

---

# Improving Natural Language Processing Tasks with Human Gaze-Guided Neural Attention

---

Ekta Sood<sup>1</sup>, Simon Tannert<sup>2</sup>, Philipp Müller<sup>1</sup>, Andreas Bulling<sup>1</sup>

<sup>1</sup>University of Stuttgart, Institute for Visualization and Interactive Systems (VIS), Germany

<sup>2</sup>University of Stuttgart, Institute for Natural Language Processing (IMS), Germany

{ekta.sood, philipp.mueller, andreas.bulling}@vis.uni-stuttgart.de  
simon.tannert@ims.uni-stuttgart.de

## Abstract

A lack of corpora has so far limited advances in integrating human gaze data as a supervisory signal in neural attention mechanisms for natural language processing (NLP). We propose a novel *hybrid text saliency model* (TSM) that, for the first time, combines a cognitive model of reading with explicit human gaze supervision in a single machine learning framework. On four different corpora we demonstrate that our hybrid TSM duration predictions are highly correlated with human gaze ground truth. We further propose a novel *joint modeling approach* to integrate TSM predictions into the attention layer of a network designed for a specific upstream NLP task without the need for any task-specific human gaze data. We demonstrate that our joint model outperforms the state of the art in paraphrase generation on the Quora Question Pairs corpus by more than 10% in BLEU-4 and achieves state of the art performance for sentence compression on the challenging Google Sentence Compression corpus. As such, our work introduces a practical approach for bridging between data-driven and cognitive models and demonstrates a new way to integrate human gaze-guided neural attention into NLP tasks.

## 1 Introduction

Neural attention mechanisms have been widely applied in computer vision and have been shown to enable neural networks to only focus on those aspects of their input that are important for a given task [48, 81]. While neural networks are able to learn meaningful attention mechanisms using only supervision received for the target task, the addition of human gaze information has been shown to be beneficial in many cases [32, 58, 80, 84]. An especially interesting way of leveraging gaze information was demonstrated by works incorporating human gaze into neural attention mechanisms, for example for image and video captioning [71, 83] or visual question answering [58].

While attention is at least as important for reading text as it is for viewing images [13, 78], integration of human gaze into neural attention mechanisms for natural language processing (NLP) tasks remains under-explored. A major obstacle to studying such integration is data scarcity: Existing corpora of human gaze during reading consist of too few samples to provide effective supervision for modern data-intensive architectures and human gaze data is only available for a small number of NLP tasks. For paraphrase generation and sentence compression, which play an important role for tasks such as reading comprehension systems [23, 28, 54], no human gaze data is available.

We address this data scarcity in two novel ways: First, to overcome the low number of human gaze samples for reading, we propose a novel hybrid text saliency model (TSM) in which we combine a cognitive model of reading behavior with human gaze supervision in a single machine learning framework. More specifically, we use the E-Z Reader model of attention allocation during reading [60] to obtain a large number of synthetic training examples. We use these examples to

pre-train a BiLSTM [22] network with a Transformer [75] whose weights we subsequently refine by training on only a small amount of human gaze data. We demonstrate that our model yields predictions that are well-correlated with human gaze on out-of-domain data. Second, we propose a novel joint modeling approach of attention and comprehension that allows human gaze predictions to be flexibly adapted to different NLP tasks by integrating TSM predictions into an attention layer. By jointly training the TSM with a task-specific network, the saliency predictions are adapted to this upstream task without the need for explicit supervision using real gaze data. Using this approach, we outperform the state of the art in paraphrase generation on the Quora Question Pairs corpus by more than 10% in BLEU-4 and achieve state of the art performance on the Google Sentence Compression corpus. As such, our work demonstrates the significant potential of combining cognitive and data-driven models and establishes a general principle for flexible gaze integration into NLP that has the potential to also benefit tasks beyond paraphrase generation and sentence compression.

## 2 Related work

Our work is related to previous works on 1) NLP tasks for text comprehension, 2) human attention modeling, as well as 3) gaze integration in neural network architectures.

### 2.1 NLP tasks for text comprehension

Two key tasks in machine text comprehension are paraphrasing and summarization [8, 28, 9, 41, 24]. While paraphrasing is the task of “conveying the same meaning, but with different expressions” [9, 18, 41], summarization deals with extracting or abstracting the key points of a larger input sequence [21, 73, 33]. Though advances have helped bring machine comprehension closer to human performance, humans are still superior for most tasks [3, 79, 85]. While attention mechanisms can improve performance by helping models to focus on relevant parts of the input [57, 65, 63, 7, 27, 11], the benefit of explicit supervision through human attention remains under-explored.

### 2.2 Human attention modeling

Predicting what people visually attend to in images (saliency prediction) is a long-standing challenge in neuroscience and computer vision [4, 6, 40]. In contrast to images, most attention models for eye movement behaviors during reading are cognitive process models, i.e. models that do not involve machine learning but implement cognitive theories [17, 59, 60]. Key challenges for such models are a limited number of parameters, hand-crafted rules and thus a difficulty to adapt them to different tasks and domains, as well as the difficulty to use them as part of an end-to-end trained machine learning architectures [16, 39, 45]. One of the most influential cognitive models of gaze during reading is the E-Z Reader model [60] It assumes attention shifts to be strictly serial in nature and that saccade production depends on different stages of lexical processing. that has been successful in explaining different effects seen in attention allocation during reading [61, 62].

In contrast, learning-based attention models for text remain under-explored. Nilsson and Nivre [50] trained person-specific models on features including length and frequency of words to predict fixations and later extended their approach to also predict fixation durations [51]. The first work to present a person-independent model for fixation prediction on text used a linear CRF model [46]. A separate line of work has instead tried to incorporate assumptions about the human reading process into the model design. For example, the Neural Attention Trade-off (NEAT) language model was trained with hard attention and assigned a cost to each fixation Hahn and Keller [25]. Subsequent work applied the NEAT model to question answering tasks, showing task-specific effects on learned attention patterns that reflect human behavior [26]. Further approaches include sentence representation learning using surprisal and part of speech tags as proxies to human attention [76], attention as a way to improve time complexity for NLP tasks [68], and learning saliency scores by training for sentence comparison [66]. Our work is fundamentally different from all of these works in that we, for the first time, combine cognitive theory and data-driven approaches.

### 2.3 Gaze integration in neural network architectures

Integration of human gaze data into neural network architectures has been explored for a range of computer vision tasks [32, 69, 80, 83, 84]. Sugano and Bulling [71] were the first to use gaze as an

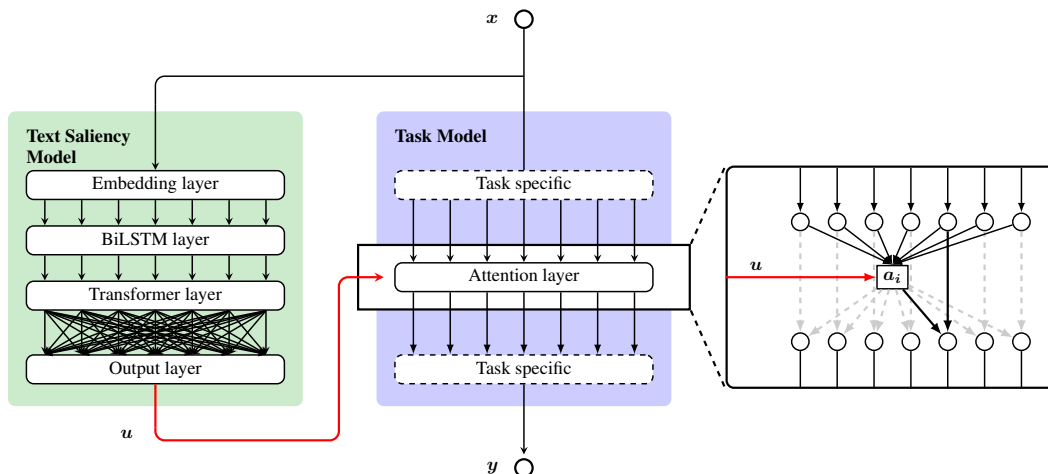


Figure 1: High-level architecture of our model. Given an input sentence  $x_1 \dots x_n$  the Text Saliency Model produces attention scores  $u_1 \dots u_n$  for each word in the input sentence  $x$ . The Task Model combines this information with the original input sentence to produce an output sentence  $y_1 \dots y_m$ .

additional input to the attention layer for image captioning, while Qiao et al. [58] used human-like attention maps as an additional supervision for the attention layer for a visual question answering task. Most previous work in gaze-supported NLP has used gaze as an input feature, e.g. for syntactic sequence labeling [36], classifying referential versus non-referential use of pronouns [82], reference resolution [30], key phrase extraction [86], or prediction of multi-word expressions [64]. Recently, Hollenstein et al. [29] proposed to build a lexicon of gaze features given word types, overcoming the need for gaze data at test time. Two recent works proposed methods inspired by multi-task learning to integrate gaze into NLP classification tasks. [37] did not integrate gaze into the attention layers but demonstrated performance improvements by adding a gaze prediction task to regularize a sentence compression model. [2] did not predict human gaze for the target task but used ground-truth gaze from another eye tracking corpus to regularize their neural attention function. In stark contrast, our work is the first to combine a cognitive model of reading and a data-driven approach to predict human gaze, to directly integrate these predictions into the neural attention layers, and to jointly train for two different tasks – generative (paraphrase generation) and classification (sentence compression).

### 3 Method

We make two contributions: A hybrid text saliency model as well as two attention-based models for paraphrase generation and text summarization employed in a novel joint modeling approach<sup>1</sup>.

#### 3.1 Hybrid text saliency model

To overcome the limited amount of eye-tracking data for reading comprehension tasks, we propose a *hybrid approach* when training our text saliency model. In the first stage of training, we leverage the E-Z Reader model [60] to generate a large amount of training data over the CNN and Daily Mail Reading Comprehension Corpus [28]. After training the text saliency model until convergence using this synthetic data, in a second training phase we fine-tune the network with real eye tracking data of humans reading from the Provo and Geco corpus [43, 14]. We used the most recent implementation of EZ Reader (Version 10.2) available from the authors' website<sup>2</sup>.

The task of text saliency is to predict fixation durations  $u_i$  for each word  $x_i$  of an input sentence. In our text saliency model, we combine a BiLSTM network [22] with a Transformer [75] (see Figure 1 for an overview). Each word  $x_i$  of the input sentence is encoded using pre-trained GloVe

<sup>1</sup>Code and other supporting material can be found at [https://perceptualui.org/publications/sood20\\_neurips/](https://perceptualui.org/publications/sood20_neurips/)

<sup>2</sup><http://www.erikdreichle.com/downloads.html>

embeddings [55]. The resulting embeddings are fed into a single-layer BiLSTM network [22] that integrates information over the whole input sentence. The outputs from the BiLSTM network are fed into a Transformer network with multi-headed self-attention [75]. In contrast to [75], we only use the encoder of the Transformer network. Furthermore, we do not provide positional encodings as input because this information is already implicitly present in the outputs of the BiLSTM layer. In preliminary experiments we found advantages in using only four layers with four attention heads each for the Transformer network in contrast to the six layers with 12 heads in the original architecture [75]. We also found the combination of a BiLSTM network with a subsequent Transformer network to yield predictions that are most similar to human data, particularly with longer sequence lengths. According to [77], this might be due to the transformer encoding coarse relational information about positions of sequence elements, while the BiLSTM better captures fine-grained word level context. The specific choice of architecture therefore allows our model to better capture the sequential context while still maintaining computational efficiency. Finally, a fully connected layer is used to obtain an attention score  $u_i$  for each input word  $x_i$  in  $x$ . We apply sigmoid nonlinearities with subsequent normalization over the input sentence to obtain a probability distribution over the sentence. As loss function we use the mean squared error.

### 3.2 Joint modeling for natural language processing tasks

To model the relationship between attention allocation and text comprehension, we integrate the TSM with two different NLP task attention-based networks in a *joint model* (see Figure 1). Specifically, we propose a modification to the Luong attention layer [44] that is a computationally light-weight but highly effective, multiplicative attention algorithm [44, 5]. We compute attention scores  $a_i$  as

$$a_i = \text{softmax}(\text{score}_T(h_i, s_j)) \quad (1)$$

using our task-specific modified score functions  $\text{score}_T$ . For the tasks of paraphrase generation and sentence compression, respectively, we propose the novel score functions

$$\text{score}_{\text{ParaGen}}(h_i, s_j) = u \odot h_i^\top W_a s_j \quad (2)$$

$$\text{score}_{\text{TextComp}}(h_i, s_j) = u \odot v_a^\top \tanh(W_a[h_i; s_j]) \quad (3)$$

Where  $h_i$  is the current hidden state,  $s_j$  are the hidden states of the encoder and  $W_a$  and  $v_a$  are learnable parameters of the attention mechanism. The outputs of the TSM model  $u$  on the input sentence are incorporated into the score function by element-wise multiplication. This way, attention scores in the upstream task network reflect word saliencies learnt from humans. In addition to that, the error signal from the upstream loss function can be propagated back to the TSM in order to adapt its' parameters to the upstream task, thereby defining an implicit loss on  $u$ . This way, the attention distribution  $u$  returned by the TSM is adapted to the specific upstream task, allowing us to incorporate and adapt a neural model of attention to tasks for which no human gaze data is available. Note, as we have two different tasks namely generative (paraphrase generation) and classification (sentence compression), we used different score functions as suggested by previous work [44].

## 4 Experiments

### 4.1 Joint model with upstream tasks

#### Evaluation details

**Datasets** We used two standard benchmark corpora to evaluate each upstream NLP task. For paraphrase generation, we used the Quora Question Pairs corpus<sup>3</sup> that consists of human-annotated pairs of paraphrased questions that were crawled from Quora. We followed the common practice of excluding negative paraphrase examples from the corpus to obtain training data for paraphrase generation [54, 23]. We split the data according to [23, 54], using either 100K or 50K examples for training, 45K examples for validation, and 4K examples for testing. For the sentence compression

<sup>3</sup><https://www.quora.com/q/quoradata/First-Quora-Dataset-Release-Question-Pairs>

Table 1: Ablation study results and comparison with the state of the art for paraphrase generation with both data splits in terms of BLEU-4 score for different training set sizes and sentence compression in terms of F1 score and compression ratio. Also shown is the number of model parameters.

Paraphrase Generation (BLEU-4)				Sentence Compression			
Method	50K	100K	Params	Method	F1	CR	Params
Baseline (Seq-to-Seq)	7.11	8.91	45M	Klerke et al. (2016)	80.9	—	—
Patro et al. (2018)	16.5	17.9	—	Baseline (BiLSTM)	81.3	0.39	12M
				Zhao et al. (2018)	<b>85.1</b>	0.39	—
No Fixation	24.62	27.81	69M	No Fixation	83.4	0.38	129M
Random TSM Init	25.26	27.11	79M	Random TSM Init	83.7	0.38	178M
TSM Weight Swap	23.43	27.60	79M	TSM Weight Swap	83.8	0.38	178M
Frozen TSM	25.73	28.26	79M	Frozen TSM	83.9	0.37	178M
Ours	<b>26.24</b>	<b>28.82</b>	79M	Ours	<b>85.0</b>	0.39	178M

task we used the Google Sentence Compression corpus [20] containing 200K sentence compression pairs that were crawled from news articles. We split the data according to [87], taking the first 1K examples as test data, and the next 1K as validation data.

**Paraphrase generation** Our first text comprehension task was paraphrase generation where, given a source sentence, the model has to produce a different target sentence with the same meaning that may have a different length. We used a sequence-to-sequence network with word-level attention that was originally proposed for neural machine translation [1]. The model consisted of two recurrent neural networks, an encoder and an attention decoder (see Figure 1). The encoder consisted of an embedding layer followed by a gated recurrent unit (GRU) [10]. The decoder produced an output sentence step-by-step given the hidden state of the encoder and the input sentence. At each output step, the encoded input word and the previous hidden state are used to produce attention weights using our modified Luong attention (see Equation 2). These attention weights are combined with the embedded input sentence and fed into a GRU to produce an output sentence. The loss between predicted and the ground-truth paraphrase was calculated over the entire vocabulary using cross-entropy.

**Sentence compression** As a second task, we opted for deletion-based sentence compression that aims to delete unimportant words from an input sentence [31, 38, 47, 12, 20]. We incorporated the attention mechanism into the baseline architecture presented in [20]. The network consisted of three stacked LSTM layers with dropout after each LSTM layer as a regularization method. The outputs of the last LSTM layer were fed through our modified Luong attention mechanism (see Equation 3) and two fully connected layers which predicted for each word whether it should be deleted. The loss between predicted and ground truth deletion mask was calculated with cross-entropy.

**Training** We used pre-trained 300-dimensional GloVe embeddings in both the TSM and the upstream task network to represent the input words [55]. We trained both upstream task models using the ADAM optimizer [35] with a learning rate of 0.0001. For paraphrase generation we used uni-directional GRUs with hidden layer size 1,024 and dropout probability of 0.2. For sentence compression we used BiLSTMs with hidden layer size 1,024 and dropout probability of 0.1.

**Metrics** The most common metric to evaluate text generative tasks is BLEU [53], which measures the n-gram overlap between the produced and target sequence. To ensure reproducibility, we followed the standard Sacrebleu [56] implementation that uses BLEU-4. For sentence compression, we followed previous works [20, 87] by reporting the F1 score as well as the compression ratio calculated as the length of the compressed sentence divided by the input sentence length measured in characters [20].

## Results and discussion

Results for our joint model on paraphrase generation and sentence compression in comparison to the state of the art are shown in Table 1. As can be seen in the table, for paraphrase generation our approach achieves a BLEU-4 score of 28.82 when using 100K training examples, clearly

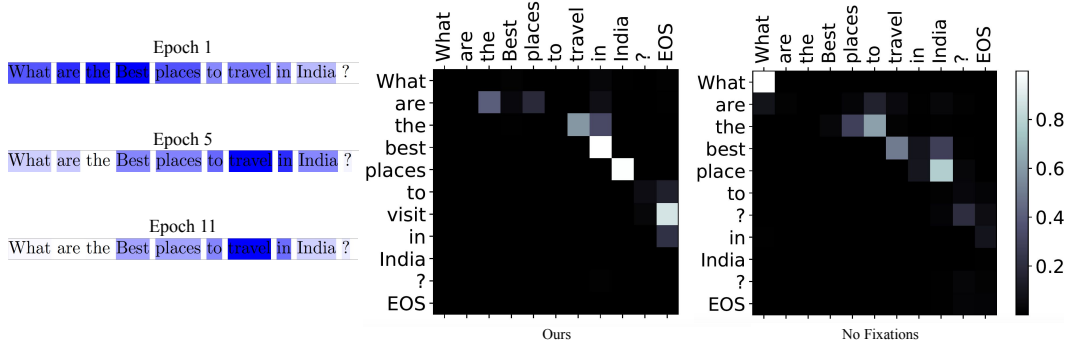


Figure 2: Paraphrase generation attention maps for both sub-networks (TSM predictions and upstream task attention) in our joint architecture. We show the TSM fixation predictions (left in blue) over epochs (last epoch is our converged models). We show the two-dimensional neural attention maps (right), showing our model and the *no fixation* model from our ablation study. The two-dimensional maps show the input sequence (horizontal axis) and the predicted sequence (vertical axis). We show the temporal TSM predictions over epochs, in order to depict how the fixation predictions change while training. The fixation predictions (for each epoch, left) are computed over words in the input sequences and then are integrated into the neural attention mechanism which in turn is used to make a prediction (vertical axis, right).

outperforming the previous state of the art for this task from [54] (17.9 BLEU-4). The same holds for 50K training examples (26.24 vs. 16.5 BLEU-4). For sentence compression, our joint model achieves a F1 score of 85.0 and a compression rate of 0.39. This is on par with the state of the art performance of 85.1 F1 score and 0.39 compression rate reported in Zhao et al. [87]<sup>4</sup>. In that work, a syntax-based language model was used to learn the syntactic dependencies between lexical items in the given input sequence. In contrast, our current method does not require any syntax-based language model, but it will be interesting to see whether it will benefit from additional syntactic information in future work. When comparing our results for sentence compression on the Google dataset to [37] we observe an increase of  $\sim 5\%$  F1 score for our method (cf. Table 1).

To further analyze the impact of our joint modeling approach, we evaluated several ablated versions of our model:

- **Baseline (Seq-to-Seq):** Stand-alone models based on a Seq-2-seq network [1] for paraphrase generation and a BiLSTM network [67] for sentence compression.
- **No Fixations:** Stand-alone upstream task network with original Luong attention (no TSM).
- **Random TSM Init:** Random initialization of the TSM instead of training on E-Z Reader and human data. Still implicit supervision by the upstream task during joint training.
- **TSM Weight Swap:** Exchange of the weights of the TSM model between tasks, i.e. sentence compression using the TSM weights obtained from the best-performing paraphrase generation model and vice versa.
- **Frozen TSM:** Training of the TSM with E-Z Reader and human gaze predictions but with frozen weights in the joint training with the upstream task, i.e. no adaptation of the TSM.

As can be seen from Table 1, all ablated models obtain inferior performance to our full model on both tasks (statistically significant at the 0.05 level). Notably, even the *No Fixation* model improves drastically over the *Seq-to-Seq baseline* for paraphrase generation, most likely due to the significant increase in network parameters. The benefit of training the *TSM* with our hybrid approach before using it in the joint model is underlined by the performance difference between the *Random TSM Init* (e.g. decrease in performance for both tasks) and our full model (e.g. best performance and differently adapted saliency predictions (see Table 1 and Figure 2)).<sup>5</sup>

<sup>4</sup>For a more detailed comparison to our model see table 1 in the supplementary material.

<sup>5</sup>Additional 1D and 2D maps over all conditions are available in the supplementary material.

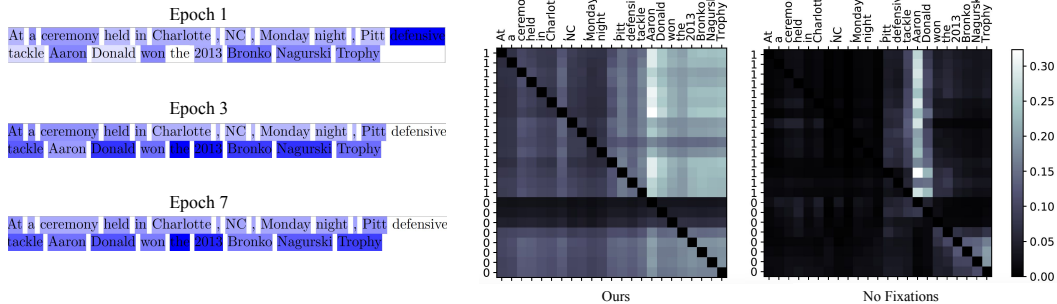


Figure 3: Sentence compression attention maps for both sub-networks (TSM predictions and upstream task attention) in our joint architecture. We show the TSM fixation predictions (left in blue) over epochs (last epoch is our converged models). We show the two-dimensional neural attention maps (right), showing our model and the *no fixation* model from our ablation study. The two-dimensional maps show the input sequence (horizontal axis) and the predicted sequence (vertical axis). We show the temporal TSM predictions over epochs, in order to depict how the fixation predictions change while training. The fixation predictions (for each epoch, left) are computed over words in the input sequences and then are integrated into the neural attention mechanism which in turn is used to make a prediction (vertical axis, right).

Most importantly, our full model achieves higher performance than the *Frozen TSM* model in all evaluations (e.g. 85.0 vs. 83.9 F1 for sentence compression), indicating that our model successfully adapts the TSM predictions during joint training. This is further underlined by the inferior performance of the *TSM Weight Swap* model: Swapping the optimal TSM weights between different upstream tasks leads to a notable performance decrease (e.g. 85.0 vs. 83.7 F1 for sentence compression), implying that the TSM model adaptation is specific to the upstream task.

To gain insights into how our joint model training adapts TSM predictions to specific upstream tasks, we analyzed the saliency predictions over time. Figure 2 shows a visualization of representative samples for both tasks over multiple training epochs. As can be seen in the left half of the figure, the adapted saliency predictions differ significantly from each other. In paraphrase generation (top) the saliency predictions focus on fewer words in the sentence within 11 epochs, specifically the word “travel” that is replaced in the correct paraphrase by “visit”. For sentence compression (bottom) the predictions continue to be spread over the whole sentence with only slight changes in the distribution over the words. This makes sense given that the task of this network is to delete as many words in the input sequence as possible while still maintaining syntactic structure and meaning.

The right half of the figure shows 2D neural attention maps of the converged models with the input sequence on the horizontal and the prediction on the vertical axis for *our* (with fixations) and the *No Fixation* model, respectively. As can be seen, our model correctly predicts the paraphrase, while the *No Fixation* model does not. Also, both the converged models neural attention weights differ with respect to allocation of probability mass. We see the *No Fixation* model densely concentrates attention towards a specific few input words (horizontal axis) when predicting several words (vertical axis). In contrast, the attention mass of our model is more spread out.

## 4.2 Pre-training of the hybrid text saliency model (TSM)

### Evaluation details

**Training datasets** Training the TSM consists of two stages: pre-training with synthetic data generated by E-Z Reader, and subsequent fine-tuning on human gaze data. For the first step, we run E-Z Reader on the CNN and Daily Mail corpus [28] consisting of 300k online news articles with on average 3.75 sentences. As recommended in Reichle et al. [60], we run E-Z Reader 10 times for each sentence to ensure stability in fixation predictions. For training we obtain a total of 7.6M annotated sentences on Daily Mail and 3.1M for CNN. For validation, we obtained 850K sentences on Daily Mail and 350K on CNN. For the second step, we used the two established gaze corpora Provo [43] and Geco [14]. Provo contains 55 short passages, extracted from different sources such as popular




Figure 4: Heatmaps showing human fixation durations in red and hybrid TSM duration predictions in blue. Here we show three different example sentences in order to depict the similarity between TSM word-level durations predictions as compared to human ground truth word-level durations

Table 2: Comparison of predicted and human ground-truth fixation durations for the different TSM conditions and corpora in terms of mean squared error (MSE), Jensen Shannon Divergence (JSD), and Spearman’s rank correlation ( $\rho$ ) between the part of speech tags based fixation distributions for model predictions and ground truth. A star indicates statistically significant  $\rho$  at  $p < 0.05$ .

Corpus	TSM			TSM w/o pre-training			TSM w/o fine-tuning		
	MSE	JSD	$\rho$	MSE	JSD	$\rho$	MSE	JSD	$\rho$
Dundee	0.063	0.39	0.99*	0.071	0.39	0.99*	0.096	0.47	-0.68
Provo + Geco	0.105	0.34	1.00*	0.112	0.36	0.99*	0.238	0.46	0.10
Provo	0.003	0.24	0.88*	0.008	0.44	0.83*	0.032	0.52	-0.25
Geco	0.118	0.35	0.99*	0.127	0.35	0.98*	0.267	0.45	-0.10
MQA-RC	0.064	0.36	0.94*	0.071	0.36	0.76*	0.083	0.42	-0.05

science magazines and fiction stories [43]. We split the data into 10K sentence pairs (pairs means sentence to human, as multiple humans read the same sentence) for train and 1K sentence pairs for validation. Geco is comprised of long passages from a popular novel [14]. We split the data into 65K sentence pairs for train and 8K sentence pairs for validation.

**Test datasets** We evaluated our model on the validation sets of the Provo and Geco corpora, as well as on the Dundee [34] and MQA-RC corpora [70]. The combined validation corpora of Provo and Geco contained 18K sentence pairs. Dundee consists of recordings from 10 participants reading 20 news articles while MQA-RC corpus is a 3-condition reading comprehension corpus using 32 documents from the MovieQA question answering dataset [72]. For our evaluation we used 1K sentence pairs from the free reading condition. This dataset is substantially different from the other eye tracking corpora because its stimuli are scraped from online sources and contain noise not found in text intended for human reading.

**Implementation details** We used pre-trained 300 dimensional GloVe word embeddings [55]. Our network has a bidirectional LSTM, with four transformer self-attention layers with four heads and hidden size of 128. The model objective is to predict normalized fixation durations for each word in the input sentence, resulting in saliency scores between 0 and 1. We used the ADAM optimizer [35] with a learning rate of 0.00001, batch size of 100, and dropout of 0.5 after the embedding layer and the recurrent layer. We pre-trained our network on synthetic training data for four epochs, and then fine-tune it on human data for 10 epochs.

**Metrics** To evaluate the TSM model, we compute mean squared error (MSE) between the predicted and ground truth fixation durations as well as the Jensen-Shannon Divergence (JSD) [42]. JSD is widely used in eye tracking research to evaluate inter-gaze agreement [49, 19, 15, 52] as, unlike Kullback-Leibler Divergence, JSD is symmetric. In addition we measured the word type predictability as it is a well-known predictor of fixation probabilities [25, 50]. We used the Stanford tagger [74] to predict part-of-speech (POS) tags for our corpora and compute the average fixation probability per tag, allowing us to compute the correlation between our model and ground truth using Spearman’s  $\rho$ .



## Results and discussion

Table 2 shows the performance of our model and ablation conditions in terms of means squared error (MSE), Jensen-Shannon-Divergence (JSD) and correlation to human ground truth. As ablation conditions we evaluate a model only trained on human data (w/o pre-train) as well as a model that is not fine-tuned on human data (w/o fine-tune), but only trained with E-Z Reader data.

Most importantly, our model is superior to- or on par with both ablation variants across all metrics and corpora, showing the importance of both the E-Z Reader pre-training as well as the fine-tuning with human data. Pre-training with data obtained from E-Z Reader is most beneficial in the case of the small Provo corpus, where we observe a reduction from 0.44 JSD to 0.24 JSD by adding the pre-training step. For the larger corpora this difference is less pronounced but still present. It is interesting to note that TSM w/o fine-tune performs consistently the worst, indicating that training on E-Z Reader data alone insufficient even though it provides benefits when combined with human data.

Using the correlations to human gaze over the POS distributions, we can compare our approach to Hahn and Keller [25] who achieved a  $\rho$  of 0.85 on the Dundee corpus, compared to a  $\rho$  of 0.99 achieved by our model. Furthermore we observe an especially large improvement in  $\rho$  as a result of E-Z Reader pre-training on the MQA-RC dataset. This dataset, unlike the other eye tracking corpora, is generated from stimuli which were scraped from online sources regarding movie plots, underlining the effectiveness of our approach in generalizing to out-of-domain data. In further analyses on the POS based correlations we observed that content words, such as adjectives, adverbs, nouns, and verbs, are more predictive than function words.<sup>6</sup> Lastly, we provide a qualitative impression of our method by comparing attention maps using our TSM predictions to ground truth human data (see Figure 4).

## 5 Conclusion

In this work we made two novel contributions towards improving natural language processing tasks using human gaze predictions as a supervisory signal. First, we introduced a novel hybrid text saliency model that, for the first time, integrates a cognitive reading model with a data-driven approach to address the scarcity of human gaze data on text. Second, we proposed a novel joint modeling approach that allows the TSM to be flexibly adapted to different NLP tasks without the need for task-specific ground truth human gaze data. We showed that both advances result in significant performance improvements over the state of the art in paraphrase generation as well as competitive performance for sentence compression but with a much less complex model than the state of the art. We further demonstrated that this approach is effective in yielding task-specific attention predictions. Taken together, our findings not only demonstrate the feasibility and significant potential of combining cognitive and data-driven models for NLP tasks – and potentially beyond – but also how saliency predictions can be effectively integrated into the attention layer of task-specific neural network architectures to improve performance.

---

<sup>6</sup>Detailed POS distributions are available in the supplementary material.

## Broader Impact

We identified a number of potential benefits and risks of our approach.

**Potential benefits** Our approach could benefit interactive applications that quantify, support, or enhance reading behavior without the need for special-purpose eye tracking equipment. Specifically, we see potential for e-learning applications in which our approach could be used to qualify reader actions and provide feedback to encourage improvement in reading comprehension. In addition, we see potential for our approach to be used as a key component in diagnostic tools to identify atypical eye movement behaviors correlated to cognitive impairments such as learning disabilities (e.g. such as Autism Spectrum Disorder or ADHD) or early onset detection of neuro-degenerative diseases which impact eye movement control (e.g Parkinsons disease). In addition, our hybrid approach could be useful for researchers building computational models of cognition, specifically geared towards combining traditional cognitive process models with neural networks in order to build a model which better emulates human cognitive processes – potentially allowing for increase in parameters and task complexity for further more robust models of human behavior. Lastly, our joint learning approach might prove useful to machine learning researchers who aim to implement artificial systems that more similarly model human behavior and thus perform more similarly to humans on currently challenging machine comprehension tasks. Our method could be used, for example, to automatically generate paraphrases for more natural playing experience in computer games (e.g., more varied utterances of non-player characters); or to aid in making a news text faster to read by compressing it before presenting it to the user.

**Potential risks** Despite these potential future benefits, we also identified a few risks. Given the nature of the underlying E-Z reader model, it is likely that reading behavior found in atypical individuals currently cannot be predicted. This could consequently negatively impact the quality of the provided feedback, guidance or diagnostic analyses. We also see a potential risk in that our improved paraphrase generation model could be used to disguise plagiarized texts. Finally, improvements in sentence compression could make users grow accustomed to simplified and potentially less balanced information if they increasingly read compressed news articles.

## Acknowledgments and Disclosure of Funding

E. Sood was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy - EXC 2075 – 390740016; S. Tannert was supported by IBM Research AI through the IBM AI Horizons Network; P. Müller and A. Bulling were funded by the European Research Council (ERC; grant agreement 801708). Additional revenues related to, but not supporting, this work: Scholarship by Google for E. Sood. We would like to thank the following people for their helpful insights and contributions: Sean Papay, Pavel Denisov, Prajit Dhar, Manuel Mager, Diego Frassinelli, Fabian Koegel and Keerthana Jaganathan.

## References

- [1] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate. In *Proc. International Conference on Learning Representations*, 2015.
- [2] Maria Barrett, Joachim Bingel, Nora Hollenstein, Marek Rei, and Anders Søgaard. Sequence classification with human attention. In *Proc. Conference on Computational Natural Language Learning*, pages 302–312, 2018.
- [3] Matthias Blohm, Glorianna Jagfeld, Ekta Sood, Xiang Yu, and Ngoc Thang Vu. Comparing attention-based convolutional and recurrent neural networks: Success and limitations in machine reading comprehension. In *Proc. Conference on Computational Natural Language Learning*, pages 108–118, 2018. doi: 10.18653/v1/K18-1011.
- [4] Ali Borji and Laurent Itti. State-of-the-art in visual attention modeling. *IEEE transactions on pattern analysis and machine intelligence*, 35(1):185–207, 2012.

- [5] Denny Britz, Anna Goldie, Minh-Thang Luong, and Quoc Le. Massive exploration of neural machine translation architectures. *arXiv preprint arXiv:1703.03906*, 2017.
- [6] Zoya Bylinskii, Adrià Recasens, Ali Borji, Aude Oliva, Antonio Torralba, and Frédo Durand. Where should saliency models look next? In *Proc. European Conference on Computer Vision*, pages 809–824. Springer, 2016.
- [7] Ziqiang Cao, Wenjie Li, Sujian Li, Furu Wei, and Yanran Li. AttSum: Joint learning of focusing and summarization with neural attention. In *Proc. International Conference on Computational Linguistics: Technical Papers*, pages 547–556, December 2016.
- [8] Danqi Chen, Jason Bolton, and Christopher D. Manning. A thorough examination of the CNN/daily mail reading comprehension task. In *Proc. Annual Meeting of the Association for Computational Linguistics*, pages 2358–2367, August 2016. doi: 10.18653/v1/P16-1223.
- [9] Eunah Cho, He Xie, and William M Campbell. Paraphrase generation for semi-supervised learning in nlu. In *Proc. Workshop on Methods for Optimizing and Evaluating Neural Language Generation*, pages 45–54, 2019.
- [10] Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using RNN encoder–decoder for statistical machine translation. In *Proc. Conference on Empirical Methods in Natural Language Processing*, pages 1724–1734, 2014. doi: 10.3115/v1/D14-1179.
- [11] Kyunghyun Cho, Aaron Courville, and Yoshua Bengio. Describing multimedia content using attention-based encoder-decoder networks. *IEEE Transactions on Multimedia*, 17(11):1875–1886, 2015.
- [12] James Clarke and Mirella Lapata. Global inference for sentence compression: An integer linear programming approach. *Journal of Artificial Intelligence Research*, 31:399–429, 2008.
- [13] Elena Commodari and Maria Guarnera. Attention and reading skills. *Perceptual and Motor Skills*, 100(2):375–386, 2005.
- [14] Uschi Cop, Nicolas Dirix, Denis Drieghe, and Wouter Duyck. Presenting geco: An eyetracking corpus of monolingual and bilingual sentence reading. *Behavior Research Methods*, 49(2): 602–615, 2017.
- [15] Alan Davies, Gavin Brown, Markel Vigo, Simon Harper, Laura Horseman, Bruno Splendiani, Elspeth Hill, and Caroline Jay. Exploring the relationship between eye movements and electrocardiogram interpretation accuracy. *Scientific reports*, 6:38227, 2016.
- [16] Włodzisław Duch, Richard J. Oentaryo, and Michel Pasquier. Cognitive architectures: Where do we go from here? In *Proc. Conference on Artificial General Intelligence*, page 122–136, NLD, 2008. ISBN 9781586038335.
- [17] Ralf Engbert, Antje Nuthmann, Eike M Richter, and Reinhold Kliegl. Swift: A dynamical model of saccade generation during reading. *Psychological Review*, 112(4):777, 2005.
- [18] Anthony Fader, Luke Zettlemoyer, and Oren Etzioni. Paraphrase-driven learning for open question answering. In *Proc. Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1608–1618, 2013.
- [19] Rui Fang, Joyce Y Chai, and Fernanda Ferreira. Between linguistic attention and gaze fixations in multimodal conversational interfaces. In *Proc. International Conference on Multimodal Interfaces*, pages 143–150, 2009.
- [20] Katja Filippova, Enrique Alfonseca, Carlos A Colmenares, Łukasz Kaiser, and Oriol Vinyals. Sentence compression by deletion with lstms. In *Proc. Conference on Empirical Methods in Natural Language Processing*, pages 360–368, 2015.
- [21] Simone Frintrop, Erich Rome, and Henrik I Christensen. Computational visual attention systems and their cognitive foundations: A survey. *ACM Transactions on Applied Perception (TAP)*, 7(1):6, 2010.

- [22] Alex Graves and Jürgen Schmidhuber. Frameworkwise phoneme classification with bidirectional lstm and other neural network architectures. *Neural Networks*, 18(5-6):602–610, 2005.
- [23] Ankush Gupta, Arvind Agarwal, Prawaan Singh, and Piyush Rai. A deep generative framework for paraphrase generation. In *Proc. Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- [24] Vishal Gupta and Gurpreet Singh Lehal. A survey of text summarization extractive techniques. *Journal of Emerging Technologies in Web Intelligence*, 2(3):258–268, 2010.
- [25] Michael Hahn and Frank Keller. Modeling human reading with neural attention. In *Proc. Conference on Empirical Methods in Natural Language Processing*, pages 85–95, 2016. doi: 10.18653/v1/D16-1009.
- [26] Michael Hahn and Frank Keller. Modeling task effects in human reading with neural attention. *CoRR*, abs/1808.00054, 2018. URL <http://arxiv.org/abs/1808.00054>.
- [27] Sadid A Hasan, Bo Liu, Joey Liu, Ashequl Qadir, Kathy Lee, Vivek Datla, Aaditya Prakash, and Oladimeji Farri. Neural clinical paraphrase generation with attention. In *Proc. Clinical Natural Language Processing Workshop*, pages 42–53, 2016.
- [28] Karl Moritz Hermann, Tomas Kocisky, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. Teaching machines to read and comprehend. In *Proc. Advances in Neural Information Processing Systems*, pages 1693–1701, 2015.
- [29] Nora Hollenstein, Maria Barrett, Marius Troendle, Francesco Bigioli, Nicolas Langer, and Ce Zhang. Advancing nlp with cognitive language processing signals. *arXiv preprint arXiv:1904.02682*, 2019.
- [30] Ryu Iida, Masaaki Yasuhara, and Takenobu Tokunaga. Multi-modal reference resolution in situated dialogue by integrating linguistic and extra-linguistic clues. In *Proc. International Joint Conference on Natural Language Processing*, pages 84–92, 2011.
- [31] Hongyan Jing. Sentence reduction for automatic text summarization. In *Proc. Applied Natural Language Processing Conference*, pages 310–315, 2000.
- [32] Nour Karessli, Zeynep Akata, Bernt Schiele, and Andreas Bulling. Gaze embeddings for zero-shot image classification. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, pages 4525–4534, 2017.
- [33] Divyansh Kaushik and Zachary C Lipton. How much reading does reading comprehension require? a critical investigation of popular benchmarks. In *Proc. Conference on Empirical Methods in Natural Language Processing*, pages 5010–5015, 2018.
- [34] Alan Kennedy and Joël Pynte. Parafoveal-on-foveal effects in normal reading. *Vision research*, 45(2):153–168, 2005.
- [35] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *Proc. International Conference on Learning Representations*, 2015.
- [36] Sigrid Klerke and Barbara Plank. At a glance: The impact of gaze aggregation views on syntactic tagging. In *Proc. Beyond Vision and Language: Integrating Real-World Knowledge*, pages 51–61, 2019.
- [37] Sigrid Klerke, Yoav Goldberg, and Anders Søgaard. Improving sentence compression by learning to predict gaze. In *Proc. Conference of the North American Chapter of the Association for Computational Linguistics*, pages 1528–1533, 2016. doi: 10.18653/v1/N16-1179.
- [38] Kevin Knight and Daniel Marcu. Summarization beyond sentence extraction: A probabilistic approach to sentence compression. *Artificial Intelligence*, 139(1):91–107, 2002.
- [39] Iuliia Kotseruba and John K Tsotsos. 40 years of cognitive architectures: core cognitive abilities and practical applications. *Artificial Intelligence Review*, pages 1–78, 2018.

- [40] Matthias Kümmerer, Lucas Theis, and Matthias Bethge. Deep gaze i: Boosting saliency prediction with feature maps trained on imagenet. In *In International Conference on Learning Representations*, pages 1–12, 2015.
- [41] Zichao Li, Xin Jiang, Lifeng Shang, and Hang Li. Paraphrase generation with deep reinforcement learning. In *Proc. Conference on Empirical Methods in Natural Language Processing*, pages 3865–3878, October–November 2018. doi: 10.18653/v1/D18-1421.
- [42] Jianhua Lin. Divergence measures based on the shannon entropy. *IEEE Transactions on Information theory*, 37(1):145–151, 1991.
- [43] Steven G Luke and Kiel Christianson. The provo corpus: A large eye-tracking corpus with predictability norms. *Behavior Research Methods*, 50(2):826–833, 2018.
- [44] Minh-Thang Luong, Hieu Pham, and Christopher D Manning. Effective approaches to attention-based neural machine translation. In *Proc. Conference on Empirical Methods in Natural Language Processing*, pages 1412–1421, 2015.
- [45] Wei Ji Ma and Benjamin Peters. A neural network walks into a lab: towards using deep nets as models for human behavior. *arXiv preprint arXiv:2005.02181*, 2020.
- [46] Franz Matthies and Anders Søgaard. With blinkers on: Robust prediction of eye movements across readers. In *Proc. Conference on Empirical Methods in Natural Language Processing*, pages 803–807, 2013.
- [47] Ryan McDonald. Discriminative sentence compression with soft syntactic evidence. In *Proc. Conference of the European Chapter of the Association for Computational Linguistics*, 2006.
- [48] Volodymyr Mnih, Nicolas Heess, Alex Graves, et al. Recurrent models of visual attention. In *Proc. Advances in Neural Information Processing Systems*, pages 2204–2212, 2014.
- [49] Saleh Mozaffari, Pascal Klein, Jouni Viiri, Sheraz Ahmed, Jochen Kuhn, and Andreas Dengel. Evaluating similarity measures for gaze patterns in the context of representational competence in physics education. In *Proc. ACM Symposium on Eye Tracking Research & Applications*, pages 1–5, 2018.
- [50] Mattias Nilsson and Joakim Nivre. Learning where to look: Modeling eye movements in reading. In *Proc. Conference on Computational Natural Language Learning*, pages 93–101, 2009.
- [51] Mattias Nilsson and Joakim Nivre. Towards a data-driven model of eye movement control in reading. In *Proc. Workshop on Cognitive Modeling and Computational Linguistics*, pages 63–71. Association for Computational Linguistics, 2010.
- [52] Catharine Oertel and Giampiero Salvi. A gaze-based method for relating group involvement to individual engagement in multimodal multiparty dialogue. In *Proc. International Conference on Multimodal Interaction*, pages 99–106, 2013.
- [53] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. Bleu: a method for automatic evaluation of machine translation. In *Proc. Annual Meeting of Association for Computational Linguistics*, pages 311–318, 2002.
- [54] Badri Narayana Patro, Vinod Kumar Kurmi, Sandeep Kumar, and Vinay Namboodiri. Learning semantic sentence embeddings using sequential pair-wise discriminator. In *Proc. International Conference on Computational Linguistics*, pages 2715–2729, 2018.
- [55] Jeffrey Pennington, Richard Socher, and Christopher D Manning. Glove: Global vectors for word representation. In *Proc. Conference on Empirical Methods in Natural Language Processing*, pages 1532–1543, 2014.
- [56] Matt Post. A call for clarity in reporting bleu scores. In *Proc. Conference on Machine Translation*, pages 186–191, 2018.

- [57] Aaditya Prakash, Sadid A. Hasan, Kathy Lee, Vivek Datla, Ashequl Qadir, Joey Liu, and Oladimeji Farri. Neural paraphrase generation with stacked residual LSTM networks. In *Proc. International Conference on Computational Linguistics*, pages 2923–2934, December 2016.
- [58] Tingting Qiao, Jianfeng Dong, and Duanqing Xu. Exploring human-like attention supervision in visual question answering. In *Proc. Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- [59] Keith Rayner. Eye movements in reading and information processing. *Psychological bulletin*, 85(3):618, 1978.
- [60] Erik D Reichle, Alexander Pollatsek, Donald L Fisher, and Keith Rayner. Toward a model of eye movement control in reading. *Psychological review*, 105(1):125, 1998.
- [61] Erik D Reichle, Tessa Warren, and Kerry McConnell. Using ez reader to model the effects of higher level language processing on eye movements during reading. *Psychonomic bulletin & review*, 16(1):1–21, 2009.
- [62] Erik D Reichle, Simon P Liversedge, Denis Drieghe, Hazel I Blythe, Holly SSL Joseph, Sarah J White, and Keith Rayner. Using ez reader to examine the concurrent development of eye-movement control and reading skill. *Developmental Review*, 33(2):110–149, 2013.
- [63] Tim Rocktäschel, Edward Grefenstette, Karl Moritz Hermann, Tomáš Kociský, and Phil Blunsom. Reasoning about entailment with neural attention. In *Proc. International Conference on Learning Representations*, 2016.
- [64] Omid Rohanian, Shiva Taslimipoor, Victoria Yaneva, and Le An Ha. Using gaze data to predict multiword expressions. In *Proc. International Conference Recent Advances in Natural Language Processing*, pages 601–609, Varna, Bulgaria, September 2017. doi: 10.26615/978-954-452-049-6\_078.
- [65] Alexander M. Rush, Sumit Chopra, and Jason Weston. A neural attention model for abstractive sentence summarization. In *Proc. Conference on Empirical Methods in Natural Language Processing*, pages 379–389, September 2015. doi: 10.18653/v1/D15-1044.
- [66] Krasen Samardzhiev, Andrew Gargett, and Danushka Bollegala. Learning neural word saliency scores. In *Proc. Joint Conference on Lexical and Computational Semantics*, pages 33–42, 2018. doi: 10.18653/v1/S18-2004.
- [67] Mike Schuster and Kuldip K Paliwal. Bidirectional recurrent neural networks. *IEEE transactions on Signal Processing*, 45(11):2673–2681, 1997.
- [68] Minjoon Seo, Sewon Min, Ali Farhadi, and Hannaneh Hajishirzi. Neural speed reading via skim-RNN. In *Proc. International Conference on Learning Representations*, 2018.
- [69] Iaroslav Shcherbatyi, Andreas Bulling, and Mario Fritz. GazeDPM: Early integration of gaze information in deformable part models. *arXiv preprint arXiv:1505.05753*, 2015.
- [70] Ekta Sood, Simon Tannert, Diego Frassinelli, Andreas Bulling, and Ngoc Thang Vu. Interpreting attention models with human visual attention in machine reading comprehension. In *Proc. ACL SIGNLL Conference on Computational Natural Language Learning (CoNLL)*. Association for Computational Linguistics, 2020.
- [71] Yusuke Sugano and Andreas Bulling. Seeing with humans: Gaze-assisted neural image captioning. *arXiv preprint arXiv:1608.05203*, 2016.
- [72] Makarand Tapaswi, Yukun Zhu, Rainer Stiefelhagen, Antonio Torralba, Raquel Urtasun, and Sanja Fidler. MovieQA: Understanding Stories in Movies through Question-Answering. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, 2016.
- [73] Oguzhan Tas and Farzad Kiyani. A survey automatic text summarization. *PressAcademia Procedia*, 5(1):205–213, 2007.

- [74] Kristina Toutanova, Dan Klein, Christopher D. Manning, and Yoram Singer. Feature-rich part-of-speech tagging with a cyclic dependency network. In *Proc. Conference of the North American Chapter of the Association for Computational Linguistics*, page 173–180, 2003. doi: 10.3115/1073445.1073478.
- [75] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Proc. Advances in Neural Information Processing Systems*, pages 5998–6008, 2017.
- [76] Shaonan Wang, Jiajun Zhang, and Chengqing Zong. Learning sentence representation with guidance of human attention. In *Proc. International Joint Conference on Artificial Intelligence*, pages 4137–4143, 2017.
- [77] Zhiwei Wang, Yao Ma, Zitao Liu, and Jiliang Tang. R-transformer: Recurrent neural network enhanced transformer. *arXiv:1907.05572*, 2019.
- [78] Jeremy M Wolfe and Todd S Horowitz. Five factors that guide attention in visual search. *Nature Human Behaviour*, 1(3):1–8, 2017.
- [79] Menglin Xia, Ekaterina Kochmar, and Ted Briscoe. Automatic learner summary assessment for reading comprehension. In *Proc. Conference of the North American Chapter of the Association for Computational Linguistics*, pages 2532–2542, 2019. doi: 10.18653/v1/N19-1261.
- [80] Jia Xu, Lopamudra Mukherjee, Yin Li, Jamieson Warner, James M Rehg, and Vikas Singh. Gaze-enabled egocentric video summarization via constrained submodular maximization. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, pages 2235–2244, 2015.
- [81] Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhudinov, Rich Zemel, and Yoshua Bengio. Show, attend and tell: Neural image caption generation with visual attention. In *Proc. International Conference on Machine Learning*, pages 2048–2057, 2015.
- [82] Victoria Yaneva, Le An Ha, Richard Evans, and Ruslan Mitkov. Classifying referential and non-referential it using gaze. In *Proc. Conference on Empirical Methods in Natural Language Processing*, pages 4896–4901, Brussels, Belgium, October-November 2018. doi: 10.18653/v1/D18-1528.
- [83] Youngjae Yu, Jongwook Choi, Yeonhwa Kim, Kyung Yoo, Sang-Hun Lee, and Gunhee Kim. Supervising neural attention models for video captioning by human gaze data. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, pages 490–498, 2017.
- [84] Kiwon Yun, Yifan Peng, Dimitris Samaras, Gregory J Zelinsky, and Tamara L Berg. Studying relationships between human gaze, description, and computer vision. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, pages 739–746, 2013.
- [85] Sheng Zhang, Xiaodong Liu, Jingjing Liu, Jianfeng Gao, Kevin Duh, and Benjamin Van Durme. Record: Bridging the gap between human and machine commonsense reading comprehension. *arXiv preprint arXiv:1810.12885*, 2018.
- [86] Yingyi Zhang and Chengzhi Zhang. Using human attention to extract keyphrase from microblog post. In *Proc. Annual Meeting of the Association for Computational Linguistics*, pages 5867–5872, 2019.
- [87] Yang Zhao, Zhiyuan Luo, and Akiko Aizawa. A language model based evaluator for sentence compression. In *Proc. Annual Meeting of the Association for Computational Linguistics*, pages 170–175, 2018.