] A. Nuthmann and J. M. Henderson, "Object-based attentional selection in scene viewing," *Journal of vision*, vol. 10, no. 8, p. 20, 2010.

[24] Y. Li, X. Hou, C. Koch, J. M. Rehg, and A. L. Yuille, "The secrets of salient object segmentation," in *Proc. CVPR*, 2014, pp. 280–287.

[25] J. Xu, M. Jiang, S. Wang, M. S. Kankanhalli, and Q. Zhao, "Predicting human gaze beyond pixels," *Journal of vision*, vol. 14, no. 1, p. 28, 2014.

[26] S. Ramanathan, H. Katti, N. Sebe, M. Kankanhalli, and T.-S. Chua, "An eye fixation database for saliency detection in images," in *Proc. ECCV*, 2010, pp. 30–43.

[27] Y. Sugano, Y. Matsushita, and Y. Sato, "Graph-based joint clustering of fixations and visual entities," *ACM Transactions on Applied Perception (TAP)*, vol. 10, no. 2, p. 10, 2013.

[28] T. Toyama, T. Kieninger, F. Shafait, and A. Dengel, "Gaze guided object recognition using a head-mounted eye tracker," in *Proc. ETRA*, 2012, pp. 91–98.

[29] D. Damen, T. Leelasawassuk, and W. Mayol-Cuevas, "You-do, i-learn: Egocentric unsupervised discovery of objects and their modes of interaction towards video-based guidance," *Computer Vision and Image Understanding*, vol. 149, pp. 98 – 112, 2016.

[30] K. Yun, Y. Peng, D. Samaras, G. J. Zelinsky, and T. Berg, "Studying relationships between human gaze, description, and computer vision," in *Proc. CVPR*. IEEE, 2013, pp. 739–746.

[31] I. Shcherbatyi, A. Bulling, and M. Fritz, "GazeDPM: Early integration of gaze information in deformable part models," *arXiv preprint arXiv:1505.05753*, 2015.

[32] S. Mathe and C. Sminchisescu, "Actions in the eye: dynamic gaze datasets and learnt saliency models for visual recognition," *IEEE TPAMI*, vol. 37, no. 7, pp. 1408–1424, 2015.

[33] S. Ramanathan, V. Yanulevskaya, and N. Sebe, "Can computers learn from humans to see better?: inferring scene semantics from viewers' eye movements," in *Proc. ACMMM*, 2011, pp. 33–42.

[34] M. I. Coco and F. Keller, "Integrating mechanisms of visual guidance in naturalistic language production," *Cognitive processing*, vol. 16, no. 2, pp. 131–150, 2015.

[35] ——, "Scan patterns predict sentence production in the cross-modal processing of visual scenes," *Cognitive Science*, vol. 36, no. 7, pp. 1204–1223, 2012.

[36] M. Ranzato, "On learning where to look," *arXiv preprint arXiv:1405.5488*, 2014.

[37] J. Ba, V. Mnih, and K. Kavukcuoglu, "Multiple object recognition with visual attention," in *Proc. ICLR*, 2015.

[38] K. Gregor, I. Danihelka, A. Graves, and D. Wierstra, "DRAW: A recurrent neural network for image generation," in *Proc. ICML*, 2015.

[39] D. Bahdanau, K. Cho, and Y. Bengio, "Neural machine translation by jointly learning to align and translate," in *Proc. ICLR*, 2015.

[40] K. Cho, A. Courville, and Y. Bengio, "Describing multimedia content using attention-based encoder-decoder networks," *arXiv preprint arXiv:1507.01053*, 2015.

[41] J. Devlin, H. Cheng, H. Fang, S. Gupta, L. Deng, X. He, G. Zweig, and M. Mitchell, "Language models for image captioning: The quirks and what works," in *Proc. ACL*, 2015.

[42] J. Donahue, L. A. Hendricks, S. Guadarrama, M. Rohrbach, S. Venugopalan, K. Saenko, and T. Darrell, "Long-term recurrent convolutional networks for visual recognition and description," in *Proc. CVPR*, 2015.

[43] H. Fang, S. Gupta, F. Iandola, R. Srivastava, L. Deng, P. Doll'ar, J. Gao, X. He, M. Mitchell, J. Platt *et al.*, "From captions to visual concepts and back," in *Proc. CVPR*, 2015.

[44] A. Karpathy and L. Fei-Fei, "Deep visual-semantic alignments for generating image descriptions," in *Proc. CVPR*, 2015.

[45] J. Mao, W. Xu, Y. Yang, J. Wang, Z. Huang, and A. Yuille, "Learning like a child: Fast novel visual concept learning from sentence descriptions of images," in *Proc. ICCV*, 2015.

[46] O. Vinyals, A. Toshev, S. Bengio, and D. Erhan, "Show and tell: A neural image caption generator," in *Proc. CVPR*, 2015.

[47] X. Chen and C. Lawrence Zitnick, "Mind's eye: A recurrent visual representation for image caption generation," in *Proc. CVPR*, 2015, pp. 2422–2431.

[48] Y. Ushiku, M. Yamaguchi, Y. Mukuta, and T. Harada, "Common subspace for model and similarity: Phrase learning for caption generation from images," in *Proc. ICCV*, 2015, pp. 2668–2676.

[49] Q. You, H. Jin, Z. Wang, C. Fang, and J. Luo, "Image captioning with semantic attention," in *Proc. CVPR*, 2016.

[50] J. Zhang and S. Sclaroff, "Exploiting surroundedness for saliency detection: a boolean map approach," *IEEE TPAMI*, 2015.

[51] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," in *Proc. ICLR*, 2015.

[52] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein *et al.*, "Imagenet large scale visual recognition challenge," *IJCV*, pp. 1–42, 2014.

[53] L. Wang, S. Guo, W. Huang, and Y. Qiao, "Places205-vggnet models for scene recognition," *arXiv preprint arXiv:1508.01667*, 2015.

[54] B. Zhou, A. Lapedriza, J. Xiao, A. Torralba, and A. Oliva, "Learning deep features for scene recognition using places database," in *Proc. NIPS*, 2014, pp. 487–495.

[55] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick, "Microsoft COCO: Common objects in context," in *Proc. ECCV*, 2014, pp. 740–755.

[56] Y.-W. Chao, Z. Wang, R. Mihalcea, and J. Deng, "Mining semantic affordances of visual object categories," in *Proc. CVPR*, 2015.

[57] S. Bird, E. Klein, and E. Loper, *Natural language processing with Python*. O'Reilly Media Inc., 2009.

[58] M. D. Zeiler and R. Fergus, "Visualizing and understanding convolutional networks," in *Proc. ECCV*, 2014, pp. 818–833.

[59] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.

[60] R. Pascanu, C. Gulcehre, K. Cho, and Y. Bengio, "How to construct deep recurrent neural networks," in *Proc. ICLR*, 2014.

[61] D. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.

[62] K. Papineni, S. Roukos, T. Ward, and W.-J. Zhu, "Bleu: a method for automatic evaluation of machine translation," in *Proc. ACL*, 2002, pp. 311–318.

[63] X. Chen, H. Fang, T.-Y. Lin, R. Vedantam, S. Gupta, P. Dollar, and C. L. Zitnick, "Microsoft coco captions: Data collection and evaluation server," *arXiv preprint arXiv:1504.00325*, 2015.

[64] C.-Y. Lin, "Rouge: A package for automatic evaluation of summaries," in *Text summarization branches out: Proceedings of the ACL-04 workshop*, vol. 8, 2004.

[65] M. Denkowski and A. Lavie, "Meteor universal: Language specific translation evaluation for any target language," in *Proceedings of the EACL 2014 Workshop on Statistical Machine Translation*, vol. 6, 2014.

[66] R. Vedantam, C. L. Zitnick, and D. Parikh, "Cider: Consensus-based image description evaluation," *arXiv preprint arXiv:1411.5726*, 2014.

[67] A. L. Yarbus, *Eye movements and vision*. Springer, 1967.

[68] N. Liu, J. Han, D. Zhang, S. Wen, and T. Liu, "Predicting eye fixations using convolutional neural networks," in *Proc. CVPR*, 2015, pp. 362–370.

[69] X. Huang, C. Shen, X. Boix, and Q. Zhao, "Salicon: Reducing the semantic gap in saliency prediction by adapting deep neural networks," in *Proc. ICCV*, 2015, pp. 262–270.