

# Belief-based Generation of Argumentative Claims

Milad Alshomary   Wei-Fan Chen   Timon Gurcke   Henning Wachsmuth

Department of Computer Science  
Paderborn University, Paderborn, Germany  
<first name>.<last name>@upb.de

## Abstract

When engaging in argumentative discourse, skilled human debaters tailor claims to the audience’s beliefs to construct effective arguments. Recently, the field of computational argumentation witnessed extensive effort to address the automatic generation of arguments. However, existing approaches do not perform any audience-specific adaptation. In this work, we aim to bridge this gap by studying the task of *belief-based claim generation*: Given a controversial topic and a set of beliefs, generate an argumentative claim tailored to the beliefs. To tackle this task, we model the people’s prior beliefs through their stances on controversial topics and extend state-of-the-art text generation models to generate claims conditioned on the beliefs. Our automatic evaluation confirms the ability of our approach to adapt claims to a set of given beliefs. In a manual study, we also evaluate the generated claims in terms of informativeness and their likelihood to be uttered by someone with a respective belief. Our results reveal the limitations of modeling users’ beliefs based on their stances. Still, they demonstrate the potential of encoding beliefs into argumentative texts, laying the ground for future exploration of audience reach.

## 1 Introduction

According to [van Eemeren and Houtlosser \(1999\)](#), debaters engaging in argumentative discourse, aimed to resolve a disagreement, design their next argumentative move considering the topical potential, the audience demand, and appropriate presentational devices. [Feinberg and Willer \(2015\)](#) stress based on the moral foundation theory ([Godden, 2010](#)) how phrasing arguments to fit the audience’s morals leads to a better agreement. For example, in a debate on former US president *Donald Trump*, potential topics could have been *immigration, health care plans, tax plans*, etc. However, knowledge

about the audience being middle-class workers would have restricted the selection to *Trump’s tax plans*. Appropriate usage of presentational devices may have then put a con argument as follows:

**Example** “*Donald Trump was a bad president. He did nothing but hurt the poor and middle class. His tax plan benefited only rich people who could afford it.*”

There is a recent growth of interest in argument generation as a subfield of computational argumentation. Several tasks have been proposed, including claim negation ([Bilu et al., 2015](#); [Hidey and McKeown, 2019](#)), counterargument generation ([Hua et al., 2019](#)), and conclusion generation ([Alshomary et al., 2020](#)). While some research considers argumentative strategies when delivering arguments ([Wachsmuth et al., 2018](#); [El Baff et al., 2019](#)), no one has worked on adapting arguments to user beliefs yet. Our goal is to bridge this gap.

In this work, we propose to extend argument generation technologies with the ability to encode beliefs. This not only reflects the process by which humans reason, but it also allows controlling the output, to better reach the audience. In particular, we introduce the task of *belief-based claim generation*: Given a controversial topic and a representation of a user’s beliefs, generate a claim relevant to the topic and matches the beliefs.

To approach this task, we first model user beliefs by their stances (pro or con) on a set of controversial topics, and then extend two state-of-the-art text generation approaches by conditioning their output on a specific set of beliefs. One approach builds on [Li et al. \(2016\)](#), equipping a sequence-to-sequence (Seq2Seq) model with a context vector representing the given stances. The other approach controls the output of a pre-trained argumentative language model (LM) using the algorithm of [Dathathri et al. \(2020\)](#) to ensure resembling the user’s beliefs. We

study the given task empirically on the *debate.org* dataset of Durmus and Cardie (2018). The dataset contains users’ arguments on various controversial topics and their stances towards the most popular topics on the website, called the *big issues*. For our purposes, we use these big issues as the controversial topics, and we model beliefs by the user’s stances towards them.

In our automatic evaluation, we compare both models against their unconditioned correspondents (i.e., the same models without knowledge about a user). We assess the generated claims in terms of the similarity to the ground truth and the likelihood of carrying textual features that reflect users’ stances on big issues. Our results suggest that using users’ beliefs significantly increases the effectiveness of the Seq2Seq and LM in most cases. Moreover, a stance classifier trained on claims generated by the conditioned LM achieves the best-averaged accuracy across all big issues.

In a subsequent manual evaluation, we find that claims generated by the conditioned LM are more informative regarding the topic. In terms of predicting stance from generated claims, we analyze the limitations of our approach in detail, which lie in the belief encoding step. By avoiding these limitations, we find that the generated claims enable the annotators to predict correctly a stance on a given big issue in 45% of the cases (26% incorrectly). These results demonstrate the applicability of encoding a user’s beliefs into argumentative texts, enabling future research on the effect of belief-based argumentative claims on audiences.

The contribution of this work is threefold<sup>1</sup>:

- A new task, *belief-based claim generation*.
- An approach to model and match users’ beliefs in the generation of arguments.
- Empirical evidence of the applicability of encoding beliefs into argumentative texts.

## 2 Related Work

Early research on argument generation aimed to create argumentative texts starting from a symbolic representation (Zukerman et al., 2000; Grasso et al., 2000; Carenini and Moore, 2006). Conceptually, those approaches all had a similar architecture consisting of three main phases: text planning, sentence planning, and realization (Stede et al., 2018).

<sup>1</sup>Code can be found under: <http://www.github.com/webis-de/eacl21-belief-based-claim-generation>

While they included a user model to a certain extent and aimed to generate convincing arguments, they were still performed on a limited scale.

With the tremendous advances of NLP and machine learning since then, research has begun to address different tasks in the realm of argument generation, showing promising results. Hua et al. (2019) proposed a neural network-based framework for generating counter-arguments. Both Bilu et al. (2015) and Hidey and McKeown (2019) addressed the task of claim negation, using a rule-based and a neural approach respectively. Also, Sato et al. (2015) proposed an approach to argument generation based on sentence retrieval, in which, given a topic, a set of paragraphs covering different aspects is generated. However, these approaches are agnostic to the target audience.

Chen et al. (2018) modified the political bias of (often claim-like) news headlines using style transfer, accounting for general political sides (left and right) at least. Moreover, Wachsmuth et al. (2018) modeled rhetorical strategies in argument synthesis conceptually, but its computational realization (El Baff et al., 2019) considers the audience implicitly only, using a language model approach to select and arrange argumentative discourse units that are phrased in an argument.

In the field of conversational AI, researchers have utilized machine translation techniques to tackle the task of dialog generation (Ritter et al., 2011). Li et al. (2016) worked on augmenting sequence-to-sequence models by learning persona vectors from the given data. In a similar fashion, one of our approaches extends such a model by a context vector representing a user’s belief. Here, however, we deal with argumentative text.

Progress in the field of text generation has been made due to the availability of large pre-trained language models (Devlin et al., 2018; Solaiman et al., 2019). While these models excel in generating coherent texts, ensuring a generated text possesses a certain property is not straightforward. Some research tackled this limitation, offering ways to better control the output (Keskar et al., 2019; Ziegler et al., 2019). One of the most flexible of such approaches is by Dathathri et al. (2020), which does not require fine-tuning for each controlling theme. Their algorithm conditions the output of a language model to contain certain properties defined by a discriminative classifier or a bag-of-words. One of our approaches makes use of this algorithm to condi-

tion the output of an argumentative language model on a bag-of-words that represents a user’s beliefs. A recent relevant work by Schiller et al. (2020) deals with the generation of aspect-controlled arguments. Similar to us, the authors utilize a pre-trained language model to generate arguments on a specific topic, with a controlled stance and aspect. Their focus is on topical aspects of arguments, though, and their approach based on Keskar et al. (2019) is limited to a predefined set of topics and aspects.

### 3 Task

Due to the importance of audience in argumentation when aiming for persuasiveness (van Eemeren and Houtlosser, 1999), and due to the fact that humans comply to certain morals that shape their beliefs and affect their reasoning (Godden, 2010; Feinberg and Willer, 2015), we introduce the audience’s beliefs as a new dimension to the argument generation process in this work. For this, we propose a new task, belief-based claim generation:

*Given a controversial topic and a representation of the audience’s beliefs, generate a claim that is both relevant to the topic and matches the beliefs.*

We focus this task on generating *claims* rather than full arguments to keep it simple and because claims denote the main units from which arguments are built. As shown by Feinberg and Willer (2015), better agreement is achieved when arguments are framed with respect to audience’s beliefs. Therefore, we argue that studying the mentioned task will enable argumentation technology, knowing its audience, to generate more convincing arguments, bridging the gap between disagreeing parties.

#### 3.1 Data

To study the proposed task, a dataset is needed in which information about users revealing their beliefs as well as their arguments on various topics are given. Here, we build upon the dataset introduced by Durmus and Cardie (2018), which was collected from *debate.org*, an online platform where users can engage in debates over controversial topics and share their profiles. The dataset contains users’ arguments as answers to topic questions and engagement in debates, along with various user information, including a user’s self-specified stances (pro or con) on up to 48 predefined popular controversial topics, called *big issues*.

Dataset	# Claims	# Topics	# Users
Training set	41 288	22 241	5 189
Validation set	5 028	2 450	2 509
Test set	5 154	2 728	2 512
Full dataset	51 470	27 419	5 189

Table 1: Number of claims, topics, and users in each of the training, validation, and test set of the data used in this paper.

In our dataset, for the task at hand, we keep only users who have at least three arguments and stated their stance on at least one of the big issues. For those, we collected their arguments along with the topics and stances. In total, the dataset contains around 51k claims, on 27k topics from 5k users. We randomly split the dataset per topic into 10% test and 90% training. 10% of the latter are used as the validation set. Statistics are given in Table 1.

To develop approaches to the belief-based claim generation task, we need training data where claims can be identified as such. Since claim detection is not our focus, we preprocess all data using the claim detection approach of Chakrabarty et al. (2019). In particular, we score the likelihood of each sentence being a claim, and only keep the one with the highest score as the user’s claim on the topic. To evaluate the model, we created a sample of 100 arguments, and two annotators decided whether the extracted sentence represents a claim on the given topic or not. In terms of full agreement, the model extracted claims correctly in 81% of the cases, the Cohen’s  $\kappa$  inter-annotator agreement being 0.3. We note that this preprocessing step produces some noise in the data, mainly affecting the training of our Seq2Seq model below.

## 4 Approach

To study our research question, we propose and compare two approaches that build on top of known techniques for text generation. Both approaches rely on modeling users’ beliefs via their stances on big issues. The first is an extension of the Seq2Seq model (Sutskever et al., 2014), where the user’s stances are encoded as a context vector, while the second conditions the output of a pre-trained argumentative language model via a bag-of-words, constructed based on stances on big issues.

## 4.1 Seq2Seq-based Model

Given a topic, as a sequences of words  $T = (w_1, w_2, \dots, w_n)$ , a user vector  $\vec{U} \in \{0, 1\}^k$  with  $k$  being the number of big issues, and a claim as a sequence of words  $C = (w_1, w_2, \dots, w_m)$ , first an LSTM-based encoder consumes the input topic and produces a hidden state  $\vec{h}$ , which is used to initialize the LSTM-based decoder. The user vector  $\vec{U}$  is projected into a new embedding space via a feed forward network with a learned weight matrix  $W_U$ , producing a new vector,  $\vec{V}$ :

$$\vec{V} = \sigma(W_U \cdot \vec{U})$$

Following Li et al. (2016), our  $\vec{V}$  is served as their speaker embedding in the model. The difference between the speaker model in Li et al. (2016) and this model is that the vector  $\vec{V}$  is not explicitly predefined but rather learned from the data, while in our model it is already predefined as a binary vector representing the user’s stances on big issues. By augmenting the Seq2Seq model with a context user vector, the model is supposed to capture the correlation between users’ stances on big issue and the corresponding claims. Once the correlation is learned, the model can generate a claim utilizing not only the topic, but also the stances on big issues of the target user, which reflect the beliefs.

## 4.2 Conditioned Language Model

In this approach, we represent a user’s stances on big issues as a bag-of-words. We then use the topic as a prompt for a pre-trained argumentative language model (LM) to synthesize a claim conditioned using the algorithm of Dathathri et al. (2020). The synthesis process is illustrated in Figure 1.

**Argumentative Language Model** Since we aim to generate *claims* in particular, a standard LM is not enough. To model argumentative language, we take a LM pre-trained on general language and fine-tune it on a large set of arguments (in our experiments, we use the corpus of Ajjour et al. (2019)). The result is an LM that is able to generate argumentative text.

**Belief-based Bag-of-words** Next, we build a bag-of-words that represents the beliefs of a user. We learn this from the user’s stances on the big issues. For example, a user pro *abortion* would likely be pro *choice*. Hence, words such as *right* and *choice* are candidates to be included in their

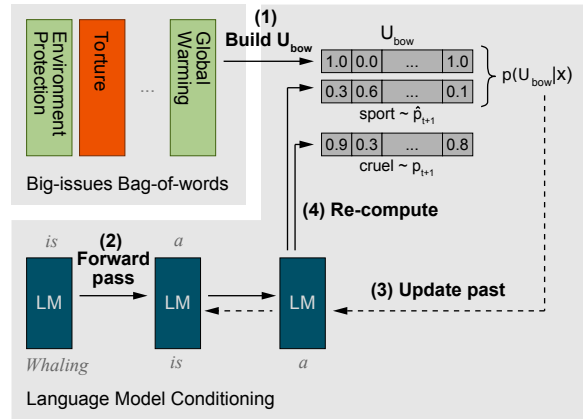


Figure 1: The synthesis process of the conditioned LM on the topic “Whaling”, given a user who is pro *environmental protection* and *global warming* and con *torture*. Steps: (1) Building  $U_{bow}$ , based on stances (2) Forward pass through the LM to generate a token, *sport* (3) Updating the LM history  $H_t$ , based on  $p(U_{bow}|x)$ , and (4) Generating from the new history  $\hat{H}_t$  a new token *cruel*.

belief-based bag-of-words. To this end, we first build two bag-of-words representations for each big issue, one for the pro and for the con side. For a user, we then construct a belief-based bag-of-words based on their stances on big issues.

To build a representative pro and con bag-of-words for each big issue, we follow the topic signature approach of Lin and Hovy (2000). Given a big issue, we first collect from some corpus of arguments three sets: relevant pro arguments  $R_{pro}$ , relevant con arguments  $R_{con}$ , and a random set of non-relevant arguments  $\hat{R}$ . For each relevant set, we then compute a likelihood ratio for all its words with respect to  $\hat{R}$  and keep only words with a score higher than a specific threshold  $\tau$ , resulting in two sets of words,  $W_{pro}$  and  $W_{con}$ . Since a word may appear in both sets, we remove it from the set where it occurs fewer times. Finally, we sort words according to their likelihood ratio and keep in both  $W_{pro}$  and  $W_{con}$  the top  $k$  words, forming the final pro and con bag-of-words respectively.

**Claim Generation** Given a user (represented by stances on big issues) and a topic, we construct a belief-based bag-of-words (Step 1 in Figure 1):

$$U_{bow} = W_1 \cup W_2 \cup \dots \cup W_n$$

where  $W_i$  is the pro bag-of-words if the stance is pro and the con bag-of-words otherwise. Then, we use the topic as a prompt and the user’s bag-of-words  $U_{bow}$  to condition the generated claim (see Figure 1). In particular, given a transformer-

based LM (Vaswani et al., 2017), a token  $x_{t+1}$  is generated at each time step as follows:

$$o_{t+1}, H_{t+1} = LM(x_t, H_t)$$

$$x_{t+1} \sim p_{t+1} = \text{Softmax}(W \cdot o_{t+1})$$

where  $H_t$  represents the history of the LM. Using the algorithm of Radford et al. (2019), called Plug and Play LM (PPLM), an update to the past,  $\Delta H$ , is computed to control the generated claim, based on the sum of the log likelihood  $p(U_{bow}|x)$  of all words in the belief-based bag-of-words. Then the new history,  $\hat{H}_t = H_t + \Delta H_t$ , is used as in the previous equations to draw a new distribution  $\hat{p}_{t+1}$ , of which a new token is sampled. To ensure fluency in the generated text,  $\Delta H$  is further modified to ensure a high log-likelihood  $p(x)$  with respect to the LM. More details on the algorithm can be found in the work of Radford et al. (2019).

In short, through fine-tuning an LM on argumentative text, we tune it to generate claims. Using the topic as a prompt, we ensure that the claim is on the topic. Finally, the PPLM represents beliefs, modeled as a bag-of-words  $U_{bow}$ , in the claim.

## 5 Automatic Evaluation

In this section, we evaluate whether utilizing user’s beliefs as input, modeled as stances on big issues, leads to claims that better match the ground-truth claims and reveal the input stances on big issues.

### 5.1 Experimental Setup

On one hand, we compute the BLEU and METEOR scores of the generated claims with respect to the ground-truth claims. On the other hand, we compute the likelihood that the generated claims possess textual features that reflect the input user’s beliefs. We do so by measuring the accuracy of predicting user’s stances on big issues given the generated claims. We compute this accuracy for each of the 48 big issues individually and report the results for all of them. To this end, we carry out the following three steps for a given approach.

First, we generate claims for all given users and topics in the test dataset. Second, we keep only instances in which users have a stance (pro/con) on the tested big issue, and split the filtered dataset into training and test. Finally, we train a simple TF-IDF based linear classifier on the training set to predict the stance on the big issue given the text of the claim. The accuracy of the classifier on the test

Approach	BLEU-1	BLEU-3	METEOR
S2S-baseline	18.2%	0.44%	<b>16%</b>
S2S-model	* <b>18.4%</b>	* <b>0.46%</b>	<b>16%</b>
LM-baseline	9.6%	<b>0.26%</b>	8%
LM-conditioned	* <b>12.0%</b>	0.16%	* <b>11%</b>

Table 2: BLEU and METEOR scores of the claims of each evaluated approach compared to the ground-truth claims. Values marked with \* are significantly better than the respective baseline at  $p < .05$  (student’s  $t$ -test).

split then quantifies the likelihood of the generated claims possessing textual features that reflect the stance on the corresponding big issue.

### 5.2 Implementation Details

In the following, we give implementation details of our approaches and the corresponding baselines:

**Seq2seq-based Model** Based on the OpenNMT framework (Klein et al., 2017), the encoder and decoder are each two-layer LSTMs of hidden size 512 with GloVe word embeddings of size 300. Users’ stances on big issues are represented as a one-hot encoded vector, and then projected into 16 dimensions space through a one-layer dense neural network. We trained the model with the Adagrad optimizer (batch size 16) and refer to it as *S2S-model*.

**Conditioned Language Model** We constructed the pro/con relevant argument sets ( $R_{pro}, R_{con}$ ) by querying the respective big issue from the API provided by Ajour et al. (2019) and extracting pro/con arguments from the top 60 results. For the non-relevant argument set ( $\hat{R}$ ), we used the same corpus (Ajour et al., 2019) and randomly selected 100 arguments. We eliminated all words with a score under  $\tau = 10$  and finally kept the top  $k = 25$  words from each set ( $R_{pro}, R_{con}$ ) to represent the bag-of-words.<sup>2</sup>

To model the argumentative language, we fine-tuned the GPT-2 model on the corpus of Ajour et al. (2019), which contains around 400k arguments. The fine-tuning was performed using the transformers framework (Wolf et al., 2019). We used the topic as a prompt to trigger the generation process. However, since some topics are phrased as a question (e.g, “is abortion wrong?”), we extracted the noun phrase from the topic and used it as a prompt. For conditioning the generated claim,

<sup>2</sup>We refrained from tuning the parameters here since we do not have a ground truth.

Approach	Abortion	Death penalty	Gay Marriage	Drug legaliz.	Global warming	Environm. protection	Medical mariju.	Smok. ban	Minim. wage	Border fence	All 48 big issues
Ground-truth	0.49	0.59	0.55	0.55	0.55	0.55	0.50	0.53	0.48	0.62	0.52
S2S-baseline	0.49	0.48	<b>0.52</b>	0.45	0.51	0.51	0.57	<b>0.53</b>	0.53	0.46	0.50
S2S-model	0.55	<b>0.55</b>	0.45	0.45	0.51	<b>0.58</b>	0.57	<b>0.53</b>	0.49	<b>0.52</b>	0.51
LM-baseline	0.48	0.50	0.54	0.49	0.54	0.56	0.51	0.45	0.59	0.46	0.50
LM-conditioned	<b>*0.58</b>	<b>*0.53</b>	0.45	<b>0.56</b>	<b>*0.61</b>	<b>0.58</b>	<b>0.58</b>	<b>0.53</b>	<b>0.65</b>	0.50	<b>0.54</b>
# Training	1 610	1 532	2 098	1 538	1 960	2 196	2 096	1 370	1 580	1 092	-
# Test	350	366	196	316	156	86	138	294	172	280	-

Table 3: Accuracy of each classifier trained on claims generated by the evaluated approaches to predict the stance, on the 10 most frequent big issues as well as on average over all 48 big issues. Values marked with \* are significantly better than corresponding baseline at  $p < 0.05$  according to a one-tailed Student’s  $t$ -test.

we used the PPLM implementation (Dathathri et al., 2020)<sup>3</sup>. We call this model the *LM-conditioned*.

**Baselines** To evaluate the gain of encoding user’s beliefs, we compare our two approaches to the corresponding version without stances on big issues as an input. We refer to these baselines as *S2S-baseline* and *LM-baseline* respectively<sup>4</sup>.

### 5.3 Results

Table 2 shows the results of our approaches and the baselines in terms of BLEU and METEOR.

For *S2S*, the BLEU scores of our approach are significantly better than the baseline. The *LM-conditioned* is significantly better than the baseline version in terms of BLEU-1 and METEOR. In general, the *S2S-model* has the highest scores across all measures. The reason may be that it was trained in a supervised manner on the given dataset, whereas the *LM-model* was only fine-tuned in an unsupervised way on a different argument corpus.

Regarding the encoding of user stances, Table 3 shows the accuracy of a linear classifier trained to predict the stance from the claims generated by each approach as well as from the ground-truth, on average and on the 10 most frequent big issues. A complete table with all big issues can be found in the appendix.

The best average accuracy across all the big issues is achieved by the *LM-model* (0.54). Compared to the corresponding baselines, the *LM-model* and the *S2S-model* generated claims that boosted the accuracy of the stance classifier on 33 (69%) and 21 (44%) of all big issues respec-

tively. Overall, in 20 of the big issues, the best accuracy was achieved on the claims generated by the conditioned LM, compared to only nine big issues for the *S2S-model*. This indicates that the *LM-conditioned* can better encode a user’s beliefs, modeled as stances on big issues, into generated claims.

## 6 Manual Evaluation

To obtain more insights into belief-based claim generation, we let users manually evaluate the output of the given approaches. Upon inspecting a sample of generated claims by our approaches, we noticed that the *LM-conditioned* produces more fluent and informative texts. Accordingly, we focused on the *LM-conditioned* and its baseline in the evaluation, where we conducted two user studies. The goal of the first was to assess the quality of the big-issue bag-of-words collected automatically, while the second targeted the output of the *LM-model*, its baseline, and a variant that utilizes a manually refined bag-of-words.

### 6.1 Automatic Collection of Bag-of-words

To keep the manual annotation effort manageable, we evaluated only the top-10 big issues. Two authors of this paper categorized each word in the pro/con bag-of-words of the corresponding big issue into five categories, *c1*–*c5*:

- c1: Word irrelevant to the big issue.
- c2: Relevant word, wrong stance.
- c3: Relevant word, both stances possible.
- c4: Relevant word, correct stance.
- c5: Very relevant word, correct stance.

<sup>3</sup>step-size=0.15 and the repetition-penalty=1.2

<sup>4</sup>A baseline that uses the corresponding bag-of-words of the targeted topic to guide the generation wouldn’t be valid, since we don’t have information on the user’s stance on this targeted topic.

Approach	Overall			Relatedness Level 4			Relatedness Level 3			Relatedness Level-2		
	True	False	Undec.	True	False	Undec.	True	False	Undec.	True	False	Undec.
LM-baseline	44%	34%	<b>22%</b>	50%	50%	<b>0%</b>	55%	31%	14%	<b>27%</b>	20%	<b>53%</b>
LM-conditioned	37%	32%	31%	35%	38%	27%	59%	41%	<b>0%</b>	13%	<b>13%</b>	74%
LM-cond. (manual)	<b>45%</b>	<b>26%</b>	28%	<b>50%</b>	<b>31%</b>	19%	<b>61%</b>	<b>28%</b>	11%	25%	18%	56%
Ground Truth	42%	30%	28%	38%	42%	19%	64%	27%	9%	27%	19%	54%

Table 4: Manual Evaluation: Percentage of cases for each approach where the majority of annotators predicted the stance of a generated claim on the given big issue correctly (true), incorrectly (false), or could not decide it (Undec.). The overall scores and those for each topic/big-issue relation level are listed.

Words	Irrelevant		Relevant		Very Relevant
	c1	c2	c3	c4	c5
Pro	14%	10%	36%	34%	6%
Con	36%	2%	34%	26%	2%

Table 5: Distribution of the pro/con bag-of-words, averaged across the top-10 big issues, over the five considered categories: c2 means wrong stance, c3 words that fit both stances, and c3 and c4 represent correct stance.

Examples can be found in the appendix. To compute inter-annotator agreement, three big issues were annotated by both annotators, resulting in Cohen’s  $\kappa$  of 0.45, reflecting moderate agreement. Afterwards, only one annotator continued the annotations for the other big issues.

Table 5 shows the distribution of words over categories, averaged across the 10 big issues. For the pro bag-of-words, around 40% of the words are relevant and reflect the right stance, while 36% are relevant but could be used in arguments from both stances. For the con bag-of-words, however, the percentages are lower (28% and 34% respectively). A considerable proportion of words belong to categories c1 and c2, which creates noise that could confuse the conditioning process of the LM. Hence, we also consider a variant of the conditioned LM that uses only relevant words from c4 and c5.

## 6.2 Claim Generation

We evaluate the effectiveness in terms of whether a given generated claim reveals the stance of the given user on a specific big issue as well as how informative the claim is regarding the given topic.

Since not all topics are directly related to the big issues that can be revealed in the generated claims, we manually annotated the relatedness of the top frequent 200 topics in the test dataset to the most frequent 10 big issues, and created the evaluation sample accordingly. In particular, two authors of this paper scored the relatedness of each pair of

topics and big issues on a scale from 1 to 4:

- 4: Topic and big issue are the same. Example: *"gay marriage should be legalized"* and *"gay marriage"*
- 3: A stance on the topic likely affects the stance on the big issue. Example: *"killing domestic abusers"* and *"death penalty"*
- 2: A stance on the topic may affect the stance on the big issue. Example: *"morality"* and *"abortion"*
- 1: Topic and big issue are not related. Example: *"do aliens exist?"* and *"abortion"*

The two annotators had a Cohen’s  $\kappa$  agreement of 0.54. Around 97.4% of all pairs got score 1, 1.1% score 2, 0.8% score 3, and 0.7% score 4. The small percentage of cases that can be evaluated reflects a limitation in the designed evaluation study. However, it still allows us to evaluate the effectiveness of our approach for different levels of relatedness. Given the annotated pairs, we randomly selected 10 pairs from levels 2, 3, and 4 each. For each pair, we then collected all claims on the topic from the test set, where the author specifies a stance on the corresponding big issue. We randomly select 30 claims each, resulting in an evaluation sample of 90 instances.

We used the crowdsourcing platform *MTurk*<sup>5</sup> for evaluation. For each instance, we showed a topic, a claim, and the corresponding big issue to three annotators. The annotators had to perform two tasks: (1) to predict the stance of the user on the corresponding big issue from the text of the claim, and (2) to rate the claim’s informativeness regarding the topic on a scale from 1 to 3.

Table 4 shows the percentage of cases in which the majority of annotators predicted the stance correctly (true), incorrectly (false), or could not decide

<sup>5</sup>A crowd sourcing platform: <https://www.mturk.com/>

Approach	Overall	Level 4	Level 3	Level 2
LM-baseline	1.8	2.5	1.9	1.4
LM-conditioned	2.1	2.3	2.5	1.5
LM-cond. (manual)	2.0	2.3	2.2	1.5
Ground Truth	2.0	1.9	1.8	2.2

Table 6: Manual Evaluation: Mean informativeness of the claims generated by each approach with regard to the topic (1–3, higher is better). The overall scores and those for each topic/big-issue relation level are listed.

about the stance (undec.) from the generated claim. Across the whole sample (Overall), the claims generated by *LM-conditioned (manual)*, the model conditioned on the refined bag-of-words, most often allowed to predict the stance correctly (45%). We thus attribute the low effectiveness of the *LM-model* to the noise generated by the automatic collection of big-issues’ bag-of-words, especially seeing that the effectiveness gets better across all levels when eliminating this noise.

Analyzing each relatedness level individually yields more insights. For relatedness level 4, where the topic is the same as the big issue, the *LM-conditioned (manual)* generated claims where the majority of the cases with known stance were correct (63%). In level 3, we observe that both versions of our approach outperform the baseline in producing claims that express the correct stance on the corresponding big issue with percentages of 59% and 68% respectively. Finally, at relation level 2, which represents a weak relation between topics and big issues, predicting the stance seems to become hard, as indicated by high percentages of undecided cases. We believe that the weak relatedness made the annotators guess the stance in some cases, leading to unreliable annotations.

Table 6 shows the average score of all approaches regarding the informativeness of the generated claims. Here, both versions of our approach achieved better scores than the baseline, matching the ground-truth score. We believe that the low scores of the ground-truth claims stem from the noise generated in the claim detection step.

**Error Analysis** Table 7 shows some cases from our evaluation. In case (a), even though the model generated a claim focusing on *fundamental rights*, the majority of the annotators was undecided about the stance. Case (b) shows the effect of noise in the bag-of-words: the generated claim reflects a pro stance on abortion, while the ground-truth stance is

(a) <i>Topic</i> : is abortion ok <i>Big issue</i> : (Con) Abortion			<i>Level</i> : 4
Approach	Claim	Stance	
LM-cond.	abortion rights groups argue that the right to abortion is a fundamental human right.	Undec.	
LM-baseline	abortionists are not the only ones who are against abortion. There are many other people who are against abortion	Undec.	
(b) <i>Topic</i> : abortion is right or wrong <i>Big issue</i> : (Con) Abortion			<i>Level</i> : 4
Approach	Claim	Stance	
LM-cond.	abortion right or wrong is a personal choice. I believe that abortion should be legal because it’s the woman’s body and she has every right to do what ever she wants.	Pro	
LM-cond. (manual)	abortion right or wrong. I believe that abortion is murder and should be illegal in all cases except when the mother’s life is at stake, rape, incest etc.	Con	
(c) <i>Topic</i> : whaling <i>Big issue</i> : (Pro) Environmental Protection			<i>Level</i> : 3
Approach	Claim	Stance	
LM-cond.	whaling is a very dangerous and cruel practice. It causes great harm to the environment, animals that are caught in it have their lives cut short by being eaten alive.	Pro	
LM-baseline	whaling is a sport. It is a sport because it involves a lot of physical exertion and skill.	Undec.	
(d) <i>Topic</i> : should murder be legalized? <i>Big issue</i> : (Pro) Death Penalty			<i>Level</i> : 2
Approach	Claim	Stance	
LM-cond.	should murder be legalized? I think that abortion should not be legal. Abortion is killing a baby, and it’s wrong to kill an innocent human being.	Undec.	
LM-baseline	should murder be legalized? I think so. I think that it should be legalized because it is a good thing	Undec.	

Table 7: A selection of claims generated by the different evaluated approaches for the different association levels between topic and big issue discussed in the text.

con. This is avoided in the claim generated by *LM-conditioned (manual)*. Case (c) shows a working example of which our approach correctly generated a claim on *whaling* from an environmental perspective when conditioned as such. Case (d) is a level 2 example, indicating limitation in our evaluation, namely, the generated claim reveals a stance on abortion, but we asked about death penalty.



## 7 Conclusion

In this paper, we have proposed to equip argument generation technology with the ability to encode beliefs for two reasons: first, it reflects the human process of synthesizing arguments, and second, it gives more control on the generated arguments leading to a better reach of the audience. For this purpose, we have presented the task of belief-based claim generation. Concretely, we studied the research questions of how to model a user’s beliefs as well as how to encode them when generating an argumentative text. We have modeled users’ beliefs via their stances on big issues, and used them as an extra input in our approaches.

Our automatic evaluation has provided evidence of the applicability of encoding beliefs into argumentative texts. In manual studies, we found that limitations in the effectiveness of our approach stem from noise produced by the automatic collection of a bag-of-words. The findings of this paper lay the ground to investigate the role of beliefs in generating arguments that reach their audience.

We point out that ethical issues arise, when tuning arguments to affect specific people, such as attempts to manipulate them. While the task and settings considered here are rather too fundamental to already make these issues critical, future work should pay attention to them. Our goal is to develop systems that bring people together.

## Acknowledgments

We thank the anonymous reviewers for their helpful feedback. This work was partially supported by the German Research Foundation (DFG) within the Collaborative Research Center “On-The-Fly Computing” (SFB 901/3) under the project number 160364472.

## References

Yamen Ajjour, Henning Wachsmuth, Johannes Kiesel, Martin Potthast, Matthias Hagen, and Benno Stein. 2019. Data acquisition for argument search: The args.me corpus. In *Proceedings of the 42nd Edition of the German Conference on Artificial Intelligence*, page 48?59.

Milad Alshomary, Shahbaz Syed, Martin Potthast, and Henning Wachsmuth. 2020. [Target inference in argument conclusion generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4334–4345, Online. Association for Computational Linguistics.

Yonatan Bilu, Daniel Hershcovich, and Noam Slonim. 2015. [Automatic claim negation: Why, how and when](#). In *Proceedings of the 2nd Workshop on Argumentation Mining*, pages 84–93, Denver, CO. Association for Computational Linguistics.

Giuseppe Carenini and Johanna D Moore. 2006. Generating and evaluating evaluative arguments. *Artificial Intelligence*, 170(11):925–952.

Tuhin Chakrabarty, Christopher Hidey, and Kathleen McKeown. 2019. Imho fine-tuning improves claim detection. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 558–563.

Wei-Fan Chen, Henning Wachsmuth, Khalid Al Khatib, and Benno Stein. 2018. [Learning to flip the bias of news headlines](#). In *Proceedings of the 11th International Conference on Natural Language Generation*, pages 79–88. Association for Computational Linguistics.

Sumanth Dathathri, Andrea Madotto, Janice Lan, Jane Hung, Eric Frank, Piero Molino, Jason Yosinski, and Rosanne Liu. 2020. [Plug and play language models: A simple approach to controlled text generation](#). In *International Conference on Learning Representations*.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.

Esin Durmus and Claire Cardie. 2018. [Exploring the role of prior beliefs for argument persuasion](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1035–1045, New Orleans, Louisiana. Association for Computational Linguistics.

Frans H van Eemeren and Peter Houtlosser. 1999. Strategic manoeuvring in argumentative discourse. *Discourse studies*, 1(4):479–497.

Roxanne El Baff, Henning Wachsmuth, Khalid Al Khatib, Manfred Stede, and Benno Stein. 2019. [Computational argumentation synthesis as a language modeling task](#). In *Proceedings of the 12th International Conference on Natural Language Generation*, pages 54–64, Tokyo, Japan. Association for Computational Linguistics.

Matthew Feinberg and Robb Willer. 2015. From gulf to bridge: When do moral arguments facilitate political influence? *Personality and Social Psychology Bulletin*, 41(12):1665–1681.

David M Godden. 2010. The importance of belief in argumentation: Belief, commitment and the effective resolution of a difference of opinion. *Synthese*, 172(3):397–414.

- Floriana Grasso, Alison Cawsey, and Ray Jones. 2000. Dialectical argumentation to solve conflicts in advice giving: a case study in the promotion of healthy nutrition. *International Journal of Human-Computer Studies*, 53(6):1077–1115.
- Christopher Hidey and Kathleen McKeown. 2019. Fixed that for you: Generating contrastive claims with semantic edits. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1756–1767.
- Xinyu Hua, Zhe Hu, and Lu Wang. 2019. Argument generation with retrieval, planning, and realization. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2661–2672.
- Nitish Shirish Keskar, Bryan McCann, Lav R Varshney, Caiming Xiong, and Richard Socher. 2019. Ctrl: A conditional transformer language model for controllable generation. *arXiv preprint arXiv:1909.05858*.
- Guillaume Klein, Yoon Kim, Yuntian Deng, Jean Senellart, and Alexander Rush. 2017. [OpenNMT: Open-source toolkit for neural machine translation](#). In *Proceedings of ACL 2017, System Demonstrations*, pages 67–72, Vancouver, Canada. Association for Computational Linguistics.
- Jiwei Li, Michel Galley, Chris Brockett, Georgios Spithourakis, Jianfeng Gao, and Bill Dolan. 2016. [A persona-based neural conversation model](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 994–1003, Berlin, Germany. Association for Computational Linguistics.
- Chin-Yew Lin and Eduard Hovy. 2000. [The automated acquisition of topic signatures for text summarization](#). In *COLING 2000 Volume 1: The 18th International Conference on Computational Linguistics*.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI Blog*, 1(8):9.
- Alan Ritter, Colin Cherry, and William B. Dolan. 2011. [Data-driven response generation in social media](#). In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*, pages 583–593, Edinburgh, Scotland, UK. Association for Computational Linguistics.
- Misa Sato, Kohsuke Yanai, Toshinori Miyoshi, Toshihiko Yanase, Makoto Iwayama, Qinghua Sun, and Yoshiki Niwa. 2015. End-to-end argument generation system in debating. In *Proceedings of ACL-IJCNLP 2015 System Demonstrations*, pages 109–114.
- Benjamin Schiller, Johannes Daxenberger, and Iryna Gurevych. 2020. Aspect-controlled neural argument generation. *arXiv preprint arXiv:2005.00084*.
- Irene Solaiman, Miles Brundage, Jack Clark, Amanda Askell, Ariel Herbert-Voss, Jeff Wu, Alec Radford, and Jasmine Wang. 2019. Release strategies and the social impacts of language models. *arXiv preprint arXiv:1908.09203*.
- M. Stede, J. Schneider, and G. Hirst. 2018. *Argumentation Mining*.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*, pages 3104–3112.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008.
- Henning Wachsmuth, Manfred Stede, Roxanne El Baff, Khalid Al Khatib, Maria Skeppstedt, and Benno Stein. 2018. [Argumentation synthesis following rhetorical strategies](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 3753–3765. Association for Computational Linguistics.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, et al. 2019. Transformers: State-of-the-art natural language processing. *arXiv preprint arXiv:1910.03771*.
- Daniel M Ziegler, Nisan Stiennon, Jeffrey Wu, Tom B Brown, Alec Radford, Dario Amodei, Paul Christiano, and Geoffrey Irving. 2019. Fine-tuning language models from human preferences. *arXiv preprint arXiv:1909.08593*.
- Ingrid Zukerman, Richard McConachy, and Sarah George. 2000. [Using argumentation strategies in automated argument generation](#). In *INLG'2000 Proceedings of the First International Conference on Natural Language Generation*, pages 55–62, Mitzpe Ramon, Israel. Association for Computational Linguistics.