From Argumentation to Deliberation: Perspectivized Stance Vectors for Fine-grained (Dis)agreement Analysis

Moritz Plenz[♣♡] **Philipp Heinisch**[♠♡] **Janosch Gehring**[♣]
Philipp Cimiano[♠] **Anette Frank**[♣]

[♣] Heidelberg University [♠] Bielefeld University
{plenz,gehring,frank}@cl.uni-heidelberg.de
{pheinisch,cimiano}@techfak.uni-bielefeld.de

Abstract

Debating over conflicting issues is a necessary first step towards resolving conflicts. However, intrinsic perspectives of an arguer are difficult to overcome by persuasive argumentation skills. Proceeding from a debate to a *deliberative process*, where we can identify actionable options for resolving a conflict requires a deeper analysis of arguments and the perspectives they are grounded in – as it is only from there that one can derive mutually agreeable resolution steps. In this work we develop a framework for a *deliberative analysis of arguments* in a computational argumentation setup. We conduct a fine-grained analysis of *perspectivized stances* expressed in the arguments of different arguers or stakeholders on a given issue, aiming not only to identify their opposing views, but also shared perspectives arising from their attitudes, values or needs. We formalize this analysis in *Perspectivized Stance Vectors* that characterize the individual perspectivized stances of all arguers on a given issue. We construct these vectors by determining *issue-* and *argument-specific concepts*, and predict an arguer's stance relative to each of them. The vectors allow us to measure a *modulated (dis)agreement* between arguers, structured by perspectives, which allows us to identify actionable points for conflict resolution, as a first step towards deliberation.[1]

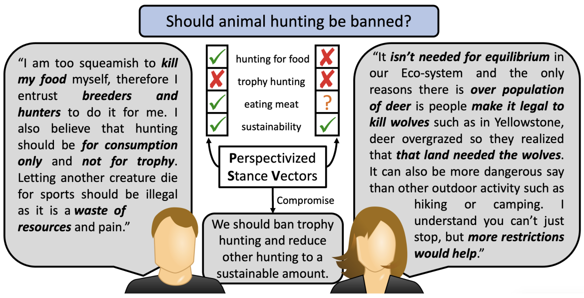

Figure 1: Example PSVs for '*Animal Hunting*'.

1 Introduction

Diverse stakeholders exchange their opinions and arguments on social media, news, debating portals and other private or public discussion formats. Often, they are in strong opposition, leaving little room for a consensus that could resolve the conflict. While argument mining technology has concentrated on analysing and generating arguments that can support arguers in winning a debate (Habernal and Gurevych, 2016; Wang et al., 2017; Wachsmuth et al., 2018), so far there has been limited interest in identifying points in opposing positions that bear a chance for consensual resolution of the conflict. Identifying points that offer a chance for resolution requires fine-grained analysis of the stances expressed by different stakeholders, to understand on which specific aspects they disagree and on which they actually might agree, and which of these are crucial for their mutual (dis)agreement.

This requires an analysis of the *perspectives* an arguer has on an issue – which may be grounded in their values, attitudes or specific goals and needs (Falk et al., 2024; Kiesel et al., 2022; Alshomary et al., 2022). We aim to compare arguments within a given issue based on their expressed perspectives, which means that we require a fixed set of perspectives for each issue. Issue-specific 'frames' are commonly used to group and analyze arguments from a given issue (Opitz et al., 2021; Heinisch et al., 2022), which makes them promising for modeling perspectives. Following Plenz et al. (2024), we use a data-driven approach to extract issue-specific frames from the commonsense knowledge graph ConceptNet (Speer et al., 2017), meaning that *concepts* (i.e., nodes from ConceptNet) form our basis for perspectives. To support a *deliberative analysis of arguments*, we develop tools to i) determine *relevant concepts* that characterize differ-

[♡]Authors contributed equally.
[1]Data and code are available at https://github.com/Heidelberg-NLP/PSV.

Findings of the Association for Computational Linguistics: NAACL 2025, pages 1525–1553
April 29 - May 4, 2025 ©2025 Association for Computational Linguistics

ent perspectives arguers may have on an issue and ii) what *stance* arguers express towards a certain perspective with a given argument. Our rationale is that by determining on which specific perspectives arguers *agree* or *disagree* in a debate, one may be able to identify points for achieving consensual agreement. In the following we thus use the terms *perspective* and *concept* interchangeably, although we want to note that perspectives may be formalized differently in future work.

Towards this goal, our work presents a new approach to construct a fine-grained representation of arguments that characterizes the *perspectivized stances* arguers express on a given issue, in so-called *Perspectivized Stance Vectors* (PSVs). A PSV is formalized as a vector of *stance values* towards issue-specific concepts (perspectives) – the so-called *signature*. Comparing multiple PSVs can reflect *opposing*, *agreeing* and *orthogonal* stances of *different strengths* for *different perspectives*, offering ways to identify potential anchors for deliberation processes. This goes beyond conventional stance classification, which only allows to identify conflicts at a binary level – instead, our analysis allows for more fine-grained assessments.

Figure 1 illustrates the concept of PSVs applied to two arguments on the issue "*Should animal hunting be banned?*" Choosing four example key perspectives from the debate, the arguer on the left is clearly PRO *hunting for food*, *sustainability* and *eating meat*, but AGAINST *trophy hunting*. The position of the arguer can thus be represented as a 4-dimensional vector, where the dimensions correspond to the above-mentioned perspectives. The vector for the second arguer will instead be: PRO: *sustainability*; AGAINST: *hunting for food*, *trophy hunting*; and NEUTRAL: *eating meat*.

The vectors show opposing stance on *hunting for food*, but agreement for *trophy hunting* (AGAINST) and *sustainability* (PRO). *Eating meat* does not show agreement nor disagreement, and hence is considered orthogonal. Agreeing dimensions could offer an entry point and basis for resolving the conflict by emphasizing shared positions, and aiming to find consensual solutions on points of opposite stances. Hence, "*We should ban trophy hunting and reduce other hunting to a sustainable amount*" could be a viable resolution.

In summary, our contributions are:

i) We formalize **P**erspectivized **S**tance **V**ectors (PSVs) as a structured representation of arguments to enable a *deliberative analysis*.
ii) This includes three subtasks: To construct PSVs we need to i) select issue-specific *signature concepts* and ii) classify the corresponding perspectivized *stances values* for a given argument. To identify (dis)agreement we need to iii) *aggregate PSVs*.
iii) We run experiments on deliberative issues from PAKT (a structured argumentation corpus; Plenz et al., 2024) and compare to a manually annotated evaluation set, showing our individual modules' performances.
iv) Our evaluations include case studies as proof-of-concept on how a perspectivized analysis of acceptability can support deliberation.

2 Related Work

Overall stance. PSVs are a fine-grained representation of stances. Classifying the overall stance of arguments towards a topic is a core task in argument mining (Bar-Haim et al., 2017; Kobbe et al., 2020; Luo et al., 2020). However, assigning an argument a global stance tag (e.g., PRO, CON and NEUTRAL or UNRELATED) lacks expressivity: it divides sets of arguments into only a couple of groups, neglecting crucial nuances. The task of same-side classification (predict whether two arguments share their overall stance) in Hou and Jochim (2017); Körner et al. (2021) does not address this problem either. Further, it does not unveil the underlying reasons why arguments share a stance.

To counter this issue, prior work incorporated background knowledge, by including *reasoning paths* to *explain*, e.g., for which reasons a premise supports or attacks a conclusion (Paul et al., 2020), or to generate an *explanation graph* for a premise-conclusion pair that explains the stance of the argument (Saha et al., 2021; Saadat-Yazdi et al., 2023; Plenz et al., 2023b). We build on this work by including concepts from a commonsense resource to define the PSV signature concepts.

Perspectives in argumentation. Our work is related to Barrow et al. (2021), who rely on graphs to represent arguments and their relationships as a basis to detect viewpoints. They proposed so-called *syntopical graphs* that model pairwise textual relationships between claims to enable a better reconstruction of latent viewpoints in a collection, thereby making points of (dis)agreement within the collection explicit. In a similar way, PSVs enable the detection of (dis)agreement. But in addition, PSVs can detect *orthogonality*, i.e., cases where a pair of arguments is not related to each other.

Our work is also related to the analysis of framing in argumentation (Heinisch et al., 2022, 2023; Otmakhova et al., 2024), where emphasized aspects are automatically detected. Specifically, our work is related to the idea that frames can be issue- or topic-specific and thus need to be identified in a bottom-up fashion for each topic. Ajjour et al. (2019) present an unsupervised approach that induces frames by clustering arguments from an issue. Ruckdeschel and Wiedemann (2022), by contrast, present a topic-specific framing approach, with the limitation of training a classifier for each topic separately – which then cannot be applied to new topics. Our unsupervised approach to signature induction follows Plenz et al. (2024), who find that the knowledge base ConceptNet (Speer et al., 2017) provides suitable, often implicit concepts relevant for argument interpretation.

Finally, our task is related to aspect-based sentiment analysis (ABSA; Cabello and Akujuobi, 2024; Wang et al., 2024; Brauwers and Frasincar, 2022; Hoang et al., 2019), which aims for a fine-grained view on which aspects are target of a certain sentiment. While ABSA is typically applied to reviews, we aim for a fine-grained, perspectivized analysis of arguments in deliberation, by detecting argument-related stances towards specific concepts.

Our framework supporting a deliberative analysis of arguments thus brings together and combines methods from viewpoint detection, framing and aspect-based sentiment analysis. We combine these methods in a novel way for deliberation support, by pinpointing conflicting perspectives and concepts between argument and stakeholders, with the aim of resolving conflicts and suggesting compromises.

Deliberation refers to a collaborative argumentative exchange where arguers hold incompatible views on an issue, which they seek to resolve by achieving a consensual decision (Felton et al., 2022). Deliberative processes are naturally framed as a dialogue (Walton et al., 2010; Snaith et al., 2010; Walton et al., 2016). For example, Al-Khatib et al. (2018) successfully classify different *strategies* of participants in deliberative discussions. Yet they do not evaluate the underlying perspectives, nor the effectiveness of these strategies for the aim of achieving an agreeable resolution of conflicts.

Deliberation is typically approached using preference frameworks that take into account the arguers' diverging desires or goals, or their normative or moral considerations (Modgil and Luck, 2009; Amgoud and Vesic, 2014; Bao et al., 2017). We do not focus on algorithmic resolution of conflicts, but on analyzing the arguers' perspectivized viewpoints to quantify dimensions of (dis)agreement – which future work may extend with reasoning processes, to derive potential resolutions.

Bergmann et al. (2018) provide overviews of debates to make decision makers aware of arguments and opinions on relevant topics. Using a Case-Based Reasoning approach, they compute similarity between arguments to retrieve or cluster similar arguments. This allows them to synthesize new arguments – by extrapolating from and combining existing arguments. While they focus on grouping similar arguments, we aim for an aggregated representation of debates in terms of perspectivized stances that reflect diverging and unified viewpoints of relevant stakeholder groups.

Some recent work leverages LLMs to model deliberative processes. E.g., Bakker et al. (2022) investigate if LMs can support humans in finding agreement on conflicting issues. They task LLMs to expand a corpus from a set of human-elicited questions and opinions on moral and political issues, and train a reward model to rate LM-generated consensus statements. They report high performance of LLMs generating consensual statements. However the evaluations do not report detailed statistics, and since the data (worth £46,000) is not made public, it remains unclear if the evaluation involves notable conflicts to start with.[2]

Our work is of smaller scale, but relies on arguments from a curated and accessible debate portal. In contrast to their work – which is elusive on *which* divisive arguments a consensus statement is meant to resolve – we explicitly represent arguments as stance vectors along conceptual perspectives, from which we compute highly interpretable acceptability scores as a basis for finding consensual solutions for conflicting arguments. Since our method is interpretable, this can increase trust, and thereby usability, compared to purely generative methods.

[2]For evaluation, opinions are clustered by topic (*not* by the original issues). Fig. 2B of the work splits the data into divisive and non-divisive questions within a group, where only 50% were found to be divisive. The win rates for their model over baselines are similar between (non)divisive questions, and the analysis does not detail agreement score differences between positioned vs. agreement statements for the divisive subset. Without access to the data nor detailed analyses, it is unclear whether the re-clustered data involves notable conflicts.

3 Perspectivized Stance Vectors

We introduce **P**erspectivized **S**tance **V**ectors (PSVs), a new representation to record the perspectives expressed in or underlying an argument, with the aim to detect and measure agreement and conflicts between pairs or a set of arguments on a given issue. To construct a PSV, we need to define its *signature* and corresponding *stance values*. Given a debate topic dt, the signature is determined by a list of concepts $\{c_i\}_{i=1}^n$ relevant for topic dt.

Given an argument a on topic dt, we abstract a PSV $\vec{v_a}$ from argument a by determining a stance s_i for each concept c_i, where s_i represents the stance the arguer expresses with argument a towards the concept c_i. We choose stance values $s_i \in [-1, 0, 1]$ to represent a stance of *against*, *neutral*, or *in favor*, respectively. We formalize PSVs as either n-dimensional vectors of stance values s_i, or $(n \times 3)$ dimensional matrices where each entry represents the probability of concept c_i to belong to one of the three stance value classes:

$$\begin{aligned} \vec{s} &\in \mathbb{S}^n \quad && \text{where} \quad \mathbb{S} = \{-1, 0, 1\} \\ \vec{p} &\in \mathbb{P}^{n \times 3} \quad && \text{where} \quad \mathbb{P} = [0, 1]. \end{aligned} \tag{1}$$

If the exact representation is not relevant, we use $\vec{v}$ to denote a general PSV. When comparing pairs of PSVs $(\vec{v}_{a_1}, \vec{v}_{a_2})$ for arguments (a_1, a_2), aligned vs. opposing dimensions indicate agreement or disagreement, respectively. Dimensions with neutral stance labels in $(\vec{v}_{a_1}, \vec{v}_{a_2})$ indicate *orthogonality*, as the arguers neither agree nor disagree.

We next describe our methods to construct PSVs, including their signatures and stance values in §3.1. We then describe how to aggregate or compare PSVs to obtain predictions for agreement, orthogonality and disagreement between arguments and which specific perspectives cause them (§3.2).

3.1 PSV Construction

Below we show how to construct PSVs given a topic dt and a set of arguments $\mathcal{A}_{dt}$ on that topic.

3.1.1 *Signature concepts* for Debate Topic

As a signature for PSVs we are interested in general – and potentially conflicting – concepts that capture the perspectives of diverse arguers towards a topic.

Following Plenz et al. (2023b, 2024), we align arguments to the commonsense knowledge graph ConceptNet (Speer et al., 2017). First, we split arguments into individual sentences, then we select for each sentence the most similar concepts (i.e., nodes in ConceptNet). We connect these concepts with weighted shortest paths that maximize semantic similarity to the argument. Concepts along such paths have been shown to cover relevant aspects of the given text, while maintaining high precision (Plenz et al., 2023b; Fu and Frank, 2024). These nodes form a set of concepts C_a that reflects the given argument $a \in \mathcal{A}_{dt}$.

To obtain conflicting concepts, we split the arguments $\mathcal{A}_{dt}$ by their overall stance towards the topic into $\mathcal{A}_{dt}^+$ and $\mathcal{A}_{dt}^-$. For each concept $c_i \in \cup_{a \in \mathcal{A}_{dt}} C_a$ and debate topic stance $s_{dt} \in \{+, -\}$ we compute the stance-specific frequency

$$f_{c_i}^{s_{dt}} = \frac{\#\{a \,|\, a \in \mathcal{A}_{dt}^{s_{dt}} \text{ and } c_i \in C_a\}}{\max\left(1, \#\{a \,|\, a \in \mathcal{A}_{dt}^{s_{dt}}\}\right)}. \tag{2}$$

We normalize the frequencies $f_{c_i}^{s_{dt}}$ for a given concept and stance by subtracting the frequencies of the concept with the opposite stance: $f_{c_i}^+ - f_{c_i}^-$ and $f_{c_i}^- - f_{c_i}^+$ for PRO and CON stance, respectively. To avoid redundancy, we remove concepts with duplicate lemmas. Finally, we take the top-k PRO and CON concepts, resulting in $2 \cdot k$ concepts in total.

Optionally, the resulting concepts can be filtered to obtain smaller and more concise PSVs: we either remove hypernyms of other concepts in order to further *reduce redundancy*, or remove concepts that ChatGPT judges as irrelevant for the topic, to *increase relevancy*. Refer to §A.1 for more details.

Ideally, the signature concepts are *relevant* to the debate topic and of appropriate *granularity*: If a signature concept is too fine-grained, then only very few (or no) arguments will evoke it. Similarly, if a signature concept is too coarse-grained then arguments may not have a clear stance towards it – consider for example the concept "hunting" for the first argument in Fig. 1. Here, "hunting for food" and "trophy hunting" would form a better signature. Hence, we aim for signature concepts with an intermediate granularity. In our experiments (§4.2) we asses the (i) *relevance* to the debate topic and (ii) *granularity* of selected signature concepts.

3.1.2 *Perspectivized stances* for Arguments

We develop different models to compute the stance value s_i for a given argument and signature concept c_i. Here we provide an overview of the methods, please refer to §A.2 for more in-depth descriptions.

Baseline. Our Baseline assigns each argument a and concept c_i the stance value 0 if the concept is *not* in the concept graph, i.e., $s_i = 0$ for $c_i \notin C_a$.

Method		Agreement [+]	Orthogonal [$\varnothing$]	Disagreement [$-$]
$\mathcal{S}$	Stance Value	$\delta\big(s_i^1, s_i^2\big)$	–	$1 - \delta\big(s_i^1, s_i^2\big)$
$\mathcal{S}_0$	Stance Value (Consid. Neut.)	$\delta\big(s_i^1, s_i^2\big)\,\big(1 - \delta\big(s_i^1, 0\big)\big)$	$\min\Big(\sum_{j=1}^{2} \delta\Big(s_i^j, 0\Big), 1\Big)$	$\big(1 - \delta\big(s_i^1, s_i^2\big)\big)\,\big(1 - \delta\big(s_i^1, 0\big)\big)\,\big(1 - \delta\big(s_i^2, 0\big)\big)$
$\mathcal{S}_D$	Stance Value (Difference)	$\mathcal{S}_0^+\big(s_i^1, s_i^2\big) - \mathcal{S}_0^-\big(s_i^1, s_i^2\big)$	–	$\mathcal{S}_0^-\big(s_i^1, s_i^2\big) - \mathcal{S}_0^+\big(s_i^1, s_i^2\big)$
$\mathcal{P}$	Stance Prob.	$\big(p_i^1 \odot p_i^2\big) \cdot [1,1,1]^T$	–	$1/2\,\big\lvert p_i^1 - p_i^2 \big\rvert \cdot [1,1,1]^T$
$\mathcal{P}_0$	Stance Prob. (Consid. Neut.)	$\big(p_i^1 \odot p_i^2\big) \cdot [1,0,1]^T$	$\big(p_i^1 \odot p_i^2\big) \cdot [0,1,0]^T$	$1/2\,\big\lvert p_i^1 - p_i^2 \big\rvert \cdot [1,0,1]^T$
$\mathcal{P}_D$	Stance Prob. (Difference)	$\mathcal{P}_0^+\big(p_i^1, p_i^2\big) - \mathcal{P}_0^-\big(p_i^1, p_i^2\big)$	–	$\mathcal{P}_0^-\big(p_i^1, p_i^2\big) - \mathcal{P}_0^+\big(p_i^1, p_i^2\big)$

Table 1: Aggregation Methods. $s_i^j \in \{-1, 0, 1\}$ is the stance value of argument j towards concept i and $p_i^j \in [0,1]^{(1\times 3)}$ are the corresponding probabilities. The Kronecker delta $\delta(x, y)$ is 1 if $x = y$ and 0 else. $\odot$ is element-wise multiplication. §A.3 discusses the formulas in more detail.

Else, for concepts that are in the argument graph C_a, the *debate topic stance* $s_{dt} \in \{+1, -1\}$ is assigned for concept c_i.

`RoBERTa`. We construct a synthetic dataset by automatically adapting sentiment (Sobhani et al., 2016; Mohammad et al., 2017) and human-value (Mirzakhmedova et al., 2024) detection datasets to our task. On this data we apply *transfer learning* using a RoBERTa model that has been fine-tuned for sentiment analysis (Loureiro et al., 2022), by further training it on our synthetic dataset. We emphasize that this synthetic data is exclusively used to finetune `RoBERTa`, which is not our best-performing system – and our remaining experiments are independent of this synthetic data. We refer to appendix A.2 for more details.

`GPT4o`. We also prompt `GPT4o` to predict whether an argument a is against, neutral or in favor of a concept c_i. We apply zero- and few-shot prompting, with two hand-crafted samples for the latter (A.2).

3.2 Computing *Acceptability* Scores

Standard stance classification allows us to predict whether two arguments agree or disagree on a debate topic. Using PSVs, we can now detect and predict agreement on a more fine-grained level. E.g., there is a *partial agreement* between the arguers in Fig. 1: Both parties are against *trophy* hunting while they are in favor of *sustainability*. Such *partial agreements* are instrumental to find compromises between arguers who *disagree* on a topic.

***Perspectivized Acceptability* Scores.** Our hypothesis is that i) arguers *agree* or *disagree* on the concepts c_i, depending on whether their arguments express the same perspectivized stances s_i towards c_i within the debated topic. Yet, ii) if at least one of two arguments has a neutral perspective towards concept c_i, then the arguers neither agree nor disagree on the concept, meaning they are *orthogonal*. Following this intuition we design several aggregation methods to detect perspectivized *agreement*, *orthogonality* and *disagreement* between arguments at concept level.

The aggregation functions are shown in Table 1: we group them depending on whether they use discrete stance values ($\mathbb{S} = \{-1, 0, 1\}$) or continuous probabilities for the respective classes ($\mathbb{P} = [0, 1]$). The functions return discrete or continuous predictions, respectively. Further, the aggregation can consider the special role of neutral stance values, as outlined above. Finally, for agreement and disagreement, we can consider the opposing class as an inhibiting factor. Hence, we also design functions that take the difference between agreement and disagreement. We refer to these functions as, e.g., $\mathcal{S}_0^-\big(s_i^1, s_i^2\big)$ to denote a disagreement score (superscript "$-$") between arguments a_1 and a_2 regarding concept i (function parameter) under a stance value-based aggregation function ($\mathcal{S}$) that considers the role of neutral stances (subscript 0).

For example, $\mathcal{S}^+\big(s_i^1, s_i^2\big) = 1$ iff the stance values are equal $s_i^1 = s_i^2$, while $\mathcal{S}_0^+\big(s_i^1, s_i^2\big) = 1$ and $\mathcal{S}_D^+\big(s_i^1, s_i^2\big) = 1$ iff the stance values are equal and non-zero $s_i^1 = s_i^2 \in \{-1, +1\}$. Compared to $\mathcal{S}_0^+\big(s_i^1, s_i^2\big), \mathcal{S}_D^+\big(s_i^1, s_i^2\big)$ also distinguishes between agreement scores of 0 and -1. §A.3 explains the functions in more detail.

Acceptability scores between pairs of arguments. By now we discussed how to calculate contributions of individual perspectives towards agreement, orthogonality or disagreement. To obtain an overall acceptability score for an argument pair, we average perspectivized stance values of all dimensions n of a PSV. While future work could investigate the effects of making the contributions of each perspective learnable, for the scope of this paper we restrict ourselves to unsupervised aggregation methods.

4 PSV-Experiments

Given a set of arguments on an issue, our approach first finds signature concepts, then computes perspectivized stances which yields PSVs and finally aggregates PSVs to obtain acceptability scores. In this section, we empirically assess the quality of each of these steps by comparing to human annotations. Where possible, we augment our manual evaluation with automatic evaluations that do not require human labels. Section 5 presents a complementary case study.

4.1 Experimental setup

Data. We conduct our analyses and case study using PAKT (Plenz et al., 2024), a debate resource that presents issues as binary questions, and answers to these questions as arguments for either stance. The arguments, on avg. 7 sentences long, discuss an author's points without elaborating on the entire issue. This makes PAKT well-suited for our purposes. Fig. 1 shows two shortened example arguments from PAKT. For our case study we further enrich PAKT with stakeholder groups (for details see §A.4).

Annotation. To assess the quality of our methods, three annotators labeled data from 5 different evenly represented topics: 300 topic-level annotations to evaluate PSV signatures, 500 argument-level annotations to evaluate PSV values and 1,500 annotations for pairs of arguments to evaluate our methods to predict acceptability scores on debate topics. We collect annotations from all three annotators for the topic *Should animal hunting be banned?* to estimate inter-annotator agreement. For most subtasks we achieve moderate to high agreement as shown by Tab. 5, despite the high subjectivity in argumentative tasks. §B.1 presents more details on the annotation procedure.

4.2 Analyzing PSV Construction Methods

First, we analyze how best to construct PSVs. This includes i) the selection of perspectives for the PSV signature and ii) how to predict PSV values.

PSV signature. Signature perspectives should be *relevant* to the topic. Also, if a perspective is too *general* (or too *specific*), it will be evoked by almost all (or no) arguments. Neither is useful for comparing arguments – hence, we check whether our signature concepts' granularity is in-between.

		all	-hyp.	-irrel.	-both
Avg	#	30.0	22.0	16.0	11.0
Relevance	P	90.0	90.9	**93.8**	90.9
	R	**100.0**	74.1	55.6	37.0
	F1	**94.7**	81.6	69.8	52.6
Granularity	P	53.3	54.5	**75.0**	72.7
	R	**100.0**	75.0	75.0	50.0
	F1	69.6	63.2	**75.0**	59.3

Table 2: Analysis of Signature Perspectives, evaluated against human annotation of the *Relevance* and appropriate *Granularity* of concepts. We present the unfiltered selection (*all*), as well as three filtering methods: *hyp* (hypernym) filtering, ChatGPT-based filtering for *irrel*evance, and a combination of *both*. *Avg* shows the effects on the avg. nb. of perspectives (PSV dimension). We measure *P*recision, *R*ecall and *F1*-score.

Tab. 2 presents the *relevance* and *granularity* scores evaluated against our human annotation.

Relevance. Unfiltered perspectives (*all*) are mostly relevant (Prec: 90.0%). This can be increased to 93.8% by filtering for irrelevance with ChatGPT – at the cost of discarding other relevant perspectives (Rec: 55.6%). The unfiltered approach yields the maximum F1-score of 94.7%.

Granularity. Without filtering (*all*) only 53.3% of perspectives have appropriate granularity. Filtering for irrelevance raises precision to 75.0%, while discarding some appropriate perspectives (Rec: 75.0%), and yields the overall highest granularity F1-score of 75.0%. This indicates that no filtering or irrelevance-filtering should yield best performances for granularity, depending on whether recall or precision is more important. Note, however, that hypernym filtering aims to reduce redundancies in the selected perspectives, which we do not assess in our human annotation. Hence, hypernym filtering might perform well in practice although our human annotation does not capture its advantages.

PSV stance values. Table 3 shows the performance of our PSV stance prediction methods compared to our annotations. `GPT4o` (zero-shot) is the best approach with 50.2% macro F1 across all perspectives. With few-shot prompting the performance reduces to 49.0%, but still outperforms our baseline and fine-tuned models. Overall the performances increase when considering only perspectives with appropriate granularity. These perspectives are less ambiguous and hence easier to annotate. Again, the `GPT4o`-based methods perform the best, achieving 56.3% and 56.4% macro F1 for zero- and few-shot, respectively.

For our remaining analyses we will use `GPT4o` (zero-shot) because of its overall best performance and simplicity. The confusion matrix (cf. Fig. 6 in §B.2) reveals that `GPT4o` (zero-shot) performs well for positive and negative perspectives, but often misclassifies neutral ones as negative.

Method	all	appropriate	negative	neutral	positive
Baseline	38.4	41.3	20.1	**71.0**	24.1
1-RoBERTa	36.5	37.9	45.5	28.5	35.4
3-RoBERTa	43.2	43.9	**46.6**	47.0	35.9
GPT4o (0-shot)	**50.2**	56.3	46.0	44.7	**59.9**
GPT4o (few-shot)	49.0	**56.4**	45.3	43.0	58.6

Table 3: PSV stance prediction. Scores are macro F1, evaluated on *all* perspectives, or on perspectives with *appropriate* granularity (c.f. §3.1.1). We also report F1 scores for individual classes across all annotated perspectives.

4.3 Evaluating Aggregation Methods

Unless stated otherwise, the reported results use 100 dimensional PSVs w/o filtering and `GPT4o` (0-shot) for perspectivized stance prediction.

4.3.1 Global acceptability: *Partial agreement* among argument pairs of *opposite stance*

In our manual annotation, argument pairs of opposite stances had been annotated for global acceptability, i.e., being in i) partial agreement, ii) agreement, iii) disagreement or for being iv) orthogonal to each other. Since only two argument pairs were annotated with "agreement" (which is expected, since all annotated argument pairs are of opposite stance), we group "agreement" and "partial agreement" to form one agreement class.

Tab. 4 (top) shows the global performance of our methods. We report ROC-AUC, which allows us to compare our discrete and continuous aggregation methods using the same metric. We observe that detecting agreement is the most difficult: the best aggregation method ($\mathcal{S}_D$) achieves 0.60 AUC. For orthogonality and disagreement $\mathcal{P}_0$ performs best with 0.69 and 0.70 AUC, respectively.

Fig. 8 (§B.3) shows scores for increasing lengths of PSVs, i.e., increasing number of perspectives the PSVs cover. For agreement and disagreement the performance is better for longer PSVs, while orthogonality is best with shorter PSVs. A closer look at the ROC curves (Fig. 7) reveals that shorter PSVs enable higher true positive rates for orthogonality. Fig. 8 also shows the impact of filtering

	Mode	Agreement	Orthogonal	Disagreement
Global	$\mathcal{S}$	0.59	–	0.67
	$\mathcal{S}_0$	0.55	0.66	0.68
	$\mathcal{S}_D$	**0.60**	–	0.64
	$\mathcal{P}$	0.57	–	0.67
	$\mathcal{P}_0$	0.54	**0.69**	**0.70**
	$\mathcal{P}_D$	0.58	–	0.62
Perspectivized	$\mathcal{S}$	0.54	–	0.63
	$\mathcal{S}_0$	0.58	0.64	0.70
	$\mathcal{S}_D$	0.57	–	0.67
	$\mathcal{P}$	0.52	–	0.73
	$\mathcal{P}_0$	**0.62**	**0.75**	**0.76**
	$\mathcal{P}_D$	0.56	–	0.72
	w/o PSV	**0.62**	0.66	0.69
Same Stance	$\mathcal{S}$	0.85	–	*0.85
	$\mathcal{S}_0$	0.79	*0.56	*0.84
	$\mathcal{S}_D$	0.86	–	*0.86
	$\mathcal{P}$	0.86	–	*0.86
	$\mathcal{P}_0$	0.80	0.55	***0.88**
	$\mathcal{P}_D$	0.86	–	*0.86

Table 4: Evaluation of aggregation methods in ROC-AUC. **Top**: evaluation on human annotation for *Global* acceptability scores (§4.3.1). **Middle**: evaluation on human annotation for *Perspectivized* acceptability scores (§4.3.2). **Bottom**: evaluation on *Same Stance* prediction (§4.3.3). For fields marked with * lower acceptability scores mean the arguments are from the same stance.

signature perspectives. Filtering by relevance shortens PSVs and improves agreement scores.

4.3.2 Perspectivized acceptability: Identifying aspects of (dis)agreement

In §4.3.1 we predicted global agreement between arguments, classifying argument pairs as a whole. However, using PSVs, we can also identify which perspectives an argument pair agrees or disagrees on. Again, we compare to our human annotation.

Tab. 4 (middle) shows the results. $\mathcal{P}_0$ consistently performs best for predicting agreeing, orthogonal and disagreeing perspectives, with 0.62, 0.75 and 0.76 AUC, respectively. Overall, the results tend to be better than global acceptability scores. This indicates that averaging over all n perspectives to obtain global scores incurs errors. A learned aggregation could alleviate this issue, but is beyond the scope of this work as it requires training data.

Ablation without PSVs. A valid concern regarding our approach is that our method reduces arguments to static vectors, which might oversimplify the nuances of deliberation. Further, it is of interest to what extent (dis)agreement scores could be predicted, in context, by a strong LLM. Thus, we also experimented with directly prompting `GPT4o` to predict the acceptability of two arguments on a specific perspective. This allows the model to di-

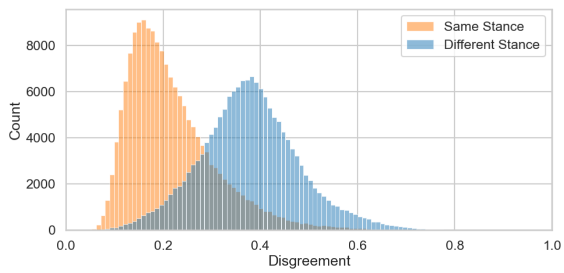

Figure 2: Disagreement ($\mathcal{P}_0$) distribution of argument pairs from the same stance or different stance.

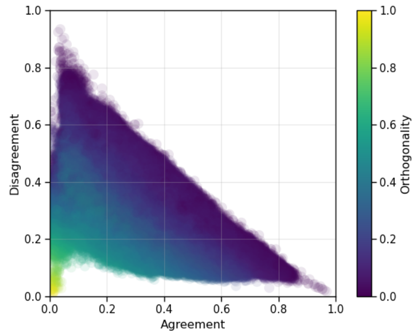

Figure 3: Acceptability scores computed with $\mathcal{P}_0$.

rectly compare the arguments, without the intermediate representation of PSVs. Note, however, that such an approach greatly diminishes *interpretability*, given the lack of a structured representation and *scalability*, as the number of comparisons scales quadratically with the number of arguments, as opposed to the linear scaling of our PSV framework.

Our prompts include a short task description, the two arguments to be compared and the list of perspectives to be evaluated. The complete prompt and further details are in §A.5. `GPT4o` obtained 0.62, 0.66 and 0.69 ROC-AUC for agreement, orthogonality and disagreement, respectively (c.f. Table 4). Surprisingly, this strong competitor falls behind our best-performing PSV aggregation $\mathcal{P}_0$. Prompting `GPT4o` with few-shot samples or chain-of-thought could potentially improve its results, but will by no means justify the loss of interpretability and scalability that is inherent to our PSV approach.

4.3.3 Evaluation with unannotated data

So far we evaluated our methods on our manually annotated data, which is naturally limited in size. To consolidate our analyses, we aim to verify our methods on larger amounts of unannotated data.

Same stance. To this end we perform same-side classification: predicting whether two arguments from the same topic share the same stance. We expect that arguments that share the same stance have higher agreement and lower disagreement scores, on average. Orthogonality is likely mostly independent of the stances from the argument pairs.

Tab. 4 (bottom) shows results for all arguments from the 5 topics – which amounts to 326,836 argument pairs. As expected, we observe higher ROC-AUC values for predictions from agreement and disagreement scores compared to predictions from orthogonal scores. Also, disagreement using $\mathcal{P}_0$ performs best. This aligns with our previous results, where i) disagreement scores were higher than agreement scores and ii) $\mathcal{P}_0$ was the best aggregation method. Fig. 2 and 9 confirm that argument pairs of the same and diverging stances form distinct distributions for disagreement with $\mathcal{P}_0$.

Acceptability correlation. As a final sanity check, in Fig. 3 we assess the correlation between acceptability scores by plotting the disagreement scores of argument pairs against their agreement scores. As expected, high agreement occurs with low disagreement and vice versa, while high orthogonality scores occur when both, agreement and disagreement, are low.

5 Case Study

For our case study we construct PSVs with a signature of 100 perspectives and `GPT4o` (zero-shot) to predict PSV stance values. As aggregation method we use $\mathcal{P}_0$. We discuss our findings for the issue "*Should animal hunting be banned?*". §C shows results for all 5 topics from our annotation.

We first look at global acceptability scores between different stakeholder groups, as shown in Figs. 4 and 11. We observe that related stakeholder groups have higher agreement scores among each other, for example *animal rights activists* and *environmentalists*. Between these two and opposing stakeholder groups, such as *hunters*, the disagreement scores are highest. Orthogonality scores are highest between the stakeholder groups *government officials*, *hunters* and *local communities*. Indeed, *government officials* and *local communities* are vague and potentially diverse stakeholder groups for the given issue, which could explain why more arguments are orthogonal to each other.

Fig. 5 and Tab. 9 show the top-3 and top-5 perspectives per acceptability score (agreement, orthogonal, and disagreement). Across all arguments,

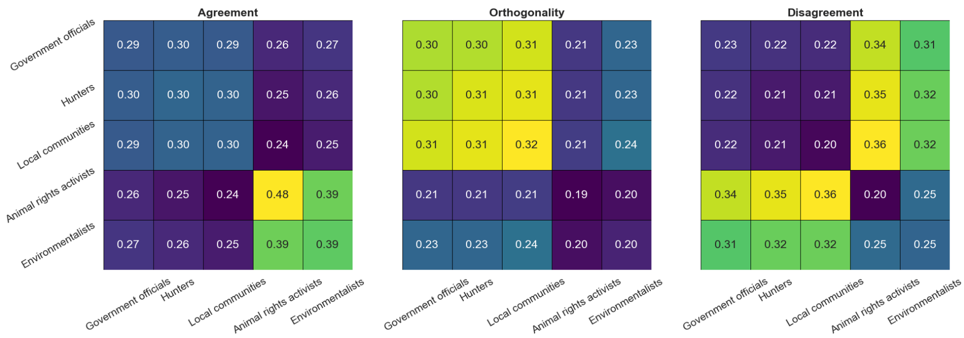

Figure 4: Agreement scores among stakeholder groups for '*Animal Hunting*'. Fig. 10 shows results for other topics.

the perspectives with highest agreement are perspectives that reflect a socially agreed 'stigma', such as *poaching*, *stabbing to death* or *people who exploit animals* – while perspectives with high disagreement reflect the overall stances in the debate (*hunt game*, *while hunting animals*). Orthogonal perspectives occur in only a few arguments, and hence are less relevant for the debated issue (*sex*, *sexual activity*, *water*).

Naturally, the agreement and disagreement perspectives depend on the overall stances of the compared arguments (cf. Tab. 9 and Fig. 12). When comparing arguments from opposing stances, the results remain in line with our findings in the previous paragraph. However, when the compared arguments have the same stance, the agreement and disagreement perspectives shift. For example *hunt game*, which is a disagreement perspective across all arguments, is agreed upon among authors who are against banning animal hunting. They also agree on *poaching*, showing that also people in favor of animal hunting disapprove poaching.[3] Disagreement perspectives can reveal differences between arguers that share a global stance: *control* and *pleasure* are not agreed upon by all arguers in favor of hunting, showing that they are in favor of hunting for different reasons.

Finally, we compare agreement and disagreement perspectives depending on the stakeholder groups we infer for the compared arguments. Here we can identify conflicts among (similar) stakeholder groups (for example *killing for food* for *animal rights activists* and *environmentalists*) as well as common ground among (conflicting) stakeholder groups (for example *poaching* or *people who exploit animals* for *hunters* and *environmentalists*).

[3]We identify that they are against poaching from the PSV stance values – agreement scores alone do not express whether the agreement stems from positive or negative stance values.

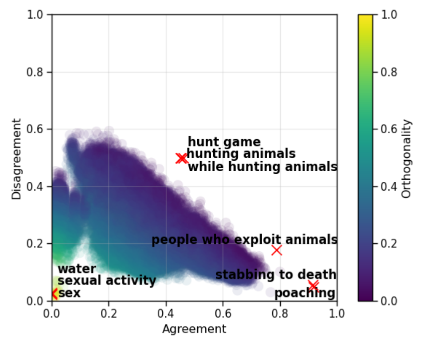

Figure 5: (Dis)agreement of selected perspectives.

Identifying such conflicts and shared understandings can help to better understand different opinions and hence, is a crucial step in deliberation.

6 Conclusions

We present **P**erspectivized **S**tance **V**ectors (PSVs) – a novel approach to represent fine-grained perspectives expressed in arguments on a debated topic. PSVs effectively identify and explain mutual (dis)-agreement between arguments and potential stakeholder groups, offering deep interpretability by revealing issue-specific perspectives driving such (dis)agreements. Identifying (dis)agreement perspectives can reveal the underlying reasons for conflicting viewpoints, and how they can potentially be resolved. Thus, we believe that our fine-grained analysis of perspectives using PSVs provides a valuable contribution to the growing field of deliberative decision making.

Acknowledgments

This work was funded by DFG, the German Research Foundation, for the project "ACCEPT: Perspectivized Argument Knowledge Graphs for Deliberation", as part of the priority program "RATIO: Robust Argumentation Machines" (SPP-1999).

Limitations

We evaluate our approach on PAKT (Plenz et al., 2024), which is limited to English arguments from a predominantly US-context. As our PSV construction partially relies on LMs, it is to be expected that the quality of individual PSVs would be lower for data from a different background. However, our aggregation method is language- and culture agnostic and thus should be robust.

Where possible we assess the quality of our approach automatically using data which is already available in large amounts (cf. §4.3.3). However, for fine-grained stance values and acceptability scores we had to rely on our manually annotated data. The annotated data covers 5 topics with 10 arguments each, which may seem like a rather small resource. However, collecting this data was a considerable annotation effort since we required annotations for argument pairs at the level of distinct perspectives. As our experiments are supported with a large-scale case study, we believe that our findings are reliable and trustworthy.

Finally, predicting aspects of a debated issue which a group of arguments / authors agrees or disagrees on is a challenging task. Reducing this task to our PSV framework might cause oversimplifications. Nonetheless, we study these structured representations of arguments for two good reasons. First, they are highly interpretable – which we believe to be important for deliberation tasks, to enhance (i) the trust of users and (ii) control for moderators. Second, large parts of our method have to be unsupervised due to a lack of training data. This makes training of end-to-end models infeasible. That being said, a more flexible end-to-end system might be able to obtain better performance in future work, for example, by creating larger amounts of partially annotated data, using our methods.

References

Yamen Ajjour, Milad Alshomary, Henning Wachsmuth, and Benno Stein. 2019. Modeling frames in argumentation. 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)*, pages 2922–2932, Hong Kong, China. Association for Computational Linguistics.

Khalid Al-Khatib, Henning Wachsmuth, Kevin Lang, Jakob Herpel, Matthias Hagen, and Benno Stein. 2018. Modeling deliberative argumentation strategies on Wikipedia. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2545–2555, Melbourne, Australia. Association for Computational Linguistics.

Milad Alshomary, Roxanne El Baff, Timon Gurcke, and Henning Wachsmuth. 2022. The moral debater: A study on the computational generation of morally framed arguments. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8782–8797, Dublin, Ireland. Association for Computational Linguistics.

Leila Amgoud and Srdjan Vesic. 2014. Rich preference-based argumentation frameworks. *International Journal of Approxoximate Reasoning*, 55(2).

Michiel Bakker, Martin Chadwick, Hannah Sheahan, Michael Tessler, Lucy Campbell-Gillingham, Jan Balaguer, Nat McAleese, Amelia Glaese, John Aslanides, Matt Botvinick, and Christopher Summerfield. 2022. Fine-tuning language models to find agreement among humans with diverse preferences. In *Advances in Neural Information Processing Systems*, volume 35, pages 38176–38189. Curran Associates, Inc.

Ziyi Bao, Kristijonas Čyras, and Francesca Toni. 2017. Abaplus: Attack reversal in abstract and structured argumentation with preferences. In *PRIMA 2017: Principles and Practice of Multi-Agent Systems*, pages 420–437, Cham. Springer International Publishing.

Roy Bar-Haim, Indrajit Bhattacharya, Francesco Dinuzzo, Amrita Saha, and Noam Slonim. 2017. Stance classification of context-dependent claims. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pages 251–261, Valencia, Spain. Association for Computational Linguistics.

Joe Barrow, Rajiv Jain, Nedim Lipka, Franck Dernoncourt, Vlad Morariu, Varun Manjunatha, Douglas Oard, Philip Resnik, and Henning Wachsmuth. 2021. Syntopical graphs for computational argumentation tasks. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1583–1595, Online. Association for Computational Linguistics.

Ralph Bergmann, Ralf Schenkel, Lorik Dumani, and Stefan Ollinger. 2018. ReCAP - Information Retrieval and Case-based Reasoning for Robust Deliberation and Synthesis of Arguments in the Political Discourse. In *Proceedings of LWDA*.

Gianni Brauwers and Flavius Frasincar. 2022. A survey on aspect-based sentiment classification. *ACM Comput. Surv.*, 55(4).

Laura Cabello and Uchenna Akujuobi. 2024. It is simple sometimes: A study on improving aspect-based sentiment analysis performance. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Bangkok, Thailand. Association for Computational Linguistics.

Neele Falk, Andreas Waldis, and Iryna Gurevych. 2024. Overview of PerpectiveArg2024 the first shared task on perspective argument retrieval. In *Proceedings of the 11th Workshop on Argument Mining (ArgMining 2024)*, pages 130–149, Bangkok, Thailand. Association for Computational Linguistics.

Mark Felton, Amanda Crowell, Merce Garcia-Mila, and Constanza Villarroel. 2022. Capturing deliberative argument: An analytic coding scheme for studying argumentative dialogue and its benefits for learning. *Learning, Culture and Social Interaction*, 36:100350. Pragmatics, Education and Argumentation.

Xiyan Fu and Anette Frank. 2024. The Mystery of Compositional Generalization in Graph-based Generative Commonsense Reasoning. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 8376–8394, Miami, Florida, USA. Association for Computational Linguistics.

Ivan Habernal and Iryna Gurevych. 2016. What makes a convincing argument? empirical analysis and detecting attributes of convincingness in web argumentation. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1214–1223, Austin, Texas. Association for Computational Linguistics.

Philipp Heinisch, Anette Frank, Juri Opitz, and Philipp Cimiano. 2022. Strategies for Framing Argumentative Conclusion Generation. In *Proceedings of the 15th International Natural Language Generation Conference*. Association for Computational Linguistics.

Philipp Heinisch, Dimitry Mindlin, and Philipp Cimiano. 2023. Unsupervised argument reframing with a counterfactual-based approach. In *Proceedings of the 10th Workshop on Argument Mining*, pages 107–119, Singapore. Association for Computational Linguistics.

Mickel Hoang, Oskar Alija Bihorac, and Jacobo Rouces. 2019. Aspect-based sentiment analysis using BERT. In *Proceedings of the 22nd Nordic Conference on Computational Linguistics*, pages 187–196, Turku, Finland. Linköping University Electronic Press.

Yufang Hou and Charles Jochim. 2017. Argument relation classification using a joint inference model. In *Proceedings of the 4th Workshop on Argument Mining*, pages 60–66, Copenhagen, Denmark. Association for Computational Linguistics.

Johannes Kiesel, Milad Alshomary, Nicolas Handke, Xiaoni Cai, Henning Wachsmuth, and Benno Stein. 2022. Identifying the human values behind arguments. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4459–4471, Dublin, Ireland. Association for Computational Linguistics.

Jonathan Kobbe, Ioana Hulpuș, and Heiner Stuckenschmidt. 2020. Unsupervised stance detection for arguments from consequences. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 50–60, Online. Association for Computational Linguistics.

Erik Körner, Gregor Wiedemann, Ahmad Dawar Hakimi, Gerhard Heyer, and Martin Potthast. 2021. On classifying whether two texts are on the same side of an argument. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 10130–10138, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.

Klaus Krippendorff. 2019. *Content Analysis: An Introduction to Its Methodology*. SAGE.

Daniel Loureiro, Francesco Barbieri, Leonardo Neves, Luis Espinosa Anke, and Jose Camacho-collados. 2022. TimeLMs: Diachronic language models from Twitter. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 251–260, Dublin, Ireland. Association for Computational Linguistics.

Yiwei Luo, Dallas Card, and Dan Jurafsky. 2020. Detecting stance in media on global warming. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 3296–3315, Online. Association for Computational Linguistics.

Nailia Mirzakhmedova, Johannes Kiesel, Milad Alshomary, Maximilian Heinrich, Nicolas Handke, Xiaoni Cai, Valentin Barriere, Doratossadat Dastgheib, Omid Ghahroodi, MohammadAli SadraeiJavaheri, Ehsaneddin Asgari, Lea Kawaletz, Henning Wachsmuth, and Benno Stein. 2024. The touché23-ValueEval dataset for identifying human values behind arguments. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 16121–16134, Torino, Italia. ELRA and ICCL.

Sanjay Modgil and Michael Luck. 2009. Argumentation based resolution of conflicts between desires and normative goals. In *Argumentation in Multi-Agent Systems*, pages 19–36, Berlin, Heidelberg. Springer Berlin Heidelberg.

Saif M. Mohammad, Parinaz Sobhani, and Svetlana Kiritchenko. 2017. Stance and sentiment in tweets. *Special Section of the ACM Transactions on Internet Technology on Argumentation in Social Media*, 17(3).

Juri Opitz, Philipp Heinisch, Philipp Wiesenbach, Philipp Cimiano, and Anette Frank. 2021. Explainable unsupervised argument similarity rating with Abstract Meaning Representation and conclusion generation. In *Proceedings of the 8th Workshop on Argu-

ment Mining, pages 24–35, Punta Cana, Dominican Republic. Association for Computational Linguistics.

Yulia Otmakhova, Shima Khanehzar, and Lea Frermann. 2024. Media framing: A typology and survey of computational approaches across disciplines. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15407–15428, Bangkok, Thailand. Association for Computational Linguistics.

Debjit Paul, Juri Opitz, Maria Becker, Jonathan Kobbe, Graeme Hirst, and Anette Frank. 2020. Argumentative Relation Classification with Background Knowledge. In *Proceedings of the 8th International Conference on Computational Models of Argument (COMMA 2020)*, volume 326 of *Frontiers in Artificial Intelligence and Applications*. Computational Models of Argument.

Moritz Plenz, Raphael Buchmüller, and Alexander Bondarenko. 2023a. Argument Quality Prediction for Ranking Documents. In *Working Notes Papers of the CLEF 2023 Evaluation Labs*, volume 3497 of *CEUR Workshop Proceedings*, Thessaloniki, Greece.

Moritz Plenz and Anette Frank. 2024. Graph language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4477–4494, Bangkok, Thailand. Association for Computational Linguistics.

Moritz Plenz, Philipp Heinisch, Anette Frank, and Philipp Cimiano. 2024. PAKT: Perspectivized Argumentation Knowledge Graph and Tool for Deliberation Analysis. In *Proceedings of the 1st International Conference on Recent Advances in Robust Argumentation Machines (RATIO-24)*. Springer.

Moritz Plenz, Juri Opitz, Philipp Heinisch, Philipp Cimiano, and Anette Frank. 2023b. Similarity-weighted construction of contextualized commonsense knowledge graphs for knowledge-intense argumentation tasks. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6130–6158, Toronto, Canada. Association for Computational Linguistics.

Victor Hugo Nascimento Rocha, Igor Cataneo Silveira, Paulo Pirozelli, Denis Deratani Mauá, and Fabio Gagliardi Cozman. 2023. *Assessing Good, Bad and Ugly Arguments Generated by ChatGPT: a New Dataset, its Methodology and Associated Tasks*, page 428–440. Springer Nature Switzerland.

Mattes Ruckdeschel and Gregor Wiedemann. 2022. Boundary detection and categorization of argument aspects via supervised learning. In *Proceedings of the 9th Workshop on Argument Mining*, pages 126–136, Online and in Gyeongju, Republic of Korea. International Conference on Computational Linguistics.

Ameer Saadat-Yazdi, Jeff Z. Pan, and Nadin Kokciyan. 2023. Uncovering implicit inferences for improved relational argument mining. In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 2484–2495, Dubrovnik, Croatia. Association for Computational Linguistics.

Swarnadeep Saha, Prateek Yadav, Lisa Bauer, and Mohit Bansal. 2021. ExplaGraphs: An explanation graph generation task for structured commonsense reasoning. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 7716–7740, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.

Mark Snaith, John Lawrence, and Chris Reed. 2010. Mixed initiative argument in public deliberation. In *Proceedings of Online Deliberation (OD 2010)*.

Parinaz Sobhani, Saif M. Mohammad, and Svetlana Kiritchenko. 2016. Detecting stance in tweets and analyzing its interaction with sentiment. In *Proceedings of the Fifth Joint Conference on Lexical and Computational Semantics (*Sem)*, Berlin, Germany.

Robyn Speer, Joshua Chin, and Catherine Havasi. 2017. Conceptnet 5.5: An open multilingual graph of general knowledge. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence and Thirty-Fifth Conference on Innovative Applications of Artificial Intelligence and Thirteenth Symposium on Educational Advances in Artificial Intelligence*, pages 4444–4451.

Henning Wachsmuth, Shahbaz Syed, and Benno Stein. 2018. Retrieval of the best counterargument without prior topic knowledge. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 241–251, Melbourne, Australia. Association for Computational Linguistics.

Douglas Walton, Katie Atkinson, Trevor Bench-Capon, Adam Wyner, and Dan Cartwright. 2010. Argumentation in the framework of deliberation dialogue. In *Arguing Global Governance*. Imprint Routledge.

Douglas Walton, Alice Toniolo, and Timothy J. Norman. 2016. Towards a richer model of deliberation dialogue: Closure problem and change of circumstances. *Argument & Computation*, 7(2-3):155–173.

Haining Wang, Kang He, Bobo Li, Lei Chen, Fei Li, Xu Han, Chong Teng, and Donghong Ji. 2024. Refining and synthesis: A simple yet effective data augmentation framework for cross-domain aspect-based sentiment analysis. In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 10318–10329, Bangkok, Thailand. Association for Computational Linguistics.

Lu Wang, Nick Beauchamp, Sarah Shugars, and Kechen Qin. 2017. Winning on the merits: The joint effects of content and style on debate outcomes. *Transactions of the Association for Computational Linguistics*, 5:219–232.

Bowen Zhang, Daijun Ding, Liwen Jing, Genan Dai, and Nan Yin. 2024. How would stance detection techniques evolve after the launch of chatgpt? *Preprint*, arXiv:2212.14548.

Xiaotian Zhang, Yanjun Zheng, Hang Yan, and Xipeng Qiu. 2023. Investigating glyph-phonetic information for Chinese spell checking: What works and what's next? In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 1–13, Toronto, Canada. Association for Computational Linguistics.

Chenye Zhao, Yingjie Li, Cornelia Caragea, and Yue Zhang. 2024. ZeroStance: Leveraging ChatGPT for open-domain stance detection via dataset generation. In *Findings of the Association for Computational Linguistics ACL 2024*, pages 13390–13405, Bangkok, Thailand and virtual meeting. Association for Computational Linguistics.

A Method

A.1 Signature

Concept selection The commonsense knowledge graphs are taken from the published data of Plenz et al. (2024). Lemmatization was performed with the `en_core_web_trf` model from Spacy. Future work could experiment with supervised concept selection, e.g., by finetuning models (Plenz and Frank, 2024) designed for knowledge graphs such as ConceptNet.

Hypernym filtering. We identify hypernyms using the NLTK implementation of WordNet. To allow for greater coverage we check for hypernyms within the lemmatized set of concepts. For each concept we only consider the first synset, and do not remove concepts which do not have a synset in WordNet.

ChatGPT-based relevance filtering. We use `ChatGPT-3.5-0125` to assign each concept with a score to reflect its *relevance for a given issue*, using the following prompt:

> We plan to compare arguments depending on which concepts they evoke. Therefore, we created a catalog of concepts for each issue. For the following concept, decide whether it is relevant for the given issue:
> 1: yes
> 2: no
> Example Annotation for issue 'gun control':
> arm themselves: 1
> control: 1
> criminals: 1
> dangerous: 1
> laws regulate who: 1
> own guns: 1
> police: 1
> politics: 1
> shooting guns: 1
> wrong: 1
>
> Issue: {*debate topic*}
> Concept: {*concept*}

The prompt is taken from our annotation guidelines. Depending on ChatGPT's output, we assign a binary label (*relevant / irrelevant*) to each concept, for each issue. The resulting labels are used as a filter to remove unrelated concepts.

To the best of our knowledge we are the first to use ChatGPT to assess the relevance of concepts for a given debate topic. Our human annotation indeed verifies that filtering with ChatGPT can boost precision for relevance (Table 2). In a more general scope, ChatGPT was successfully used for many argument classification tasks such as quality (Rocha et al., 2023; Plenz et al., 2023a) and stance (Zhao et al., 2024; Zhang et al., 2024, 2023; Plenz et al., 2023a) classification, which motivates our approach.

A.2 Stance values

`RoBERTa`. We compose a synthetic dataset using the stance dataset of Sobhani et al. (2016); Mohammad et al. (2017) and the dataset on human-values detection by Mirzakhmedova et al. (2024). In the stance dataset, which is based on annotated tweets, we select those tweets that address the annotated target, by being against, in favor, or none of those (neutral). To increase the target diversity, we map each of the six targets to a hand-crafted set of synonyms and antonyms[4]. To adapt the genre (from short tweets to more comprehensive arguments), we concatenate up to four tweets toward the same target to new instances[5]. We follow the same procedure with the human-values detection dataset on arguments, where we treat an annotated encouragement of a human value as *in favor of* a set of

[4]In case of antonyms, we switch the classes against and in favor.

[5]The aggregation of the stances is *neutral* in case of a neutralized/balanced set of stances, *in favor* if at least one stance is in favor and no stance is against, *against* if at least one stance is against and no stance in favor, else the instance is dropped.

hand-crafted associated concepts and *against* a set of hand-crafted contrastive concepts. We consider a set of hand-crafted associated concepts of human values that are labeled as not relevant for this argument as the *neutral* class.

We start from a `roberta-base` model that was already fine-tuned on the task of sentiment prediction by Loureiro et al. (2022)[6]. We randomly split the synthetic dataset into 80% training, 10% development and 10% test data, which we then use to further fine-tune the model with a learning rate of $2e-5$ and early stopping (evaluated after every 1000 processed train instances). We use a modified mean squared error as loss function:

$$\begin{aligned} l = \; & \textstyle\sum_{y_{\text{against}}} \lambda_{\text{against}} (\hat{y}^2_{\text{neutral}} + \hat{y}_{\text{for}})^2 \\ + \; & \textstyle\sum_{y_{\text{neutral}}} \lambda_{\text{neutral}} (\hat{y}_{\text{against}} + \hat{y}_{\text{for}})^2 \\ + \; & \textstyle\sum_{y_{\text{for}}} \lambda_{\text{for}} (\hat{y}_{\text{against}} + \hat{y}^2_{\text{neutral}})^2 \end{aligned} \tag{3}$$

Here, λ are hyperparameters to control the weighting between the three classes AGAINST, NEUTRAL, and FOR, where y represents the target class and $\hat{y}_c$ the predicted probability ($[0, 1]$) for class c. We choose the model that performed best across 8 runs on the synthetic test split as the final model for **P**erspectivized **S**tance **V**ector (PSV)-stance inference, once with $\lambda_{\text{against}} = \lambda_{\text{neutral}} = \lambda_{\text{for}} = 1$, denoted as *1*-`RoBERTa`, and once as an ensemble of three models which were trained with ($\lambda_{\text{against}} \in \{1, 1.33\}$, $\lambda_{\text{for}} \in \{1, 1.33\}$ and $\lambda_{\text{neutral}} = 1$), denoted as *3*-`RoBERTa`.

GPT4o. We use `gpt-4o-2024-11-20` in a zero-shot setting (`GPT4o` *(zero-shot)*) and a setting using 2 handcrafted examples showcasing that the overall sentiment of the argument does not need to correlate with the stance toward a specific concept (`GPT4o` *(few-shot)*). The specific prompt used is:

> Given a controversial topic discussed by a given argument, and an aspect, provide an one-word answer to the following question: Considering a person writing this argument, what is their attitude towards the given aspect: negative, neutral, or positive?

[6]The model is available from at https://huggingface.co/cardiffnlp/twitter-roberta-base-sentiment-latest. We also tried non-fine-tuned RoBERTa models, but they performed worse in preliminary analysis.

A.3 Aggregation Methods for *Acceptability* Scores

We first discuss the *stance value-based* aggregation methods $\mathcal{S}$, $\mathcal{S}_0$ and $\mathcal{S}_D$.

We use the notation from Tab. 1, where $s_i^j \in \{-1, 0, 1\}$ is the stance value of argument j towards concept i. Kronecker delta $\delta(x, y)$ is 1 if $x = y$ and 0 else.

Stance Value $\mathcal{S}$. For agreement ($\mathcal{S}^+(s_i^1, s_i^2) = \delta(s_i^1, s_i^2)$) we simply consider whether the arguments have the same stance value towards concept c_i – if yes, the agreement is 1 and otherwise 0. For disagreement ($\mathcal{S}^-(s_i^1, s_i^2) = 1 - \delta(s_i^1, s_i^2)$) we instead check whether the stance values are different – different values means the disagreement is 1, and otherwise it is 0.

Stance Value (Considering Neutral) $\mathcal{S}_0$. For this adaptation we consider the special role of neutral stance values: If both arguers are neutral towards a concept c_i, c_i is not a meaningful indicator of agreement. Hence, we only record a perspectivised agreement if both stance values are in favor or against the concept – but not if both are neutral: $\mathcal{S}_0^+(s_i^1, s_i^2) = \delta(s_i^1, s_i^2)\left(1 - \delta(s_i^1, 0)\right)$.

Similarly, we do not record disagreement if at least one of the arguers is neutral towards a perspective c_i. Hence, we only record disagreement if one stance value is in favor, while the other is against: $\mathcal{S}_0^-(s_i^1, s_i^2) = \left(1 - \delta(s_i^1, s_i^2)\right)\left(1 - \delta(s_i^1, 0)\right)\left(1 - \delta(s_i^2, 0)\right)$.

For Orthogonality it is enough if at least one arguer is neutral towards the given perspective c_i. Hence, the orthogonality score is 1 if at least one arguer is neutral, and otherwise 0: $\mathcal{S}_0^{\varnothing} = \min\left(\sum_{j=1}^{2} \delta\left(s_i^j, 0\right), 1\right)$

Stance Value (Difference) $\mathcal{S}_D$. For agreement the previous approach might have the limitation that disagreement and orthogonal concepts both contribute 0 when computing a global agreement score by averaging over all concepts. It could be more expressive to have disagreement concepts *reduce* the overall agreement, instead of treating them the same as orthogonal concepts. Hence, we designed $\mathcal{S}_D^+(s_i^1, s_i^2) = \mathcal{S}_0^+(s_i^1, s_i^2) - \mathcal{S}_0^-(s_i^1, s_i^2)$ to be $+1$, 0 and -1 for agreement, orthogonal and disagreement concepts, respectively. Analogously, we constructed $\mathcal{S}_D^-(s_i^1, s_i^2) = \mathcal{S}_0^-(s_i^1, s_i^2) - \mathcal{S}_0^+(s_i^1, s_i^2)$.

Stance Probability $\mathcal{P}$. Aggregation methods that are based on *stance probabilities* have similar motivations. Their advantage, compared to the above methods based on stance values, is that they return continuous values for each concept individually.

For *agreement*, we take the element-wise multiplication $\odot$ and sum the resulting values – optionally disregarding the neutrality scores. If the probability mass is similar (and not on neutral) for both arguments, then the agreement score is high. For *orthogonality* we also consider element-wise multiplication.

For *disagreement*, we take the element-wise difference between the probabilities instead, and sum up the absolute values – again, optionally without the neutral contribution. To obtain scores between 0 and 1 we normalize scores with a factor of $1/2$.

A.4 Stakeholder Annotation

Debate portals are commonly lacking information about the authors of arguments, given privacy concerns. To associate arguments with stakeholder information, we enrich our dataset by creating automatic predictions of stakeholder group information using ChatGPT, at issue and argument level.. To this end, we apply `ChatGPT-3.5-0125` for two subtasks: First, we let ChatGPT predict potential stakeholder groups at *issue-level*, by prompting ChatGPT to return a set of relevant stakeholder groups for a given topic:

> A stakeholder is a group of people who are affected by a topic. For example, the topic "Should young children have access to the internet?" has the stakeholders "Children" and "Parents". Return a list of the most important stakeholders for the topic "{*debate topic*}". Return a simple list without explanations. Limit yourself to the few most important ones.

We then use this set of stakeholder groups proposed by the model as input to a second call, to assign stakeholder types at *argument-level*. We ask ChatGPT to *select*, from the given set of possible stakeholders, those types of stakeholders that could could plausibly utter a given argument from the relevant topic, using the following prompt:

> Here is an argument from someone: '{*argument*}'. Which of these stakeholders are most likely to utter this argument: {*stakeholder set*}? Return a list of stakeholders without additional information. Multiple may apply.

We extract the stakeholder groups the model predicts to extend each argument in our dataset. This typically results in 1–3 stakeholder labels per argument.

As detailed below in §B.1, we verify the predictions obtained by ChatGPT with manual annotations. Our evaluation shows that the predicted stakeholder groups are highly consistent with human judgment, at issue and argument level, and that only a small number of stakeholder groups is missing from ChatGPT's generated list.

A.5 Ablation without PSV: Pairwise `GPT4o` prompting

We use `gpt-4o-2024-11-20`, which is the same model used to predict the perspectivized stances (c.f. §3.1.2). The prompt is

> Arguments of opposite stance can have agreements – even though they don't agree on the issue at a binary level. Similarly, arguments with the same stance can disagree. We are interested in identifying and specifying such (dis-)agreements. We will present you with two independently written arguments of opposite stance and a list of concepts.
>
> For each concept, annotate whether it is part of the agreement or disagreement:
> ␣␣␣␣1: agreement, i.e., the authors could likely find agreement regarding this concept.
> ␣␣␣␣2: neutral
> ␣␣␣␣3: disagreement, i.e., it is not likely that the authors could agree regarding this concept.
>
> Argument 1: *{argument_1}*
>
> Argument 2: *{argument_1}*
>
> Concepts: *{python_list_of_concepts}*
>
> Return your output as a list of integers, where each integer corresponds to the concept at the same index in the list of concepts. Do not include any additional information in your output.

B Experiments

B.1 Annotation

Table 5 shows the number of annotations and inter-annotator agreement (IAA) scores for different subtasks. IAA is measured using Krippendorff's α (Krippendorff, 2019).

We annotate data from 5 distinct topics:

i) Should animal hunting be banned?
ii) Do you support the death penalty?
iii) Should students get paid for good grades?
iv) Should illegal immigrants be deported?
v) Should kids have to wear school uniforms?

The first topic was annotated by all 3 annotators, to assess IAA. The remaining topics were annotated by only one annotator, allowing us to collect more annotated data. All annotators are experts in computational linguistics and argumentation in particular.

For each topic we annotate 30 signature concepts for i) their relevance with respect to the topic and ii) their granularity. These 30 concepts are the top-15 concepts per stance without any filtering. For each topic, we consider 10 arguments – 5 from each stance. For each of these arguments, we annotate a topic-specific set of 10 concepts (top-5 concepts per stance, with hypernym filtering) for the *PSV stance* (one of the values *for, against, neutral*), yielding 500 annotations in total.

To obtain annotations for *agreement between arguments*, we annotate all arguments pairs of opposite stance within a topic, i.e., 25 argument-pairs per topic, on whether there is *full* or *partial agreement*, *disagreement* or whether they are *orthogonal* to each other. Further, for each of these argument-pairs, we annotate the same *10 concepts* for whether the arguments (or rather the authors of the arguments) would *agree, disagree* or are *neutral* in relation to that concept. The annotated data as well as detailed annotation guidelines will be published upon acceptance.

Table 5 summarizes the number of annotations and shows the IAA scores. We achieve moderate to substantial agreements, except for the concept-level argument pair annotation. We aimed for concise annotation guidelines to reduce the impact of subjective interpretations, but of course they can never be avoided in polarizing debates as we deal with. In particular for the concept-level argument pair annotation there are many subjective factors: interpretations of each of the arguments, as well as what their (dis)agreement is, and interpretation of the concept. This makes this a challenging annotation task, which partially explains the low IAA. However, we note that for disagreement the IAA is $\alpha = 0.21$. Also, two annotators had an IAA of $\alpha = 0.18$ across all classes. These annotators annotated 4 out of 5 topics. Future work can potentially further increase the IAA with multiple training rounds, thereby better calibrating annotators.

Annotation task	# annotations	IAA α
Signature – relevance	150	0.42
Signature – granularity	150	0.42
PSV stances	500	0.60
Argument pairs	125	0.64
Argument pairs – Concept-level	1,250	0.03

Table 5: Number of annotations and the inter-annotator agreement (IAA) measured in Krippendorff's α for different annotation tasks. Note that Krippendorff's α is computed only on one topic, i.e., only on one fifth of the total number of annotations shown in this table.

Label	Frequency [%]
Plausible	29.5
Independent	46.4
Unlikely	24.1

Table 6: Distribution of stakeholder group labels at argument-level(ground truth via majority voting).

We also obtain validating annotations for the stakeholder group predictions that ChatGPT generated at issue- and argument level, (see §A.4).

The stakeholder groups that the model generates as potentially relevant for an issue were judged for *relevancy* at issue level. At argument level, the applicability of a stakeholder to a given argument can be labeled as *plausible*, *unlikely* or as *independent* (meaning that there is no consensus in relating the argument's content to a stakeholder group) – depending on how likely it is that members of a given group are to utter it. The distribution of these labels is displayed in Table 6.

For the five *issues* included in our annotation, ChatGPT generated on average five possibly rel-

Precision	Recall	F1
1.0	78.6	88.0

Table 7: Evaluation of ChatGPT-generated stakeholder annotations on issue-level.

Class	Prec.	Rec.	F1	IAA
Plausible	38.7	63.1	48.0	0.447
Independent	0.0	0.0	0.0	0.228
Unlikely	31.6	67.9	43.1	0.648

Table 8: Evaluation of ChatGPT's stakeholder groups prediction for arguments against human annotation. While ChatGPT was asked for binary judgements (*plausible, unlikely*), human annotators were also offered a third category *independent*. Most indicative is ChatGPT's Recall for the relevant classes *unlikely* and *plausible*, with substantial and moderate IAA (Fleiss' Kappa between three annotators), as opposed to *independent*.

evant stakeholder groups. Our annotators were asked to assign a relevancy label for each generated stakeholder group for a given issue. We also asked them to list any missing stakeholders that could be relevant to the topic. These are interpreted as missing, allowing us to assess recall. Results are displayed in Table 7. While a small amount of relevant stakeholders are missing from ChatGPT's output, all predicted stakeholder groups were validated as being relevant to the larger topic by the majority of annotators.

The validation of predicted stakeholder groups at *argument-level* was performed by human labelers. If annotators disagree on a label, the gold label is determined via majority voting. We compare this ground truth to the predictions of ChatGPT to determine whether its predictions of plausible stakeholder groups at argument level are valid or invalid.

Note, however, that ChatGPT's predictions were restricted to the binary classes *plausible* and *unlikely*, while the annotators were offered an additional label to express that an argument may be *independent* from a given stake holder group (see Table 8). Most indicative of the quality of ChatGPT's prediction is thus Recall for the relevant classes *unlikely* and *plausible*, which show substantial and moderate IAA (Fleiss' Kappa between three annotators), respectively, as opposed to *independent* which achieves only fair IAA quality.

B.2 Confusion matrices for PSV stance value prediction

Figure 6 shows the confusion matrix for `GPT4o` (zero shot).

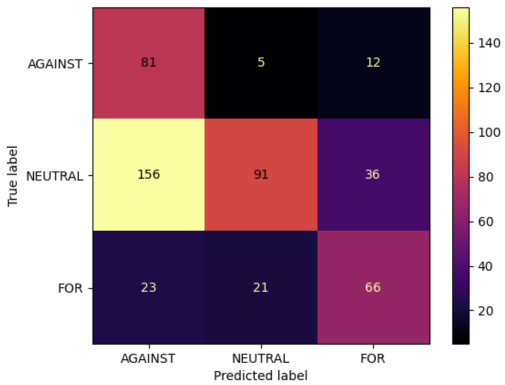

Figure 6: `GPT4o` (zero-shot) confusion matrix for stance value prediction compared to gold annotation.

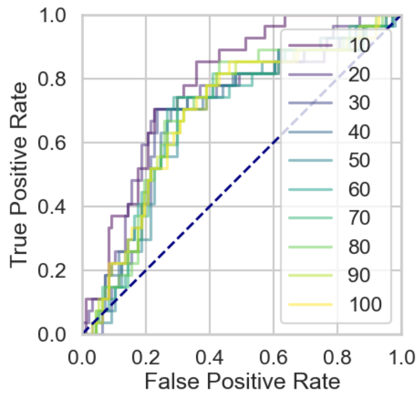

Figure 7: ROC curves for Orthogonality with $\mathcal{P}_0$ for different PSV lengths.

B.3 Impact of PSV length on Acceptability scores

Figure 8 shows the impact of the number of perspectives (i.e., the dimension or length of a PSV) on global agreement, orthogonality and disagreement prediction. Figure 7 shows corresponding ROC curves for orthogonality.

B.4 Same stance prediction

Figure 9 shows the (dis)agreement distributions for argument pairs of same and different stances.

C Case Study

This section collects plots and tables for the case study. Figure 10 presents acceptability scores between different stakeholders. Figure 11 shows argument pairs depending on whether the respective stakeholders are the same or different. Table 9 and Figure 12 show the perspectives with highest acceptability scores depending on the topic and

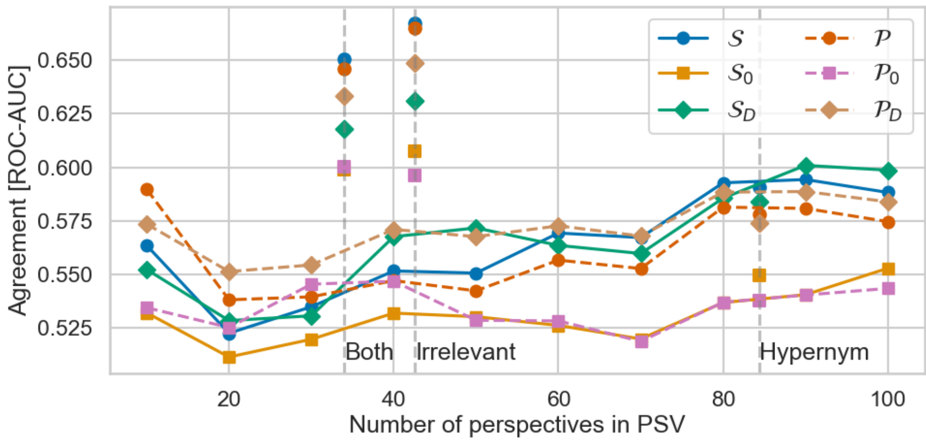

(a) Agreement.

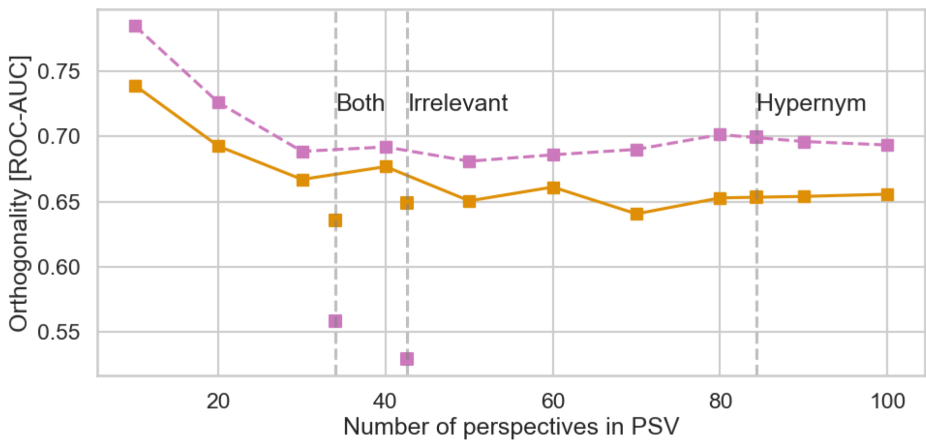

(b) Orthogonality.

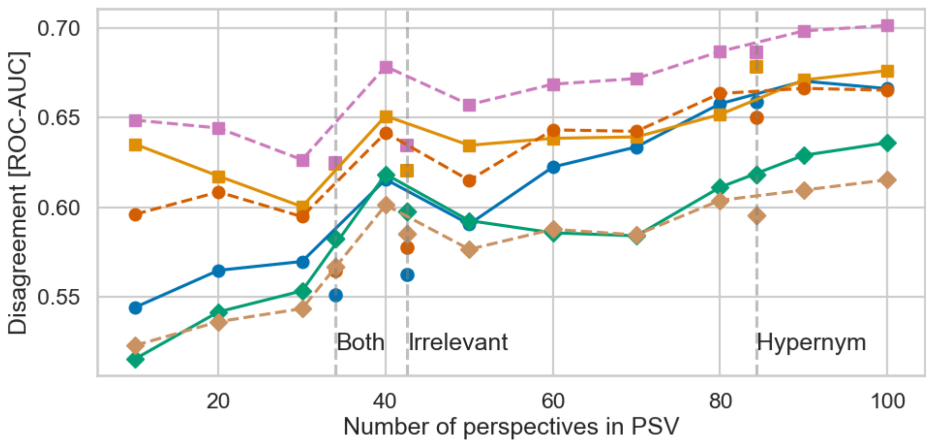

(c) Disagreement.

Figure 8: ROC-AUC scores compared to human annotation depending on PSV length. Filtering options are shown at their average PSV length across the 5 topics.

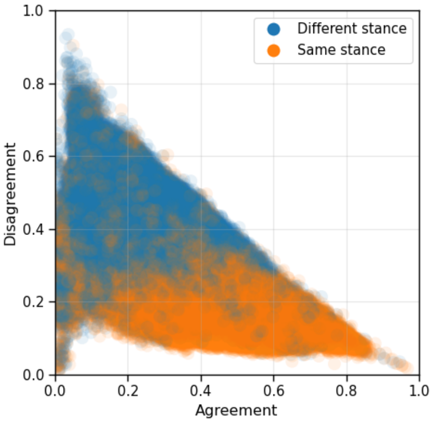

(a) All topics.

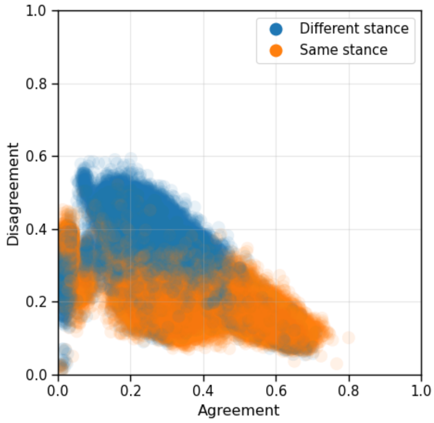

(b) Animal Hunting.

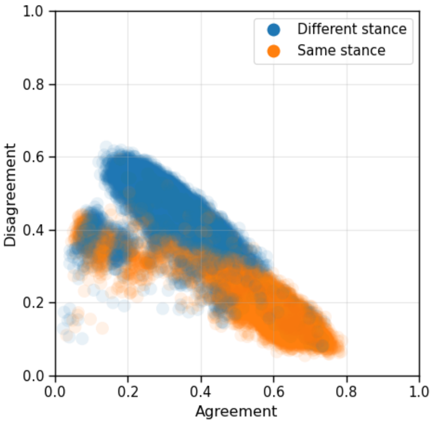

(c) Death Penalty.

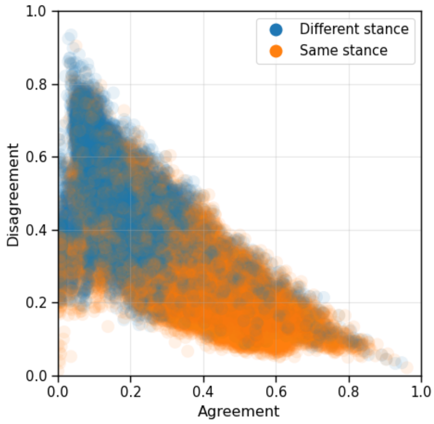

(d) Good Grades.

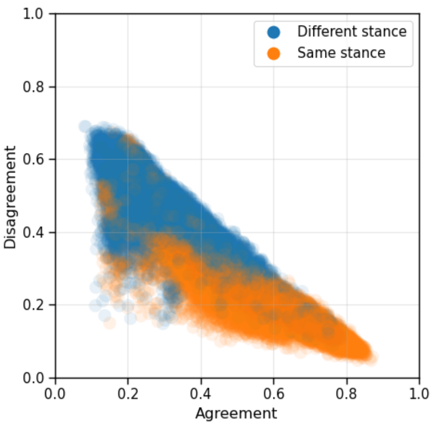

(e) Illegal Immigrants.

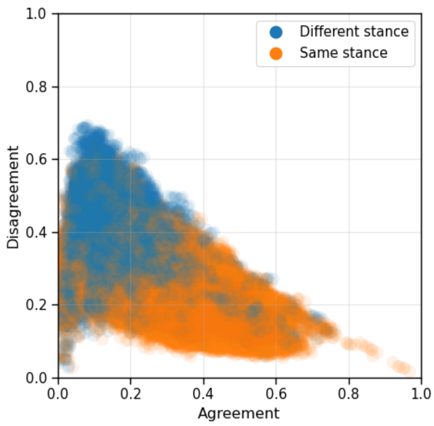

(f) School Uniforms.

Figure 9: Argument pairs depending on their (Dis)agreement scores, colored by same stance.

argument stances. Similarly, Tables 10-14 show the most prominent perspectives depending on the stakeholder groups of the arguments.

D Usage of AI assistants

We use GitHub Copilot (https://github.com/features/copilot) for speeding up programming, and ChatGPT (https://chat.openai.com) to aid with reformulations. The content of this work is our own, and not largely inspired by AI assistants.

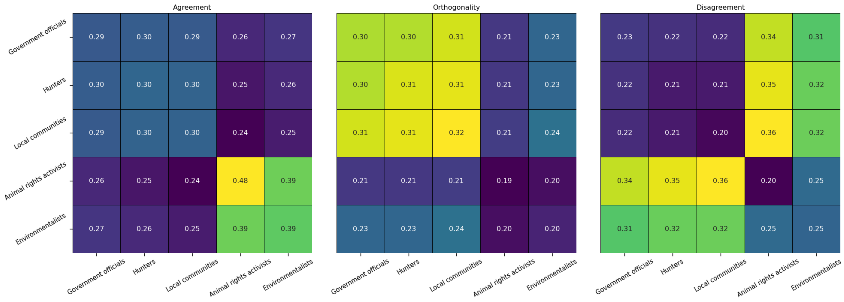

(a) Animal Hunting.

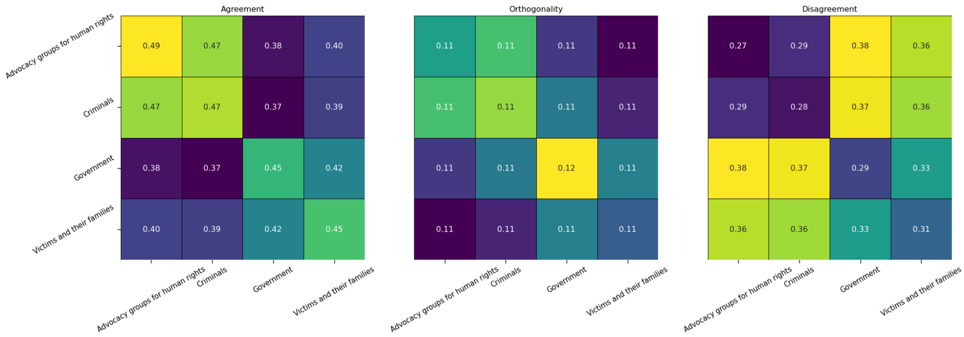

(b) Death Penalty.

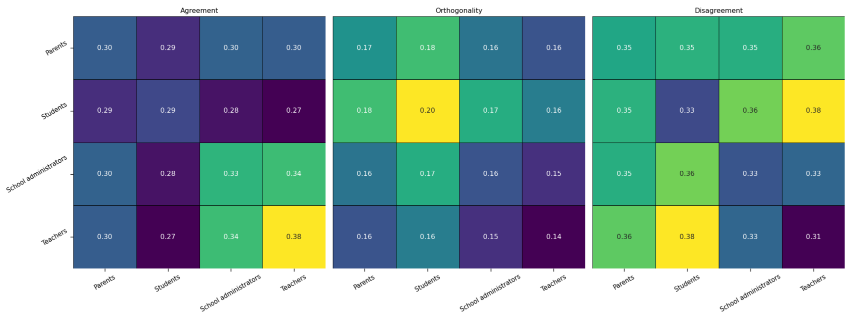

(c) Good Grades.

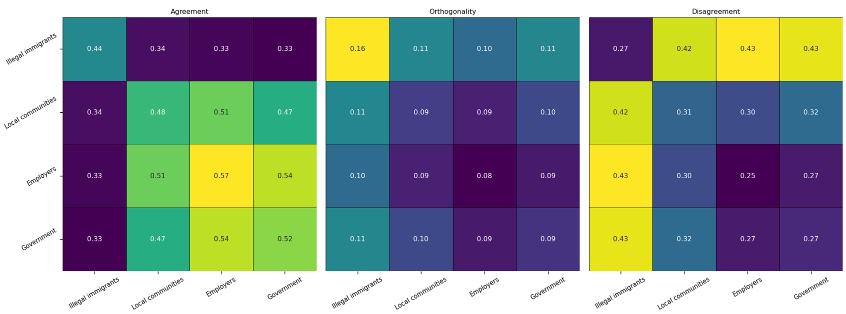

(d) Illegal Immigrants.

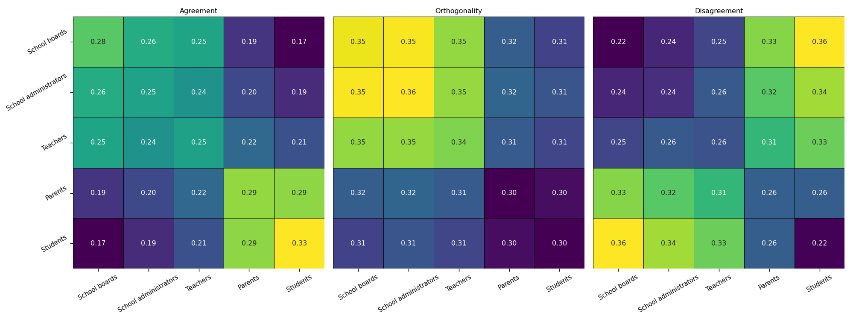

(e) School Uniforms.

Figure 10: Acceptability scores among stakeholder groups for different topics.

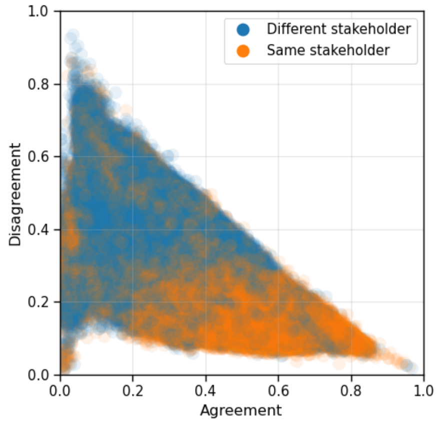

(a) All topics.

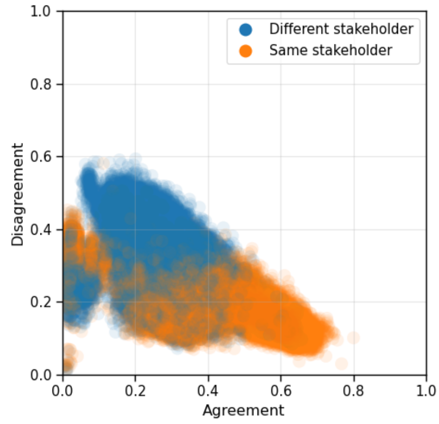

(b) Animal Hunting.

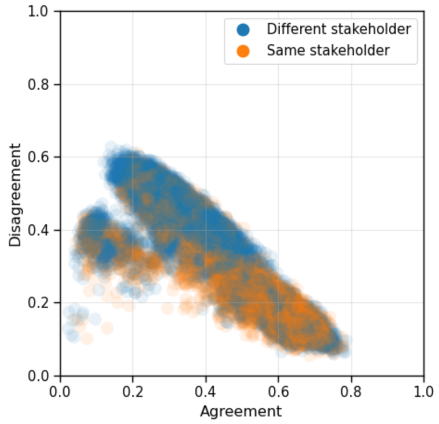

(c) Death Penalty.

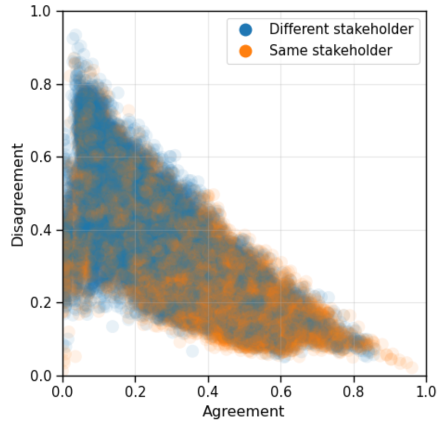

(d) Good Grades.

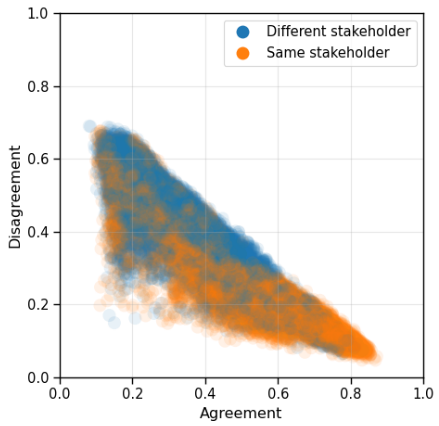

(e) Illegal Immigrants.

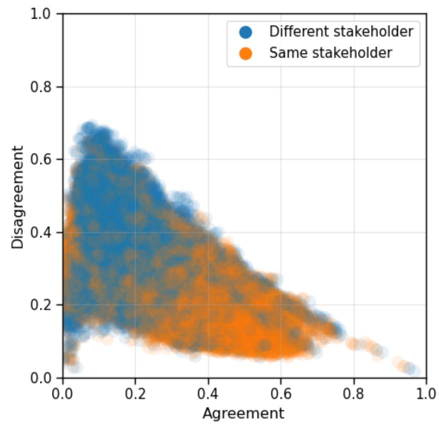

(f) School Uniforms.

Figure 11: Argument pairs depending on their (Dis)agreement scores, colored by whether the stakeholders of the two arguments are the same. As our stakeholder prediction returns a set of stakeholders for each argument, we check whether one set is a subset of the other.

		all	both pro	both con	one pro one con
Animal Hunting	Agreement	poaching, stabbing to death, people who exploit animals, evil, unnatural thing	people who exploit animals, stabbing to death, poaching, blood sport, cruelty	poaching, stabbing to death, hunt game, hunt, wrong	poaching, stabbing to death, people who exploit animals, unnatural thing, evil
	Orthogonality	sex, sexual activity, water, hiking, video game	water, hiking, sex, sexual activity, city	sex, sexual activity, video game, copulating, fly	sex, water, sexual activity, hiking, video game
	Disagreement	hunt game, while hunting animals, hunting animals, hunter, hunt	goal, kind, many wild animals, endangered species, animals and sometimes people	control, pleasure, living thing, kind, humans	hunt game, hunting animals, hunt, hunter, while hunting animals
Death Penalty	Agreement	committing crime, stupid, murdering, murder, injustice	someones who commits murder, committing crime, human right, murderers, crimes	murdering, stupid, kill, human killing, injustice	committing crime, stupid, murdering, murder, crimes
	Orthogonality	album, legs, british, running after ball, play golf	british, album, legs, running after ball, play golf	album, play golf, legs, running after ball, british	album, legs, british, running after ball, play golf
	Disagreement	capital punishment, death sentence, death penalty, right to life, face death penalty	legal, law, change, killed, kill	innocent people, change, human, imprisonment, prosecuted and sent to jail	death sentence, sentenced to death, death penalty, face death penalty, capital punishment
Good Grades	Agreement	pay off teacher, bank on failing in school, bribe, penalty, free	bank on failing in school, pay off teacher, reward, get paid, money	bribe, pay off teacher, twenty bucks for every, buying, get paid	pay off teacher, bank on failing in school, bribe, penalty, free
	Orthogonality	sleep, sports, eat, clean house, acting in play	sports, clean house, eat, acting in play, sleep	sleep, sports, eat, clean house, acting in play	sports, sleep, eat, clean house, acting in play
	Disagreement	reward, make money, get paid, value, money	satisfaction, school, fee, learning, discipline	education, further education, learn lessons well, educate, study	get paid, money, reward, make money, twenty bucks for every
Illegal Immigrants	Agreement	unnatural thing, stealing, criminals, criminal act, people who break laws	people who break laws, illegal, stealing, amnesty, exemption	turn away, deportations, deport, ejection, go home	unnatural thing, stealing, criminals, government, criminal act
	Orthogonality	video game, computing, canada, food, walking	video game, computing, canada, food, walking	computing, video game, walking, canada, food	video game, computing, canada, food, walking
	Disagreement	turn away, ouster, exile, order, return home	law, immigration law, exile, country, order	quality, america, good feelings, exemption, change of location	deportations, deport, amnesty, immigrants, immigrants people who
School Uniforms	Agreement	bad, disguise, touchy about wearing uniforms, pain, special outfit	uniform, school uniform, bad, reason, required for schools to function effectively	school uniform, uniform, touchy about wearing uniforms, required for schools to function effectively, disguise	bad, disguise, very expensive, kids clothing, change
	Orthogonality	food, mathematics, dance, painting, church	church, painting, food, fencing, biology	mathematics, food, dance, sports, plastic surgery	food, dance, mathematics, church, painting
	Disagreement	motivation, expressing yourself, reason, ideal, self esteem	special way to dress, kids clothing, clothes, clothing, changing appearance	change, fashion, expressing yourself, student, dress themselves	school uniform, uniform, required for schools to function effectively, school, ideal

Table 9: Concepts with highest acceptability scores depending on the topic and argument stances.

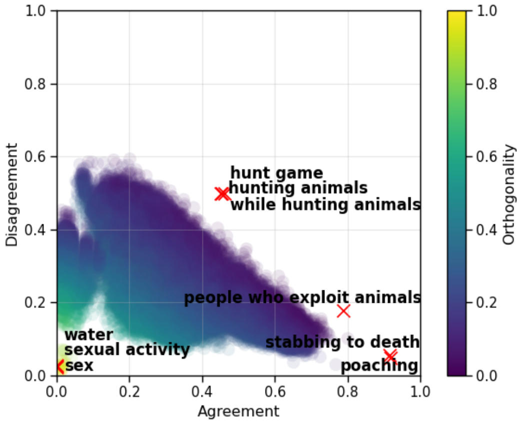

(a) All argument pairs.

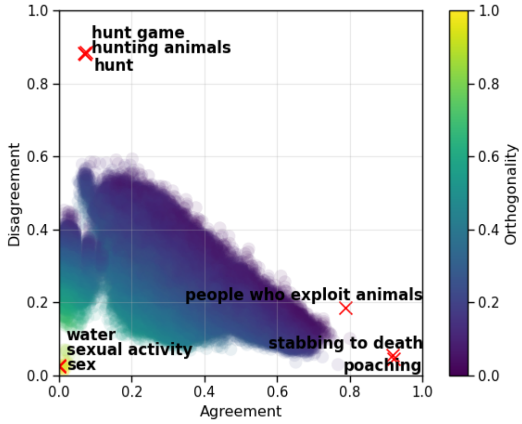

(b) Opposing Stance.

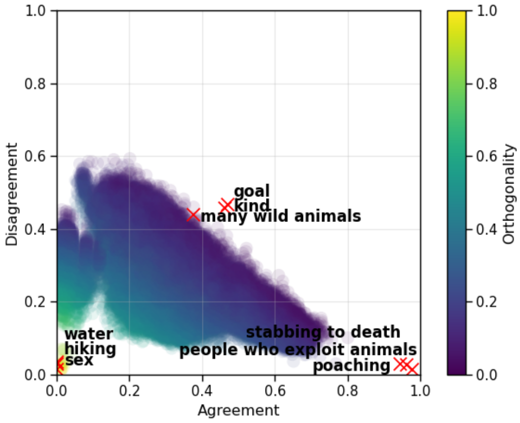

(c) Same Stance: *yes*.

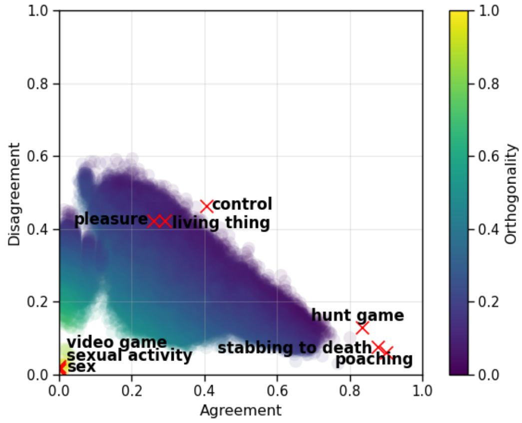

(d) Same Stance: *no*.

Figure 12: Most prominent perspectives for Animal Hunting, depending on the stances of the compared arguments.

	Government officials	Hunters	Local communities	Animal rights activists	Environmentalists
Government officials	poaching, stabbing to death, wrong, hunt game, evil	poaching, stabbing to death, hunt game, hunt, hunting animals	poaching, stabbing to death, hunt game, hunt, hunting animals	stabbing to death, poaching, evil, unnatural thing, selfish	poaching, stabbing to death, evil, wrong, selfish
Hunters	poaching, stabbing to death, hunt game, hunt, hunting animals	poaching, hunt game, hunt, hunting animals, hunter	poaching, hunt game, hunt, hunting animals, hunter	poaching, stabbing to death, evil, people who exploit animals, unnatural thing	poaching, stabbing to death, evil, wrong, people who exploit animals
Local communities	poaching, stabbing to death, hunt game, hunt, hunting animals	poaching, hunt game, hunt, hunting animals, hunter	hunt game, hunt, poaching, hunter, hunting animals	stabbing to death, poaching, unnatural thing, evil, killing people	poaching, stabbing to death, evil, wrong, killing humans
Animal rights activists	stabbing to death, poaching, evil, unnatural thing, selfish	poaching, stabbing to death, evil, people who exploit animals, unnatural thing	stabbing to death, poaching, unnatural thing, evil, killing people	stabbing to death, people who exploit animals, poaching, cruelty, blood sport	stabbing to death, people who exploit animals, poaching, blood sport, unnatural thing
Environmentalists	poaching, stabbing to death, evil, wrong, selfish	poaching, stabbing to death, evil, wrong, people who exploit animals	poaching, stabbing to death, evil, wrong, killing humans	stabbing to death, people who exploit animals, poaching, blood sport, unnatural thing	stabbing to death, poaching, people who exploit animals, blood sport, unnatural thing
Government officials	wild animal, living thing, control, animals, deer	control, living thing, pleasure, wild animal, animals	control, pleasure, living thing, kind, wild animal	hunt game, hunting animals, hunt, while hunting animals, hunter	hunt game, while hunting animals, hunting animals, hunter, hunt
Hunters	control, living thing, pleasure, wild animal, animals	control, pleasure, living thing, kind, joy	control, pleasure, living thing, kind, joy	hunt game, hunting animals, hunt, hunter, while hunting animals	hunt game, hunting animals, hunter, hunt, while hunting animals
Local communities	control, pleasure, living thing, kind, wild animal	control, pleasure, living thing, kind, joy	control, pleasure, kind, joy, living thing	hunt game, hunting animals, hunter, hunt, while hunting animals	hunt game, hunter, hunting animals, while hunting animals, hunt
Animal rights activists	hunt game, hunting animals, hunt, while hunting animals, hunter	hunt game, hunting animals, hunt, hunter, while hunting animals	hunt game, hunting animals, hunter, hunt, while hunting animals	goal, kind, many wild animals, endangered species, animals and sometimes people	goal, killing for food, outdoor activity, getting food, hunt
Environmentalists	hunt game, while hunting animals, hunting animals, hunter, hunt	hunt game, hunting animals, hunter, hunt, while hunting animals	hunt game, hunter, hunting animals, while hunting animals, hunt	goal, killing for food, outdoor activity, getting food, hunt	goal, killing for food, kind, good, getting food

Table 10: Agreement (top) and disagreement (bottom) concepts by Stakeholder for Animal Hunting.

	Advocacy groups for human rights	Criminals	Government	Victims and their families
Advocacy groups for human rights	murdering, committing crime, stupid, murder, injustice	committing crime, murdering, stupid, murder, reward	committing crime, stupid, murdering, murder, crimes	committing crime, stupid, murdering, murder, injustice
Criminals	committing crime, murdering, stupid, murder, reward	committing crime, murdering, stupid, reward, murder	committing crime, stupid, murdering, murder, crimes	committing crime, stupid, murdering, murder, crimes
Government	committing crime, stupid, murdering, murder, crimes	committing crime, stupid, murdering, murder, crimes	committing crime, crimes, someones who commits murder, stupid, murderers	committing crime, stupid, crimes, someones who commits murder, murdering
Victims and their families	committing crime, stupid, murdering, murder, injustice	committing crime, stupid, murdering, murder, crimes	committing crime, stupid, crimes, someones who commits murder, murdering	committing crime, someones who commits murder, stupid, crimes, murderers
Advocacy groups for human rights	human right, rehab, human, change, imprisonment	human right, rehab, human, compassion, change	capital punishment, death penalty, death sentence, right to life, punishable by death	death penalty, death sentence, capital punishment, punishable by death, face death penalty
Criminals	human right, rehab, human, compassion, change	rehab, human, human right, feel remorse, compassion	capital punishment, right to life, death penalty, death sentence, punishable by death	capital punishment, death penalty, punishable by death, death sentence, execute
Government	capital punishment, death penalty, death sentence, right to life, punishable by death	capital punishment, right to life, death penalty, death sentence, punishable by death	kill, human killing, legal, change, deciding criminal s fate	human killing, meant as deterrent to crime, death, kill, die
Victims and their families	death penalty, death sentence, capital punishment, punishable by death, face death penalty	capital punishment, death penalty, punishable by death, death sentence, execute	human killing, meant as deterrent to crime, death, kill, die	kill, human killing, deciding criminal s fate, hanging, change

Table 11: Agreement (top) and disagreement (bottom) concepts by Stakeholder for Death Penalty.

	Parents	Students	School administrators	Teachers
Parents	pay off teacher, bank on failing in school, bribe, penalty, free	pay off teacher, bank on failing in school, bribe, penalty, hard work	pay off teacher, bank on failing in school, bribe, free, penalty	pay off teacher, bank on failing in school, bribe, free, penalty
Students	pay off teacher, bank on failing in school, bribe, penalty, hard work	bank on failing in school, pay off teacher, penalty, hard work, work hard	pay off teacher, bank on failing in school, bribe, penalty, free	pay off teacher, bank on failing in school, bribe, penalty, free
School administrators	pay off teacher, bank on failing in school, bribe, free, penalty	pay off teacher, bank on failing in school, bribe, penalty, free	pay off teacher, bank on failing in school, bribe, free, fee	pay off teacher, bank on failing in school, bribe, free, twenty bucks for every
Teachers	pay off teacher, bank on failing in school, bribe, free, penalty	pay off teacher, bank on failing in school, bribe, penalty, free	pay off teacher, bank on failing in school, bribe, free, twenty bucks for every	pay off teacher, bribe, bank on failing in school, twenty bucks for every, spend money
Parents	get paid, reward, make money, money, feel good	get paid, reward, money, make money, twenty bucks for every	reward, value, make money, money, get paid	reward, make money, money, get paid, value
Students	get paid, reward, money, make money, twenty bucks for every	satisfaction, celebrate, twenty bucks for every, feel good, spend money	get paid, reward, make money, money, value	get paid, make money, reward, money, value
School administrators	reward, value, make money, money, get paid	get paid, reward, make money, money, value	education, better, work, school, make better world	education, better, work, school, make better world
Teachers	reward, make money, money, get paid, value	get paid, make money, reward, money, value	education, better, work, school, make better world	education, educate, learn lessons well, further education, get good grade

Table 12: Agreement (top) and disagreement (bottom) concepts by Stakeholder for Good Grades.

	Illegal immigrants	Local communities	Employers	Government
Illegal immigrants	turn away, attack, exile, go home, roadblock	unnatural thing, stealing, criminals, government, attack	unnatural thing, stealing, government, criminals, attack	unnatural thing, stealing, criminals, attack, government
Local communities	unnatural thing, stealing, criminals, government, attack	unnatural thing, stealing, criminals, criminal act, illegal	stealing, unnatural thing, criminals, people who break laws, illegal	stealing, unnatural thing, criminals, criminal act, people who break laws
Employers	unnatural thing, stealing, government, criminals, attack	stealing, unnatural thing, criminals, people who break laws, illegal	stealing, people who break laws, criminals, unnatural thing, illegal	stealing, criminals, people who break laws, unnatural thing, illegal
Government	unnatural thing, stealing, criminals, attack, government	stealing, unnatural thing, criminals, criminal act, people who break laws	stealing, criminals, people who break laws, unnatural thing, illegal	criminals, stealing, people who break laws, unnatural thing, illegal
Illegal immigrants	quality, justice, america, good feelings, change of location	amnesty, immigrants, deportations, immigrants people who, deport	immigrants, amnesty, immigrants people who, deportations, deport	amnesty, deportations, immigrants, immigrants people who, deport
Local communities	amnesty, immigrants, deportations, immigrants people who, deport	exile, turn away, order, ouster, return home	exile, order, turn away, law, ouster	exile, turn away, ouster, order, return home
Employers	immigrants, amnesty, immigrants people who, deportations, deport	exile, order, turn away, law, ouster	exile, order, law, country, citizen	exile, order, law, country, turn away
Government	amnesty, deportations, immigrants, immigrants people who, deport	exile, turn away, ouster, order, return home	exile, order, law, country, turn away	exile, law, order, turn away, immigration law

Table 13: Agreement (top) and disagreement (bottom) concepts by Stakeholder for Illegal Immigrants.

	School boards	School administrators	Teachers	Parents	Students
School boards	school uniform, uniform, bad, expressing yourself, reason	bad, uniform, school uniform, reason, expressing yourself	bad, uniform, school uniform, reason, expressing yourself	bad, very expensive, fashion, change, kids clothing	bad, very expensive, disguise, touchy about wearing uniforms, kids clothing
School administrators	bad, uniform, school uniform, reason, expressing yourself	bad, uniform, school uniform, reason, expressing yourself	bad, uniform, school uniform, reason, expressing yourself	bad, very expensive, disguise, fashion, change	bad, disguise, touchy about wearing uniforms, very expensive, kids clothing
Teachers	bad, uniform, school uniform, reason, expressing yourself	bad, uniform, school uniform, reason, expressing yourself	bad, uniform, school uniform, reason, expressing yourself	bad, very expensive, disguise, fashion, touchy about wearing uniforms	bad, disguise, touchy about wearing uniforms, pain, special way to dress
Parents	bad, very expensive, fashion, change, kids clothing	bad, very expensive, disguise, fashion, change	bad, very expensive, disguise, fashion, touchy about wearing uniforms	touchy about wearing uniforms, disguise, bad, special outfit, pain	disguise, touchy about wearing uniforms, bad, pain, special outfit
Students	bad, very expensive, disguise, touchy about wearing uniforms, kids clothing	bad, disguise, touchy about wearing uniforms, very expensive, kids clothing	bad, disguise, touchy about wearing uniforms, pain, special way to dress	disguise, touchy about wearing uniforms, bad, pain, special outfit	disguise, touchy about wearing uniforms, school uniform, uniform, required for schools to function effectively
School boards	special way to dress, clothes, kids clothing, clothing, changing appearance	special way to dress, clothing, clothes, changing appearance, kids clothing	special way to dress, clothing, special outfit, clothes, changing appearance	school uniform, uniform, required for schools to function effectively, school, improving image	school uniform, uniform, required for schools to function effectively, school, ideal
School administrators	special way to dress, clothing, clothes, changing appearance, kids clothing	special way to dress, clothing, changing appearance, special outfit, clothes	special outfit, special way to dress, clothing, changing appearance, organization	school uniform, uniform, required for schools to function effectively, school, ideal	uniform, school uniform, required for schools to function effectively, school, ideal
Teachers	special way to dress, clothing, special outfit, clothes, changing appearance	special outfit, special way to dress, clothing, changing appearance, organization	special outfit, special way to dress, clothing, changing appearance, like	school uniform, uniform, required for schools to function effectively, school, ideal	uniform, school uniform, required for schools to function effectively, school, ideal
Parents	school uniform, uniform, required for schools to function effectively, school, improving image	school uniform, uniform, required for schools to function effectively, school, ideal	school uniform, uniform, required for schools to function effectively, school, ideal	expressing yourself, change, dress themselves, motivation, reason	expressing yourself, dress themselves, change, motivation, reason
Students	school uniform, uniform, required for schools to function effectively, school, ideal	uniform, school uniform, required for schools to function effectively, school, ideal	uniform, school uniform, required for schools to function effectively, school, ideal	expressing yourself, dress themselves, change, motivation, reason	change, expressing yourself, fashion, dress themselves, student

Table 14: Agreement (top) and disagreement (bottom) concepts by Stakeholder for School Uniforms.