PERSPECTIVES GENERATION VIA MULTI-HEAD ATTENTION MECHANISM AND COMMON-SENSE KNOWLEDGE

Fatima Alkhawaldeh, Tommy Yuan and Dimitar Kazakov

Department of Computer Science, University of York, YO10 5GH, UK

ABSTRACT

Consideration of multiple viewpoints on a contentious issue is critical for avoiding bias and assisting in the formulation of rational decisions. We observe that the current model imposes a constraint on diversity. This is because the conventional attention mechanism is biased toward a single semantic aspect of the claim, whereas the claim may contain multiple semantic aspects. Additionally, disregarding common-sense knowledge may result in generating perspectives that violate known facts about the world. The proposed approach is divided into two stages: the first stage considers multiple semantic aspects, which results in more diverse generated perspectives; the second stage improves the quality of generated perspectives by incorporating common-sense knowledge. We train the model on each stage using reinforcement learning and automated metric scores. The experimental results demonstrate the effectiveness of our proposed model in generating a broader range of perspectives on a contentious subject.

KEYWORDS

Perspective generation, diversity, quality, common-sense and Multi-Head Attention Mechanism

1. INTRODUCTION

Individuals' assessments of factual truth vary due to their varying levels of subject knowledge and their linguistic abilities. Additionally, the rapid pace, enormous volume, and noise associated with data generated by users with questionable authorship and authenticity result in the emergence of emerging claims in a variety of domains, necessitating the consideration of alternative perspectives. It is not always possible to substantiate a claim with an authoritative source, especially when previously unmentioned claims are discovered. Viewing a claim through a singular lens may introduce bias. Without taking into account additional data, relying exclusively on textual information from a single source is likely to result in inaccuracies and bias. To address this issue, it is necessary to critically analyse a claim from multiple perspectives. Regrettably, there is a dearth of diverse perspectives on specific claims arising from unseen events for rapidly developing claims and responses.

An argumentative text's objective is to persuade the reader to concur with a particular conclusion. Each argument begins with a conclusion, followed by one or more supporting premises. By taking into account a diverse range of conclusions, i.e., perspectives relevant to a given claim for veracity prediction, analysing diverse arguments helps to alleviate the bias problem. To address this issue, the majority of current systems attempt to extract or generate the omitted conclusion from relevant evidence, which may contain a large amount of information about numerous topic aspects as well as additional data that supports or refutes the claim. The evidence retrieved in support of a particular claim may be associated with it, thereby reducing the model's diversity by limiting it to only visible evidence. As a result, we propose a new task for generating diverse and high-quality perspectives on a given claim about a contentious subject.

David C. Wyld et al. (Eds): NLDML, CCITT-2022
pp. 01-13, 2022. IJCI – 2022 DOI:10.5121/ijci.2022.110201

Diverse and high-quality perspectives are required for improving the prediction performance and robustness of end-uses, and thus optimising model prediction via cross-entropy loss alone is insufficient to encourage the model to generate diverse paraphrases. In relation to the cross-entropy loss Due to the fact that prediction is made against a single score, at the word level, strict sequence matching between the generated perspective and the ground truth perspective is required. Additionally, limiting perspectives to a single semantic aspect reduces diversity, as a claim may have multiple aspects. Even though common sense knowledge is critical for perspective formation, current systems overlook it. While conventional models generate generic responses, the perspectives they generate may contradict common sense.

We observe that prior work has a low-performance level, and in this proposed work, we seek to improve the performance, quality, and diversity of a state-of-the-art system in response to a specific claim. Additionally, as demonstrated by the state of art Park et al.'s model [1], it still performs poorly in terms of quality and diversity on some data sets. Park et al.'s model [1]generates N distinct perspectives and selects the one with the lowest negative log-likelihood NLL for the given reference perspective as the generated perspective. Even though they followed the multi-head attention work [2] and achieved state-of-the-art results in a perspective generation, we believe that multiple automated metrics as a reinforcement learning reward will further improve the approach to controlling perspective generation for optimization and will encourage the model to consider various factors that are necessary to improve the quality a during generation. Additionally, we hypothesise that incorporating common sense makes generated perspectives more plausible and does not violate world facts, which were not considered in the state-of-the-art model. As a result, we begin by introducing a novel technique for generating diverse and high-quality perspectives by focusing on various semantic aspects. The generated perspectives are then supplemented with a collection of common-sense facts. By utilising reinforcement learning, we can combine multiple learning objectives for model training.

The purpose of this paper is to examine the question, "Can we generate high quality multiple relevant candidate perspectives for a given claim? Our proposed model employs a multi-generator to generate perspectives. Our model develops the ability to generate multiple perspectives from its input, which is a claim, an argumentative sentence, and a reward function as an evaluator. the proposed model is trained and evaluated on datasets from Perspectrum [3]. Table 1 shows an example of our generator's generation of perspectives alongside a reference from the Perspectrum dataset.

Table 1: An example of perspectives generation

Claim: "A government should lessen the economic gap between its rich and poor citizens."	**Perspectives in the Perspectrum dataset**	**Correspondent perspectives generated by our model**
perspective 1	"True individual freedom cannot exist without economic security and independence."	True individual freedom cannot exist without monetary safety and independence
perspective 2	"The wealth gap does not allow for equality between the rich and the poor, and so it should be reduced."	The wealth hole does now not permit for equality between the rich and the poor and so it has to be lessened

2. RELATED WORKS

Numerous works have concentrated on identifying claims within the context of argument mining [4]. To accomplish this, the work in [5] demonstrates the critical nature of taking the conclusion and premises' primary objectives into account. Other studies use alternative methods for generating claims, such as opinion summarization, to capture the text's most salient points [6], [7]. While Egan et al.'s method [8] for summarising points made in online political debates relies on verbs and their syntactic arguments to identify silent information in political debates. From the premises, the conclusion (or claim) can be constructed, including its stance (for or against) the target [1], [9]. Hua & Wang [10] and Hua et al. [11] attempt to generate counterarguments to a given statement, whereas Wachsmuth et al. [12] and Hidey & McKeown [13] edited an original claim from the comments to generate new claims. Reisert et al. [14] summarise the text's main points using the Toulmin model and the relationships between the model's components and then use the summarised text to generate new claims automatically. In terms of the Toulmin argument model, Reisert et al. [14] construct complete arguments based on Toulmin's model [15], which requires that a claim be substantiated by data and justified by a warrant. They confine themselves to logical argument structure and grammatical rules to generate arguments about debate claims. Composing complex linguistic rules is challenging due to varying levels of knowledge about the language's syntactic structure and the requirement for extensive domain knowledge. In general, the task of generating perspectives relies on supporting evidence and continues to be difficult in terms of quality and diversity. In this work, we propose a novel approach to the task: Given a specific claim about a contentious issue, construct a logical set of perspectives with varying stances. Our primary contributions are as follows:

- Rather than generating a perspective based on a single semantic aspect, as traditional approaches do, we propose a Seq2Seq model with a multi-head attention mechanism that generates diverse perspectives based on the diverse semantic aspects.
- We incorporate common sense knowledge to ensure that the model does not violate known facts about the world and to improve the quality of the generated perspectives.
- We employ a reward function; multi-objective reinforcement learning produces various scores, ensuring that the generated perspectives make appropriate use of the given context and allowing control of the text generation model without relying on a single objective during the decoding process.
- Experiments show that our model outperforms several existing Seq2Seq-based perspectives models on quality and diversity metrics.

3. METHODOLOGY

In this section, we outline our generation processes and discuss how reinforcement learning can be used to further improve the technique rather than cross-entropy loss. reinforcement learning is a rapidly growing field of research that involves intelligent agents that learn to reason through Markov Decision Processes [16]. Recently, there has been significant progress in the field of reinforcement learning (DRL) for natural language processing, including relation extraction [17]and reasoning in question answering [18] and generation of paraphrases [19].

To begin, given the claim and some random words replaced with synonyms, the generator is contextualised and more diverse thanks to the wordnet corpus. Then, multiple semantic aspects are extracted, and the claim representation is learned [20][21] the modified version of the claim with new replaced synonyms words of claim with the semantic extracted aspect are used to guide the generation process to generate candidate perspectives for each semantic aspect. After that, for more information and a better-quality perspective, our models take common sense into account.

We used reinforcement learning to enforce diversity, stylistic, and quality constraints on the generated perspective. The overall architecture of our proposed model is shown in figure 1:

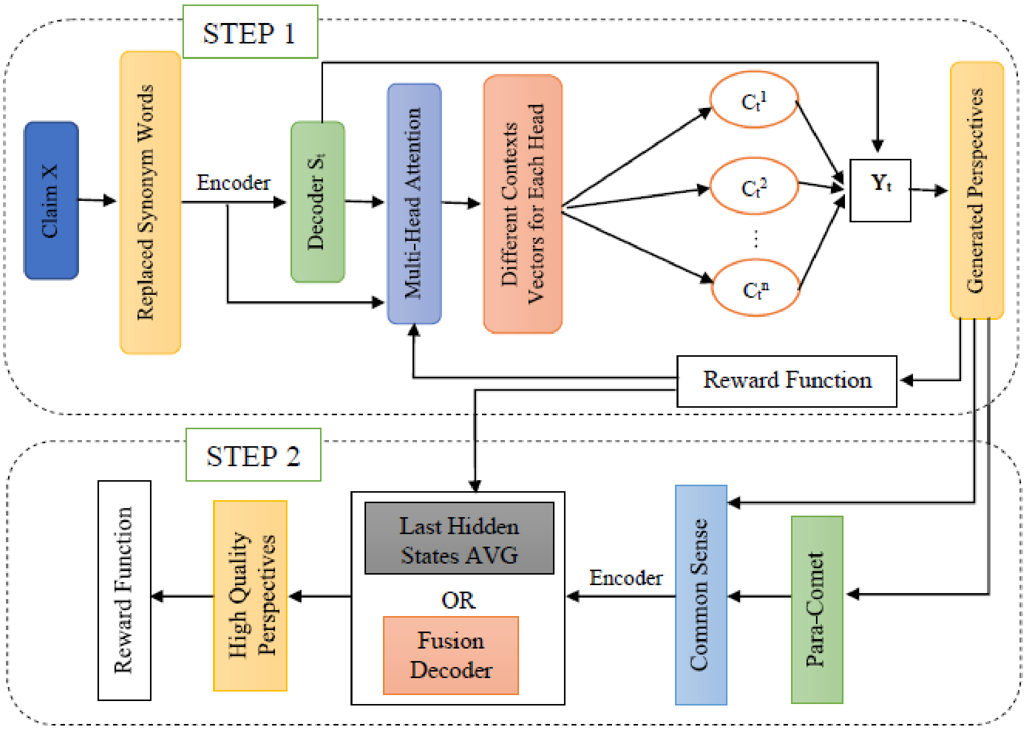

Figure 1: Proposed model architecture

3.1. First Stage: Various Semantic Aspects Considered While Generating Perspectives

3.1.1. The revised version of the claim includes new substituted synonyms

A modified version of a claim is created by replacing specific words in the claim input with synonyms at random. The claim's word sequence is used as input, with specific words being replaced by synonyms at a 60% ratio that fits the semantic aspect context.

3.1.2.Multi-head Attention Mechanism

Inspired by the multi-head attention with a Seq2Seq model [2], this paper employs a mechanism of multi-head attention to enable the generator to attend to information from different representation subspaces during the generating process, with context vectors obtained via the multi-head attention mechanism focusing on different semantic aspects of the text rather than on a single semantic aspect, as in traditional attention mechanisms, [20][21].

To generate n distinct perspectives on the claim, n distinct context vectors are created by projecting each state to multiple semantic spaces using various learnable projection matrices as in equation 1. The context vector for each head can then be produced by multiplying the encoder's hidden states by a weighted sum for all semantic spaces, the attention process is used to obtain numerous attention probability distributions over the claim words as illustrated in equations 2-5.

$$h_j^n = W_t^n . h_j \quad (1)$$

$$a_{t,j}^c = v_{cl} \cdot tanh\left(W_{cl} s_t + U_{cl} h_j^n\right) \quad (2)$$

$$\alpha_{tj}^{n} = \frac{exp\left(a_{t,j}^{n}\right)}{\sum_{j=1}^{|c|} exp\left(a_{t,j}^{n}\right)} \quad (3)$$

$$c_{t}^{n} = \sum_{i=1}^{|n|} \alpha_{t,j}^{n}\, h_{j}^{n} \quad (4)$$

$$s_t = GRU_{dec}\left(s_{t-1}, \left[c_{t-1}^{n} \quad , E_{(Y_{t-1})}\right]\right) \quad (5)$$

For each semantic space, there is W_t^nlearnable projection matrix, h_i hidden representation for each time-step for the new version of the claim word, $E_{(Y_{t-1})}$ is the previous word embeddings, st is the current state of the decoder at time step t, and $c_t^n$is Context vector for i-th head at a time step where it could be used to generate a i-th perspective that focuses on a particular semantic aspect of the claim. The hidden state of the decoder st at each time t is computed as follow, considering the previous state s_{t-1}, the previous claim context vector c_{t-1}^n and the previous word embeddings.

The probability distribution over the output vocabulary o_t, as equation 6 to decide the word which has the highest probability is computed from the context vector ct, and the decoder state st as equation 6, where $W_g^{(2)}$, $W_g^{(1)}$, $b_g^{(1)}$and $b_g^{(2)}$are learnable parameters:

$$o_t = W_g^{(2)}\left(W_g^{(1)}[s_t, c_t^n] + b_g^{(1)}\right) + b_g^{(2)} (6)$$

$p_t^{pointer}$ is used as a switch to select between [22] (a) copying words from the source text via pointing (copying a word from the input sequence by selection according to the attention distribution) or (b) generating a word from the vocabulary by selecting based on Pv as illustrated in equations 7-9 where v_{ptr}^T and $b_{pointer}$ are learnable parameters.

$$p_t^{pointer} = sigmoid\left(v_{ptr}^T\left[s_t, E_{(Y_{t-1})}, c_t^n\right] + b_{pointer}\right) (7)$$

The generation probability $p_{tj}^{gen} \in [0,1]$ for timestep t is computed as equation 8. If $p_{tj}^{gen} > 0.5$, word is copied from the input determined by the attention distribution where the attention is the highest, else the generator output is used. The probability of generating timestamp t is set to 0. 5 empirically.

$$p_{tj}^{gen} = \frac{exp(o_{tj})}{\sum_k exp(o_{tk})} (8)$$

The model then generates distribution Pv over vocabulary. P_v *is* probability distribution over all words in the vocabulary and gives us the final distribution to expect words. It concatenates the output of decoder st as the input of the output projection layer. T, it will show the details of these variables in equation 9 where W_v and $b_v$are learnable parameters.

$$P_v = softmax\left(W_v\left[s_t; E_{(Y_{t-1})}, c_t^n\right] + b_v\right) (9)$$

3.2 Second Stage: Conscious of Common-Sense Knowledge to Maintain a Higher-Quality Perspective

Common sense knowledge or world facts are required for the successful completion of a large number of natural language processing tasks [23][24]. Additional inferences based on common sense knowledge can be formed from a claim accompanied by a modified claim, hence improving the quality of the generated perspectives. To incorporate common-sense inferences into our model we rely on PARA-COMET [25]. As we have two sentences, and a modified claim with random replacements words, we feed this as an input to the trained PARA-COMET model, which generates nine common-sense relations for both sentences. For each perspective we have common-sense, for example, [perspective-1, common-sense-1, perspective-2], [perspective-1, common-sense-2, perspective-3], … [perspective-1, common-sense-n-1, perspective-n] and so on for all perspectives. So, to enhance the generated perspective-1, all common-sense relation is encoded. PARA-COMET provides a set of commonsense inferences for the 9 inferential relations from ATOMIC for each perspective, based on n-perspectives p1, p2, … pn, that is consistent with the complete narrative. To achieve that, we try two different models that consider the common-sense aiming at enhancing the quality of the generated perspectives. The first enhancer model averages the last hidden states for all related common-sense to i-th perspective and update the context vector. The second enhancer model makes use of a "fusion-in decoder"[26] that is supplemented with common-sense sentences retrieved from external knowledge.

3.2.1. Model 1 of the Enhancer

Each agent takes the encoded information $h_i^{(k)}$ from its encoder, which represents a particular generated perspective from the first stage. It considers other agents' information common-sense relations by averaging the last hidden states of other encoders $h_{m,I}^{(k)}$, to produce other important information $v^{(k)}$. An attention vector $f\left(h_i^{(k)}, v^{(k)}\right)$ is produced by considering its encoded feature $h_i^{(k)}$, previous decoder state s^{t-1} and other $v^{(k)}$. Finally, the context vector $c_{tj}^{k}$is updated based on attention distribution a_{tj}^{k}. Then apply the pointer attention method as in the first stage. The steps are as follows:

- The average of last hidden states for the encoded knowledge common-sense relations as in equation 10:

$$v^{(k)} = \frac{1}{M-1} \sum_{m \neq \alpha} h_{m,I}^{(k)} (10)$$

- Update context vector for each agent, as in equations 11, 12 and 13:

$$f\left(h_i^{(k)}, v^{(k)}\right) = v_1^T \, tanh\left(W_3 h_i^{(k)} + s^{t-1} + W_4 v^{(k)} + clt\right) (11)$$

$$a_{tj}^{k} = \frac{exp(f_{tj})}{\sum_{k=1}^{l} \, exp(f_{tk})} (12)$$

$$c_{tj}^{k} = \sum_{j=1}^{n} \quad a_{tj_i}^{k} h_i (13)$$

W_{n} are parameters of weights, b_v, v_1^T, and W_3 are learnable parameters

3.2.2. Model 2 of the Enhancer:

We use the Fusion in Decoder [26] in this enhancer model, a sequence-to-sequence model that accepts as input a previously generated perspective and a set of common-senses from a PARA-COMET [25]. It produces high-quality work that adheres to accepted world facts. Given a perspective with n-1 common senses in support, each common sense is concatenated with the perspective to produce perspectives–common sense contexts. fi = [pi; sj], where fi is encoded separately, but in the decoder, the encodings are combined to produce a higher-quality perspective.

3.3. Reward Function

As our algorithm attempts to provide varied diverse and high-quality perspectives with distinct stances. So, we use a composite score generated by averaging the specific measures to generate text under various conditions. the average of the individual metrics includes ROUGE, textual entailment, Style **control reward, stance control reward,** diversity and fluency provide a normalised score between 0 and 1. The perspectives are fed to evaluation modules

3.3.1. ROUGE Reward with Reference

To compare the degree to which the generated perspective retains context, it is rewarded using the ROUGE package's primary evaluation metric [27] and the score is then used as a reward. The ROUGE measures the unigram overlap, bigram overlap, and longest common sub-sequence between the predicted and reference [27].

- ROUGE-1: the unigram overlaps describe the overlap of each word between the candidate and reference summaries.
- ROUGE-2: bigram-overlap between the reference summary and the summary to be assessed.
- ROUGE-L: the longest common subsequence between the reference summary and the summary to be assessed.

3.3.2. Text Entailment

In terms of supporting perspectives, it should have a higher degree of entailment and a greater reward, whereas attacked perspectives should have a lower degree of entailment and a greater reward. We evaluate our generators using entailment metrics to determine whether the generated perspectives are inferable from (influenced by) the underlying claim. To measure textual entailment, we use a ranking-based loss function to train a model that generates a space embedding for claim contexts and generated perspectives [28]

3.3.3. Style Control Reward

We feed the generator with additional style embeddings and can calculate the probability of the output condition based on the style control variable. The rewarder is a convolutional neural network that has been trained to minimise cross-entropy loss in style classification so that the classifier can learn to correctly classify text styles.

3.3.4. Stance Module:

We can calculate the probability of the output condition based on the stance control variable by feeding the generator with additional stance embeddings. The classifier can learn to correctly classify text stance by using a convolutional neural network that has been trained to minimise cross-entropy loss in stance classification.

3.3.5. Fluency

It demonstrates the naturalness of the generated perspectives by measuring the grammatical correctness to increase the probability of the target sentences being used. Each generated perspective is assigned a perplexity level (PPL) by the language model. The less perplexing a perspective is, the more fluent it becomes. We propose to use GPT-2 [29] large-scale pre-trained language models for fluency which is suitable for likelihood-based fluency evaluation and conditional generation.

3.3.6. Expression Diversity

Allowing for a wide range of linguistic variations to be captured [30]. Self-BLEU is a tool we use to assess diversity, with a higher score indicating greater diversity. According to BLEU, it calculates the BLEU score for each generated sentence by comparing it to previously generated sentences. By averaging these BLEU scores (for generated sentences), a metric called Self-BLEU is created, with lower values indicating greater diversity.

4. PERSPECTRUM DATASET

The authors (Chen et al.) in [3], [31] have developed a dataset for the task of validated perspective discovery, PERSPECTRUM dataset, which is a set of claims, perspectives and evidence from online websites. In this dataset, each claim has a supported or opposed stance to the given claim supported by enough evidence. The statistic information about the PERSPECTRUM dataset is shown in table 2.

Table 2: A summary of PERSPECTRUM statistics [32]

Split	**Supporting Pairs**	**Opposing Pairs**	**Total Pairs**
Training	3603	3404	7007
Validation	1051	1045	2096
Test	1471	1302	2773
Total	**6125**	**5751**	**11876**

For example, claim A has a supported relationship with perspective A, while claim B has a refuted relationship with perspective B, perspectives are generated based on claim text. In other words, rewording claims have supported or undermined relations with perspectives that have supporting evidence.

Claim A: "A government must lessen the economic gap between its rich and poor citizens".

Perspective A:" The rich-poor gap silences the political voice of the poor".

Evidence A:" Research has also demonstrated a connection between economic inequality and political voice. The political process is far more responsive to the privilege’s claims, and the privileged are better organized and engaged in the political process than are less affluent citizens. Recent studies show that government officials are far more likely to support the wealthy's policy preferences than those of the poor. In short, there is considerable evidence to suggest that there is a growing divide between those who have wealth and political influence and those who do not. Yasmin Dawood, the new inequality: constitutional democracy and the problem of wealth, Maryland Law Review: 2007".

Claim B: "Internet access is a human right".

Perspective B:" It is a big problem; too many people are file-sharing".

Evidence B:" The plan to slow down or stop internet connections is the most economical and practical way to deal with file-sharers. Many illegal downloaders are young people, and this plan will prevent the offenders from receiving a criminal record".

Figure 2 shows an example of annotated data on PERSPECTRUM.

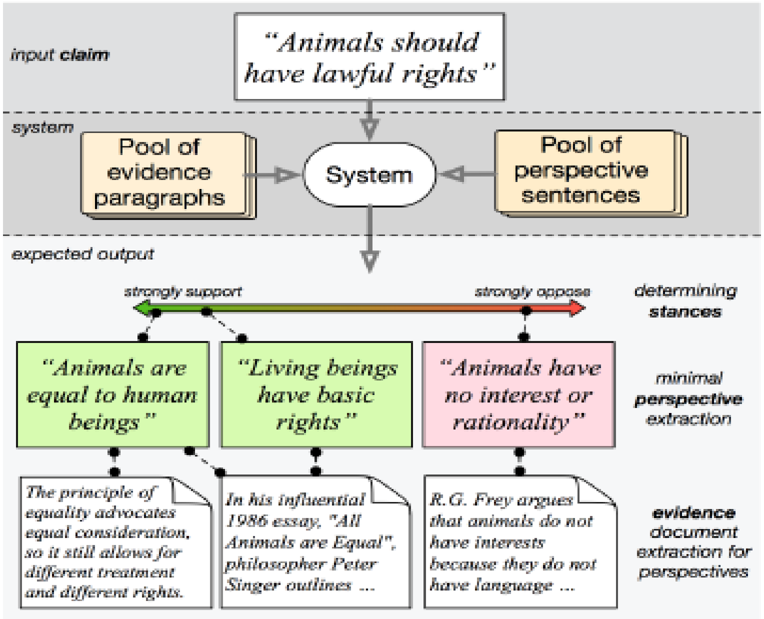

Figure 2: An example of annotated data on PERSPECTRUM [3]

5. EXPERIMENTS AND RESULTS

To evaluate our proposed model, we compare our model to Park et al.'s model [1], which was trained and evaluated on the Perspectrum dataset. The results in tables 3 and 4 show that the multi-agent model outperforms baselines in terms of automatic evaluation metrics, diversity, and quality by taking into account various semantic aspects and Common-sense knowledge. We show how leveraging the reinforcement learning reward function improves the perspectives generator model performance of a state-of-the-art model. A multi-agent model, where the decoder network learns from the different semantic aspect vectors during the decoding stage, may capture more realistic arguments than a baseline model. By pooling the common-sense knowledge of multiple agents, the multi-Agent model can capture richer data from multiple perspectives and cover a broader range of issues. Our model can generate high-quality, diverse, and multiple arguments based on the metrics results compared to baseline models. We observe that our model outperforms competitors in all metrics when BLEU score and word embedding-based metrics are used. We achieve the best performance in four metrics (Dist-1, Dist-2, and Dist-1/2-within) for diversity.

Park et al. generate claims in response to a given claim, utilising a diversity penalty to encourage the presentation of diverse perspectives. It utilises a Seq2Seq framework and introduces latent mechanisms on the assumption that each latent mechanism can be associated with a single perspective.

- BLEU-1/2: measures N-gram precision of the generated text to multiple target arguments references [27].

- Embedding Average/Greedy/Extreme: measures the semantic similarity between hypothesis and references, using a semantic representation by word embedding [34]
- Dist-1/2: computes the percentage of unique unigrams/bigrams within a sentence to measure the diversity among multiple generated texts [28].
- Dist-1/2-within [1], propose a simple metric to calculate the sum of the numbers of unique N-grams for each result that does not occur in other results) / (The sum of all generated numbers of unigrams/bigrams).

Table 3: Automatic Evaluation Results for Perspectives Generation Quality on Perspectrum Dataset

Method	**BLUE 1**	**BLUE 2**	**Embedding Average**	**Embedding Greedy**	**Embeddin g extreme**
Generator: Pointer attention - Onley one semantic aspect	0.2635	0.0684	0.6838	0.4858	0.2810
First stage of the Generator: Various semantic aspects of pointer attention	0.3127	0.0919	0.7211	0.5801	0.4139
ArgDiver [1]	0.3268	0.0964	0.8107	0.6002	0.4146
Second stage of the Generator: Using the average of last states to incorporate common sense following the first stage	0.3528	0.1027	0.8329	0.6324	0.4476
Second stage of the Generator: Using fusion decoder to incorporate common sense following the first stage	0.3618	0.1096	0.8514	0.6514	0.4526
Generator rewarded by RL function	**0.3955**	**0.1183**	**0.8801**	**0.6665**	**0.4918**

Table 4: Automatic Evaluation Results on the Diversity of Perspectives Generation on Perspectrum Dataset

Method	**Dist-1**	**Dist-2**	**Dist-1-within**	**Dist-2-within**
Generator: Pointer attention-only one semantic aspect	0.1328	0.1983	0.2814	0.4612
First stage of the Generator: Various semantic aspects of pointer attention	0.1420	0.2727	0.3529	0.6051
ArgDiver [1]	0.1585	0.2909	0.3645	0.6134
Second stage of the Generator: Using the average of last states to incorporate common sense following the first stage	0.1603	0.3086	0.4066	0.6272
Second stage of the Generator: Using fusion decoder to incorporate common sense following the first stage	0.1681	0.3126	0.4182	0.6423
Generator rewarded by RL function	**0.1703**	**0.3208**	**0.4461**	**0.7006**

6. AN ABLATION STUDY WITH AUTOMATED EVALUATION METRIC SCORES: QUALITY, AND DIVERSITY

We investigate our model in depth in this section to develop perspectives for use in an ablation study. The ablation findings are summarised in Tables 3 and 4. We begin with the pointer attention model, which concentrates on a single semantic aspect; the findings indicate that the model achieves the least performance. Our proposed generator has two stages. In the first stage, the generator's effectiveness is evaluated in terms of quality and variety using a metric that considers multiple heads of attention from the model to generate more diverse perspectives. As seen in Tables 3 and 4, on both metrics, it exceeds pointer attention, which concentrates on the same semantic aspect. In the second step, the generator incorporates semantic aspects and common-sense knowledge, and when compared to the outputs of the first stage, the latter outperforms the earlier. This confirms that incorporating world fact makes the generated perspectives more plausible compared to the ones generated in the first stage. The findings of the two independent models are used to compel the generator to combine common sense from various knowledge passages in separate encoders, demonstrating that decoder fusion performs better than taking the average of these knowledge's final hidden states. Experiments and findings reveal that the reinforcement learning-based technique is capable of effectively learning to generate diverse and high-quality paraphrases and greatly increases generation quality when compared to numerous state-of-the-art baselines.

7. Conclusion

This article discusses the generation of perspectives via the employment of several heads of attention to analyse various semantic aspects of a claim. Additionally, we addressed how to maximise the benefits of utilising various common-sense pieces of information. Our approach is capable of generating a variety of high-quality viewpoints on a given claim using a variety of distinct postures, Additionally, we conclude that our reward function advances the state of the art in perspective generation and signals the generation of viewpoints with a specific stance. We examined each step independently in our experiments. We compared our overall strategy to the state-of-the-art approach described in Park et al. [1], using automated evaluation ratings. The results demonstrate that the proposed RL is significantly more performant than a state-of-the-art perspective generative model when considering different semantic aspects and encoding aware common-sense knowledge.

References

[1] C. Park, W. Yang, and J. Park, "Generating Sentential Arguments from Diverse Perspectives on Controversial Topic," in Proceedings of the 2nd Workshop on NLP for Internet Freedom: Censorship, Disinformation, and Propaganda, 2019, vol. November 4, pp. 56–65. doi: 10.18653/v1/d19-5007.

[2] C. Tao, S. Gao, M. Shang, W. Wu, D. Zhao, and R. Yan, "Get The Point of My Utterance! Learning Towards Effective Responses with Multi-Head Attention Mechanism," in Proceedings of the 27th International Joint Conference on Artificial Intelligence, IJCAI'18:, 2018, pp. 4418–4424.

[3] S. Chen, D. Khashabi, W. Yin, C. Callison-Burch, and D. Roth, "Seeing Things from a Different Angle: Discovering Diverse Perspectives about Claims," in NAACL HLT 2019 - 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies - Proceedings of the Conference, 2019, vol. 1, pp. 542–557. doi: 10.18653/v1/n19-1053.

[4] C. Stab and I. Gurevych, "Identifying Argumentative Discourse Structures in Persuasive Essays," in EMNLP 2014 - 2014 Conference on Empirical Methods in Natural Language Processing, Proceedings of the Conference, 2014, no. October, pp. 46–56. doi: 10.3115/v1/d14-1006.

[5] M. Alshomary, S. Syed, M. Potthast, and H. Wachsmuth, "Target Inference in Argument Conclusion Generation," in Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, 2020, no. 1, pp. 4334–4345. doi: 10.18653/v1/2020.acl-main.399.

[6] Y. Bilu and N. Slonim, "Claim synthesis via predicate recycling," in Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL 2016 - Short Papers, 2016, pp. 525–530. doi: 10.18653/v1/p16-2085.

[7] L. Wang and W. Ling, "Neural network-based abstract generation for opinions and arguments," in 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL HLT 2016 - Proceedings of the Conference, 2016, pp. 47–57. doi: 10.18653/v1/n16-1007.

[8] C. Egan, A. Siddharthan, and A. Wyner, "Summarising the Points Made in Online Political Debates," in Proceedings of the 3rd Workshop on Argument Mining (ArgMining2016), 2016, pp. 134–143. doi: 10.18653/v1/w16-2816.

[9] R. Bar-Haim, I. Bhattacharya, F. Dinuzzo, A. Saha, and N. Slonim, "Stance Classification of Context-Dependent Claims," in Proceedings ofthe 15th Conference ofthe European Chapter ofthe Association for Computational Linguistics: Long Papers, 2017, vol. 1, no. April, pp. 251–261. doi: 10.18653/v1/e17-1024.

[10] X. Hua and L. Wang, "Neural argument generation augmented with externally retrieved evidence," in ACL 2018 - 56th Annual Meeting of the Association for Computational Linguistics, Proceedings of the Conference (Long Papers), 2018, vol. 1, pp. 219–230. doi: 10.18653/v1/p18-1021.

[11] X. Hua, Z. Hu, and L. Wang, "Argument generation with retrieval, planning, and realization," in ACL 2019 - 57th Annual Meeting of the Association for Computational Linguistics, Proceedings of the Conference, 2019, pp. 2661–2672. doi: 10.18653/v1/p19-1255.

[12] H. Wachsmuth, M. Stede, R. El Baff, K. Al-Khatib, M. Skeppstedt, and B. Stein, "Argumentation Synthesis following Rhetorical Strategies," in Proceedings of the 27th International Conference on Computational Linguistics, Coling, 2018, vol. August, pp. 3753–3765.

[13] C. Hidey and K. McKeown, "Fixed that for you: Generating contrastive claims with semantic edits," NAACL HLT 2019 - 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies - Proceedings of the Conference, vol. 1, pp. 1756–1767, 2019, doi: 10.18653/v1/n19-1174.

[14] P. Reisert, N. Inoue, N. Okazaki, and K. Inui, "A Computational Approach for Generating Toulmin Model Argumentation," in Proceedings of the 2nd Workshop on Argumentation Mining, 2015, vol. June, pp. 45–55. doi: 10.3115/v1/w15-0507.

[15] S. E. Toulmin, The Uses of Argument. Cambridge, UK: Cambridge University Press, 1958.

[16] V. Mnih et al., "Human-level control through deep reinforcement learning," Nature, vol. 518, no. 7540, pp. 529–533, Feb. 2015, doi: 10.1038/nature14236.

[17] P. Qin, W. Xu, and W. Y. Wang, "Robust Distant Supervision Relation Extraction via Deep Reinforcement Learning," in Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), Jul. 2018, pp. 2137–2147. doi: 10.18653/v1/P18-1199.

[18] W. Xiong, T. Hoang, and W. Y. Wang, "DeepPath: A Reinforcement Learning Method for Knowledge Graph Reasoning," in Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, Sep. 2017, pp. 564–573. doi: 10.18653/v1/D17-1060.

[19] Z. Li, X. Jiang, L. Shang, and H. Li, "Paraphrase Generation with Deep Reinforcement Learning," in Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Nov. 2018, pp. 3865–3878. doi: 10.18653/v1/D18-1421.

[20] A. Vaswani et al., "Attention Is All You Need," in Proceeding of the 31st Conference on Neural Information Processing Systems (NIPS 2017), 2017, pp. 5998–6008.

[21] Z. Lin et al., "A Structured Self-attentive Sentence Embedding," Mar. 2017. [Online]. Available: http://arxiv.org/abs/1703.03130

[22] J. Hasselqvist, N. Helmertz, and M. Kågebäck, "Query-Based Abstractive Summarization Using Neural Networks," CoRR, abs/1712.06100, 2017.

[23] B. Wang, K. Liu, and J. Zhao, "Conditional Generative Adversarial Networks for Commonsense Machine Comprehension," in Proceedings of the 26th International Joint Conference on Artificial Intelligence, IJCAI'17, 2017, pp. 4123–4129. doi: https://doi.org/10.24963/ijcai.2017/576.

[24] H. Zhou, T. Young, M. Huang, H. Zhao, J. Xu, and X. Zhu, "Commonsense Knowledge Aware Conversation Generation with Graph Attention," in Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence (IJCAI-18), 2018, pp. 4623–4629. doi: https://doi.org/10.24963/ijcai.2018/643.

[25] S. Gabriel, C. Bhagavatula, V. Shwartz, R. le Bras, M. Forbes, and Y. Choi, "Paragraph-level Commonsense Transformers with Recurrent Memory," in Proceedings of the AAAI Conference on Artificial Intelligence, Oct. 2021, vol. 35, no. 14, pp. 12857–12865. Accessed: Nov. 21, 2021. [Online]. Available: http://arxiv.org/abs/2010.01486

[26] G. Izacard and E. Grave, "Leveraging Passage Retrieval with Generative Models for Open Domain Question Answering," in Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics, Apr. 2021, vol. Main Volum, pp. 874–880. doi: 10.18653/v1/2021.eacl-main.74.

[27] C.-Y. Lin, "ROUGE: A Package for Automatic Evaluation of Summaries," in Proceedings of the ACL-04 Workshop on Text Summarization Branches Out (WAS 2004), 2004, no. July, pp. 74–81. [Online]. Available: https://www.aclweb.org/anthology/W04-1013

[28] N. Reimers and I. Gurevych, "Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks," in Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processingand the 9th International Joint Conference on Natural Language Processing, Nov. 2019, pp. 3982–3992. [Online]. Available: https://github.com/UKPLab/

[29] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, and I. Sutskever, "Language Models are Unsupervised Multitask Learners," OpenAI blog, vol. 1, no. 8, p. 9, Feb. 2019, [Online]. Available: https://github.com/codelucas/newspaper

[30] Y. Zhu et al., "Texygen: A benchmarking platform for text generation models," in 41st International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR, Jul. 2018, pp. 1097–1100. doi: 10.1145/3209978.3210080.

[31] S. Chen, D. Khashabi, C. Callison-Burch, and D. Roth, "PERSPECTROSCOPE: A Window to the World of Diverse Perspectives," in Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics: System Demonstrations, 2019, pp. 129–134. doi: 10.18653/v1/p19-3022.

[32] K. Popat, S. Mukherjee, A. Yates, and G. Weikum, "STANCY: Stance Classification Based on Consistency Cues," in Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP), 2019, pp. 6413–6418. doi: 10.18653/v1/D19-1675.

[33] K. Papineni, S. Roukos, T. Ward, and W.-J. Zhu, "BLEU: a Method for Automatic Evaluation of Machine Translation," in Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics (ACL), 2002, no. JULY, pp. 311–318. doi: 10.1002/andp.19223712302.

[34] C. W. Liu, R. Lowe, I. v. Serban, M. Noseworthy, L. Charlin, and J. Pineau, "How NOT To Evaluate Your Dialogue System: An Empirical Study of Unsupervised Evaluation Metrics for Dialogue Response Generation," in EMNLP 2016 - Conference on Empirical Methods in Natural Language Processing, Proceedings, 2016, pp. 2122–2132. doi: 10.18653/v1/d16-1230.