

# Detecting Attackable Sentences in Arguments

Yohan Jo<sup>1</sup> Seojin Bang<sup>1</sup> Emaad Manzoor<sup>2</sup> Eduard Hovy<sup>1</sup> Chris Reed<sup>3</sup>

<sup>1</sup>School of Computer Science, Carnegie Mellon University, USA

<sup>2</sup>Heinz College of Information Systems and Public Policy, Carnegie Mellon University, USA

<sup>3</sup>Centre for Argument Technology, University of Dundee, UK

{yohanj, seojinb}@cs.cmu.edu, {emaad, hovy}@cmu.edu

c.a.reed@dundee.ac.uk

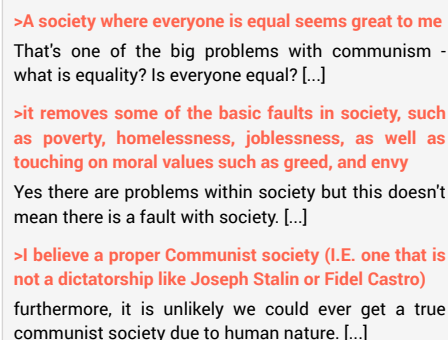
## Abstract

Finding attackable sentences in an argument is the first step toward successful refutation in argumentation. We present a first large-scale analysis of sentence attackability in online arguments. We analyze driving reasons for attacks in argumentation and identify relevant characteristics of sentences. We demonstrate that a sentence’s attackability is associated with many of these characteristics regarding the sentence’s content, proposition types, and tone, and that an external knowledge source can provide useful information about attackability. Building on these findings, we demonstrate that machine learning models can automatically detect attackable sentences in arguments, significantly better than several baselines and comparably well to laypeople.<sup>1</sup>

## 1 Introduction

Effectively refuting an argument is an important skill in persuasion dialogue, and the first step is to find appropriate points to attack in the argument. Prior work in NLP has studied argument quality (Wachsmuth et al., 2017a; Habernal and Gurevych, 2016a) and counterargument generation (Hua et al., 2019; Wachsmuth et al., 2018). But these studies mainly concern an argument’s overall quality and making counterarguments toward the main claim, without investigating what parts of an argument are attackable for successful persuasion. Nevertheless, attacking specific points of an argument is common and effective; in our data of online discussions, challengers who successfully change the original poster’s view are 1.5 times more likely to quote specific sentences of the argument for attacks than unsuccessful challengers (Figure 1). In this paper, we examine how to computationally

<sup>1</sup>Our data and source code are available at: [github.com/yohanjo/emnlp20\\_arg\\_attack](https://github.com/yohanjo/emnlp20_arg_attack)



>A society where everyone is equal seems great to me  
That’s one of the big problems with communism - what is equality? Is everyone equal? [...]  
>it removes some of the basic faults in society, such as poverty, homelessness, joblessness, as well as touching on moral values such as greed, and envy  
Yes there are problems within society but this doesn’t mean there is a fault with society. [...]  
>I believe a proper Communist society (I.E. one that is not a dictatorship like Joseph Stalin or Fidel Castro)  
furthermore, it is unlikely we could ever get a true communist society due to human nature. [...]

Figure 1: A comment to a post entitled “I believe that Communism is not as bad as everyone says”. It quotes and attacks some sentences in the post (red with “>”)

detect attackable sentences in arguments. This attackability information would help people make persuasive refutations and strengthen an argument by solidifying potentially attackable points.

To examine the characteristics of attackable sentences in an argument, we first conduct a qualitative analysis of reasons for attacks in online arguments. Our data comes from discussions in the Change-MyView (CMV) forum on Reddit. In CMV, users challenge the viewpoints of original posters (OPs), and those who succeed receive a  $\Delta$  from the OPs. In this setting, sentences that are attacked and lead to the OP’s view change are considered “attackable”, i.e., targets that are worth attacking. Admittedly, persuasion has to do with “how” to attack as well, but this is beyond the scope of this paper. We only focus on choosing proper sentences to attack, which is a prerequisite for effective persuasion.

This analysis of reasons for attacks, along with argumentation theory and discourse studies, provide insights into what characteristics of sentences are relevant to attackability. Informed by these insights, we extract features that represent relevant sentence characteristics, clustered into four categories: content, external knowledge, proposition

types, and tone. We demonstrate the effects of individual features on sentence attackability, in regard to whether a sentence would be *attacked* and whether a sentence would be attacked *successfully*.

Building on these findings, we examine the efficacy of machine learning models in detecting attackable sentences in arguments. We demonstrate that their decisions match the gold standard significantly better than several baselines and comparably well to laypeople.

To the best of our knowledge, this work is the first large-scale analysis of sentence attackability in arguments. Our contributions are as follows:

- We introduce the problem of detecting attackable sentences in arguments and release the processed data from online discussions and the external knowledge source we used.
- We analyze driving reasons for attacks in arguments and the effects of sentence characteristics on a sentence’s attackability.
- We demonstrate the performance of machine learning models for detecting attackable sentences, setting a baseline for this challenging task and suggesting future directions.

## 2 Background

The strength of an argument is a long-studied topic, dating back to Aristotle (2007), who suggested three aspects of argument persuasiveness: ethos (the arguer’s credibility), logos (logic), and pathos (appeal to the hearer’s emotion). More recently, Wachsmuth et al. (2017b) summarized various aspects of argument quality studied in argumentation theory and NLP, such as clarity, relevance, and arrangement. Some research took empirical approaches and collected argument evaluation criteria from human evaluators (Habernal and Gurevych, 2016a; Wachsmuth et al., 2017a). By adopting some of these aspects, computational models have been proposed to automatically evaluate argument quality in various settings, such as essays (Ke et al., 2019), online comments (Gu et al., 2018), and pairwise ranking (Habernal and Gurevych, 2016b). While these taxonomies help understand and evaluate the quality of an argument as a whole, little empirical analysis has been done in terms of what to attack in an argument to persuade the arguer.

What can be attacked in an argument has been studied more in argumentation theory. Particularly, Walton et al. (2008) present argumentation schemes and critical questions (CQs). Argument schemes

are reasoning types commonly used in daily argumentation. For instance, the scheme of *argument from cause to effect* has the conclusion “*B* will occur” supported by the premise “if *A* occurs, *B* will occur. In this case, *A* occurs”. Each scheme is associated with a set of CQs for judging the argument to be good or fallacious. CQs for the above scheme include “How strong is the causal generalization?” and “Are there other factors that interfere with the causal effect?” Unlike the general argument quality described in the previous paragraph, CQs serve as an evaluation tool that specify local attackable points in an argument. They have been adopted for grading essays (Song et al., 2017) and teaching argumentation skills (Nussbaum et al., 2018). Some of the sentence characteristics in our work are informed by argumentation schemes and CQs.

NLP researchers have widely studied the effectiveness of counterarguments on persuasion (Tan et al., 2016; Cano-Basave and He, 2016; Wei et al., 2016; Wang et al., 2017; Morio et al., 2019) and how to generate counterarguments (Hua et al., 2019; Wachsmuth et al., 2018). Most of the work focuses on the characteristics of counterarguments with respect to topics and styles, without consideration of what points to attack. On the other hand, some studies aimed to model the salience of individual sentences in attacked arguments by paying different degrees of attention to sentences using attention mechanism (Jo et al., 2018; Ji et al., 2018). While their approaches helped to predict the success of persuasion, it was difficult to interpret what constitute the salience or attackability of sentences. To address this limitation, we quantify and analyze the characteristics of sentences that are attacked and lead to the arguer’s view change.

## 3 Data

Here we describe how we collected and labeled our data.

### 3.1 Data Collection

We use online discussions from the Change-MyView (CMV) subreddit<sup>2</sup>. In this forum, users post their views on various issues and invite other users to challenge their views. If a comment changes the original poster (OP)’s view, the OP acknowledges it by replying to the comment with a  $\Delta$  symbol. The high quality of the discussions in this forum is maintained through several mod-

<sup>2</sup><https://www.reddit.com/r/changemyview>

eration rules, such as the minimum length of an original post and the maximum response time of OPs. As a result, CMV discussions have been used in many NLP studies (Chakrabarty et al., 2019; Morio et al., 2019; Jo et al., 2018; Musi, 2017; Wei et al., 2016; Tan et al., 2016).

We scraped CMV posts and comments written between January 1, 2014 and September 30, 2019, using the Pushshift API. We split them into a dev set (Jan 2014–Jan 2018 for training and Feb 2018–Nov 2018 for validation) and a test set (Dec 2018–Sep 2019), with the ratio of 6:2:2. We split the data by time to measure our models’ generality to unseen subjects.

As the characteristics of arguments vary across different issues, we categorized the posts into domains using LDA. For each post, we chose as its domain the topic that has the highest standard score; topics comprising common words were excluded. We tried different numbers of topics (25, 30, 35, 40) and finalized on 40, as it achieves the lowest perplexity. This process resulted in 30 domains (excluding common-word topics): media, abortion, sex, election, Reddit, human economy, gender, race, family, life, crime, relationship, movie, world, game, tax, law, money, drug, war, religion, job, food, power, school, college, music, gun, and Jewish (from most frequent to least, ranging 5%–2%).

### 3.2 Labeling Attackability

Since we are interested in which parts of a post are attacked by comments and whether the attacks lead to successful view changes, our goal here is to label each sentence in a post as *successfully attacked*, *unsuccessfully attacked*, or *unattacked*. We only consider comments directly replying to each post (top-level comments), as lower-level comments usually address the same points as their parent comments (as will be validated at the end of the section).

**Attacked vs. Unattacked:** Some comments use direct quotes with the > symbol to address specific sentences of the post (Figure 1). Each quote is matched with the longest sequence of sentences in the post using the Levenshtein edit distance (allowing a distance of 2 characters for typos). A matched text span should contain at least one word and four characters, and cover at least 80% of the quote to exclude cases where the > symbol is used to quote external content. As a result, 98% of the matched spans cover the corresponding quotes entirely. Additionally, a sentence in the post is considered to

be quoted if at least four non-stopwords appear in a comment’s sentence. For example:

Post: ... *If you do something, you should be prepared to accept the consequences.* ...

Comment: ... *I guess my point is, even if you do believe that “If you do something, you should be prepared to accept the consequences,” you can still feel bad for the victims.* ...

We considered manually annotating attacked sentences too, but it turned out to be extremely time-consuming and subjective (Appendix A). We tried to automate it using heuristics (word overlap and vector embeddings), but precision severely deteriorated. As we value the precision of labels over recall, we only use the method described in the previous paragraph. Chakrabarty et al. (2019) used the same method to collect attack relations in CMV.

**Successfully vs. Unsuccessfully Attacked:** After each sentence in a post is labeled as attacked or not, each attacked sentence is further labeled as *successfully attacked* if any of the comments that attack it, or their lower-level comments win a  $\Delta$ .

We post-process the resulting labels to increase their validity. First, as a challenger and the OP have discussion down the comment thread, the challenger might attack different sentences than the originally attacked ones and change the OP’s view. In this case, it is ambiguous which sentences contribute to the view change. Hence, we extract quotes from all lower-level comments of  $\Delta$ -winning challengers, and if any of the quotes attack new sentences, this challenger’s attacks are excluded from the labeling of *successfully attacked*. This case is not common, however (0.2%).

Second, if a comment attacks many sentences in the post and change the OP’s view, some of them may not contribute to the view change but are still labeled as *successfully attacked*. To reduce this noise, comments that have more than three quotes are excluded from the labeling of *successfully attacked*<sup>3</sup>. This amounts to 12% of top-level comments (63% of comments have only one quote, 17% two quotes, and 8% three quotes).

Lastly, we verified if quoted sentences are actually attacked. We randomly selected 500 comments and checked if each quoted sentence is purely agreed with without any opposition, challenge, or question. This case was rare (0.4%)<sup>4</sup>, so we do

<sup>3</sup>This allows our subsequent analyses to capture stronger signals for successful attacks than without this process.

<sup>4</sup>Further, this case happened in only one out of the 500 comments (0.2%), where the author agreed with 4 quoted sentences. In CMV, challengers do use concessions but hardly

Dataset		Train	Val	Test
Attacked	#posts	25,839	8,763	8,558
	#sentences	420,545	133,090	134,375
	#attacked	119,254	40,163	40,354
Successful	#posts	3,785	1,235	1,064
	#sentences	66,628	20,240	17,129
	#successful	8,746	2,718	2,288

Table 1: Data statistics. “Attacked” contains posts with at least one attacked sentence. “Successful” contains posts with at least one successfully attacked sentence.

R1	<i>S</i> is true but does not support the main claim (19%)
R2	<i>S</i> misses cases suggesting opposite judgment (18%)
R3	<i>S</i> has exceptions (17%)
R4	<i>S</i> is false (12%)
R5	<i>S</i> misses nuanced distinctions of a concept (8%)
R6	<i>S</i> is unlikely to happen (6%)
R7	<i>S</i> has no evidence (6%)
R8	<i>S</i> uses an invalid assumption or hypothetical (4%)
R9	<i>S</i> contradicts statements in the argument (4%)
R10	Other (4%)

(a) Rationales for attacking a sentence (*S*).

F1	Personal opinion (28%)
F2	Invalid hypothetical (26%)
F3	Invalid generalization (13%)
F4	No evidence (11%)
F5	Absolute statement (7%)
F6	Concession (5%)
F7	Restrictive qualifier (5%)
F8	Other (5%)

(b) Motivating factors for attacks.

Table 2: Rationales and motivating factors for attacks.

not further process this case. Table 1 shows some statistics of the final data.

## 4 Quantifying Sentence Characteristics

As the first step for analyzing the characteristics of attackable sentences, we examine driving reasons for attacks and quantify relevant characteristics.

### 4.1 Rationales and Motivation for Attacks

To analyze rationales for attacks, two authors examined quotes and rebuttals in the training data (one successful and one unsuccessful comment for each post). From 156 attacks, we identified 10 main rationales (Table 2a), which are finer-grained than the refutation reasons in prior work (Wei et al., 2016). The most common rationale is that the sentence is factually correct but is irrelevant to the main claim (19%). Counterexample-related rationales are also common: the sentence misses an example suggest-

quote the OP’s sentences just to agree.

ing the opposite judgment to the sentence’s own (18%) and the sentence has exceptions (17%).

This analysis is based on polished rebuttals, which mostly emphasize logical aspects, and cannot fully capture other factors that motivate attacks. Hence, we conducted a complementary analysis, where an undergraduate student chose three sentences to attack for each of 50 posts and specified the reasons in their own terms (Table 2b). The most common factor is that the sentence is only a personal opinion (28%). Invalid hypotheticals are also a common factor (26%). The tone of a sentence motivates attacks as well, such as generalization (13%), absoluteness (7%), and concession (5%).

### 4.2 Feature Extraction

Based on these analyses, we cluster various sentence characteristics into four categories—content, external knowledge, proposition types, and tone.<sup>5</sup>

#### 4.2.1 Content

Content and logic play the most important role in CMV discussions. We extract the content of each sentence at two levels: TFIDF-weighted *n*-grams ( $n = 1, 2, 3$ ) and sentence-level **topics**. Each sentence is assigned one topic using Sentence LDA (Jo and Oh, 2011). We train a model on posts in the training set and apply it to all posts, exploring the number of topics  $\in \{10, 50, 100\}$ .<sup>6</sup>

#### 4.2.2 External Knowledge

External knowledge sources may provide information as to how truthful or convincing a sentence is (e.g., Table 2a-R2, R3, R4, R7 and Table 2b-F4). As our knowledge source, we use kialo.com—a collaborative argument platform over more than 1.4K issues. Each issue has a main statement, and users can respond to any existing statement with pro/con statements (1-2 sentences), building an argumentation tree. Kialo has advantages over structured knowledge bases and Wikipedia in that it includes many debatable statements; many attacked sentences are subjective judgments (§4.1), so fact-based knowledge sources may have limited utility. In addition, each statement in Kialo has pro/con counts, which may reflect the convincingness of the statement. We scraped 1,417 argumentation trees and 130K statements (written until Oct 2019).

<sup>5</sup>Some rationales in Table 2a (e.g., R1 and R9) are difficult to operationalize reliably using the current NLP technology and thus are not included in our features.

<sup>6</sup>We also tried features based on semantic frames using SLING (Ringgaard et al., 2017), but they were not helpful.



For each sentence in CMV, we retrieve similar statements in Kialo that have at least 5 common words<sup>7</sup> and compute the following three features. **Frequency** is the number of retrieved statements; sentences that are not suitable for argumentation are unlikely to appear in Kialo. This feature is computed as  $\log_2(N + 1)$ , where  $N$  is the number of retrieved statements. **Attractiveness** is the average number of responses for the matched statements, reflecting how debatable the sentence is. It is computed as  $\log_2(M + 1)$ , where  $M = \frac{1}{N} \sum_{i=1}^N R_i$  and  $R_i$  is the number of responses for the  $i$ th retrieved statement. Lastly, **extremeness** is  $\frac{1}{N} \sum_{i=1}^N |P_i - N_i|$ , where  $P_i$  and  $N_i$  are the proportions (between 0 and 1) of pro responses and con responses for the  $i$ th retrieved statement. A sentence that most people would see flawed would have a high extremeness value.

### 4.2.3 Proposition Types

Sentences convey different types of propositions, such as predictions and hypotheticals. No proposition types are fallacious by nature, but some of them may make it harder to generate a sound argument. They also communicate different moods, causing the hearer to react differently. We extract 13 binary features for proposition types. They are all based on lexicons and regular expressions, which are available in Appendix C.

**Questions** express the intent of information seeking. Depending on the form, we define three features: **confusion** (e.g., *I don't understand*), **why/how** (e.g., *why ...?*), and **other**.

**Normative** sentences suggest that an action be carried out. Due to their imperative mood, they can sound face-threatening and thus attract attacks.

**Prediction** sentences predict a future event. They can be attacked with reasons why the prediction is unlikely (Table 2a-R6), as in critical questions for *argument from cause to effect* (Walton et al., 2008).

**Hypothetical** sentences may make implausible assumptions (Table 2a-R8 and Table 2b-F2) or restrict the applicability of the argument too much (Table 2b-F7).

**Citation** often strengthens a claim using authority, but the credibility of the source could be attacked (Walton et al., 2008).

**Comparison** may reflect personal preferences that are vulnerable to attacks (Table 2b-F1).

<sup>7</sup>Similarity measures based on word embeddings and knowledge representation did not help (Appendix B).

**Examples** in a sentence may be attacked for their invalidity (Walton et al., 2008) or counterexamples (Table 2a-R3).

**Definitions** form a ground for arguments, and challengers could undermine an argument by attacking this basis (e.g., Table 2a-R5).

**Personal stories** are the arguer's experiences, whose validity is difficult to refute. A sentence with a personal story has subject *I* and a non-epistemic verb; or it has *my* modifying non-epistemic nouns.

**Inclusive sentences** that mention *you* and *we* engage the hearer into the discourse (Hyland, 2005), making the argument more vulnerable to attacks.

### 4.2.4 Tone

Challengers are influenced by the tone of an argument, e.g., subjectiveness, absoluteness, or confidence (Table 2b). We extract 8 features for the tone of sentences.

**Subjectivity** comprises judgments, which are often attacked due to counterexamples (Table 2a-R2) or their arbitrariness (Table 2b-F1, Walton et al. (2008)). The subjectivity of a sentence is the average subjectivity score of words based on the Subjectivity Lexicon (Wilson et al., 2005) (non-neutral words of "weaksbj" = 0.5 and "strongsbj" = 1).

**Concreteness** is the inverse of abstract diction, whose meaning depends on subjective perceptions and experiences. The concreteness of a sentence is the sum of the standardized word scores based on Brysbaert et al. (2014)'s concreteness lexicon.

**Qualification** expresses the level of generality of a claim, where absolute statements can motivate attacks (Table 2b-R3). The qualification score of a sentence is the average word score based on our lexicon of qualifiers and generality words.

**Hedging** can sound unconvincing (Durik et al., 2008) and motivate attacks. A sentence's hedging score is the sum of word scores based on our lexicon of downtoners and boosters.

**Sentiment** represents the valence of a sentence. Polar judgments may attract more attacks than neutral statements. We calculate the sentiment of each sentence with BERT (Devlin et al., 2018) trained on the data of SemEval 2017 Task 4 (Rosenthal et al., 2017). **Sentiment score** is a continuous value ranging between -1 (negative) and +1 (positive), and **sentiment categories** are nominal (positive, neutral, and negative)<sup>8</sup>. In addition, we compute the

<sup>8</sup>We achieved an average recall of 0.705, which is higher than the winner team's performance of 0.681.

scores of **arousal** (intensity) and **dominance** (control) as the sum of the standardized word scores based on Warriner et al. (2013)’s lexicon.

## 5 Task 1: Attackability Characteristics

One of our goals in this paper is to analyze what characteristics of sentences are associated with a sentence’s attackability. Hence, in this section, we measure the effect size and statistical significance of each feature toward two labels: (i) whether a sentence is attacked or not, using the dev set of the “Attacked” dataset ( $N=553,635$ ), (ii) whether a sentence is attacked successfully or unsuccessfully, using all attacked sentences ( $N=159,417$ ).<sup>9</sup> Since the effects of characteristics may depend on the issue being discussed, the effect of each feature is estimated conditioned on the domain of each post using a logistic regression, and the statistical significance of the effect is assessed using the Wald test. For interpretation purposes, we use *odds ratio* (OR)—the exponent of the effect size.<sup>10</sup>

### 5.1 Content

Attacked sentences tend to mention big issues like gender, race, and health as revealed in topics 47, 8, and 6 (Table 3) and  $n$ -grams *life*, *weapons*, *women*, *society*, and *men* (Table 7 in Appendix E). These issues are also positively correlated with successful attacks. On the other hand, mentioning relatively personal issues (*tv*, *friends*, topic 38) seems negatively correlated with successful attacks. So do forum-specific messages (*cmv*, *thank*, topic 4).

Attacking seemingly evidenced sentences appears to be effective for persuasion when properly done. Successfully attacked sentences are likely to mention specific data (*data*, *%*) and be the OP’s specific reasons under bullet points (2. and 3.).

$n$ -grams capture various characteristics that are vulnerable to attacks, such as uncertainty and absoluteness (*i believe*, *never*), hypotheticals (*if i*), questions (*?*, *why*), and norms (*should*).

<sup>9</sup>Simply measuring the predictive power of features in a prediction setting provides an incomplete picture of the roles of the characteristics. Some features may not have drastic contribution to prediction due to their infrequency, although they may have significant effects on attackability.

<sup>10</sup>Odds are the ratio of the probability of a sentence being (successfully) attacked to the probability of being not (successfully) attacked; OR is the ratio of odds when the value of the characteristic increases by one unit (Appendix D).

	Feature	Attacked	Successful
Content	Topic47: Gender <sup>†</sup>	1.37 (***)	1.34 (***)
	Topic8: Race <sup>†</sup>	1.19 (***)	1.21 ( ** )
	Topic6: Food <sup>†</sup>	1.00 ( )	1.39 (***)
	Topic38: Movie & Show <sup>†</sup>	1.03 ( )	0.78 (***)
	Topic4: CMV-Specific <sup>†</sup>	0.16 (***)	0.36 ( ** )
Knowledge	Kialo Frequency (log2)	1.18 (***)	1.07 (***)
	Kialo Attractiveness (log2)	1.30 (***)	1.18 (***)
	Kialo Extremeness	1.51 (***)	1.19 (***)
Proposition Types	Question - Confusion <sup>†</sup>	0.97 ( )	1.29 ( * )
	Question - Why/How <sup>†</sup>	1.77 (***)	1.27 (***)
	Question - Other <sup>†</sup>	1.16 (***)	1.11 ( * )
	Citation <sup>†</sup>	0.53 (***)	1.17 ( * )
	Definition <sup>†</sup>	1.04 ( )	1.32 ( ** )
	Normative <sup>†</sup>	1.26 (***)	1.10 ( ** )
	Prediction <sup>†</sup>	1.22 (***)	1.02 ( )
	Hypothetical <sup>†</sup>	1.29 (***)	1.07 ( )
	Comparison <sup>†</sup>	1.25 (***)	1.02 ( )
	Example <sup>†</sup>	1.20 (***)	1.17 ( * )
	Personal Story <sup>†</sup>	0.70 (***)	1.09 ( ** )
	Use of <i>You</i> <sup>†</sup>	1.18 (***)	1.04 ( )
	Use of <i>We</i> <sup>†</sup>	1.24 (***)	0.98 ( )
Tone	Subjectivity <sup>‡</sup>	1.03 (***)	0.97 (***)
	Concreteness <sup>‡</sup>	0.87 (***)	0.92 (***)
	Hedges <sup>‡</sup>	1.04 (***)	1.06 (***)
	Quantification <sup>‡</sup>	0.97 (***)	1.02 ( )
	Sentiment Score <sup>‡</sup>	0.87 (***)	1.00 ( )
	Sentiment: Positive <sup>†</sup>	0.76 (***)	0.99 ( )
	Sentiment: Neutral <sup>†</sup>	0.82 (***)	1.00 ( )
	Sentiment: Negative <sup>†</sup>	1.34 (***)	1.00 ( )
	Arousal <sup>‡</sup>	1.02 (***)	0.95 (***)
Dominance <sup>‡</sup>	1.07 (***)	1.08 (***)	

Table 3: Odds ratio (OR) and statistical significance of features. An effect is positive (blue) if  $OR > 1$  and negative (red) if  $OR < 1$ . (<sup>†</sup>: binary, <sup>‡</sup>: standardized / \*:  $p < 0.05$ , \*\*:  $p < 0.01$ , \*\*\*:  $p < 0.001$ )

### 5.2 External Knowledge

The Kialo-based knowledge features provide significant information about whether a sentence would be attacked successfully (Table 3). As the frequency of matched statements in Kialo increases twice, the odds for successful attack increase by 7%. As an example, the following attacked sentence has 18 matched statements in Kialo.

*I feel like it is a parents right and responsibility to make important decisions for their child.*

The attractiveness feature has a stronger effect; as matched statements have twice more responses, the odds for successful attack increase by 18%, probably due to higher debatability.

A sentence being completely extreme (i.e., the matched sentences have only pro or con responses) increases the odds for successful attack by 19%.

As expected, the argumentative nature of Kialo

allows its statements to match many subjective sentences in CMV and serves as an effective information source for a sentence’s attackability.

### 5.3 Proposition Types

Questions, especially why/how, are effective targets for successful attack (Table 3). Although challengers do not pay special attention to expressions of confusion (see column “Attacked”), they are positively correlated with successful attack (OR=1.29).

Citations are often used to back up an argument and have a low chance of being attacked, reducing the odds by half. However, properly attacking citations significantly increases the odds for successful attack by 17%. Similarly, personal stories have a low chance of being attacked and definitions do not attract challengers’ attacks, but attacking them is found to be effective for successful persuasion.

All other features for proposition types have significantly positive effects on being attacked (OR=1.18–1.29), but only normative and example sentences are correlated with successful attack.

### 5.4 Tone

Successfully attacked sentences tend to have lower subjectivity and arousal (Table 3), in line with the previous observation that they are more data- and reference-based than unsuccessfully attacked sentences. In contrast, sentences about concrete concepts are found to be less attackable.

Uncertainty (high hedging) and absoluteness (low qualification) both increase the chance of attacks, which aligns with the motivating factors for attacks (Table 2b), while only hedges are positively correlated with successful attacks, implying the importance of addressing the arguer’s uncertainty.

Negative sentences with high arousal and dominance have a high chance of being attacked, but most of these characteristics have either no or negative effects on successful attacks.

### 5.5 Discussion

We have found some evidence that, somewhat counter-intuitively, seemingly evidenced sentences are more effective to attack. Such sentences use specific data (*data*, %), citations, and definitions. Although attacking these sentences may require even stronger evidence and deeper knowledge, arguers seem to change their viewpoints when a fact they believe with evidence is undermined. In addition, it seems very important and effective to

identify and address what the arguer is confused (confusion) or uncertain (hedges) about.

Our analysis also reveals some discrepancies between the characteristics of sentences that challengers commonly think are attackable and those that are indeed attackable. Challengers are often attracted to subjective and negative sentences with high arousal, but successfully attacked sentences have rather lower subjectivity and arousal, and have no difference in negativity compared to unsuccessfully attacked sentences. Furthermore, challengers pay less attention to personal stories, while successful attacks address personal stories more often.

## 6 Task 2: Attackability Prediction

Now we examine how well computational models can detect attackable sentences in arguments.

### 6.1 Problem Formulation

This task is cast as ranking sentences in each post by their attackability scores predicted by a regression model. We consider two types of attackability: (i) whether a sentence will be attacked or not, (ii) whether a sentence will be successfully attacked or not (attacked unsuccessfully + unattacked). For both settings, we consider posts that have at least one sentence with the positive label (Table 1).

We use three evaluation metrics. **P@1** is the precision of the first ranked sentence, measuring the model’s accuracy when choosing one sentence to attack for each post. Less strictly, **A@3** gives a score of 1 if any of the top 3 sentences is a positive instance and 0 otherwise. **AUC** measures individual sentence-level accuracy—how likely positive sentences are assigned higher probabilities.

### 6.2 Comparison Models

For machine learning models, we explore two logistic regression models to compute the probability of the positive label for each sentence, which becomes the sentence’s *attackability score*. **LR** is a basic logistic regression with our features<sup>11</sup> (Section 4) and binary variables for domains. We explored feature selection using L1-norm and regularization using L2-norm.<sup>12</sup> **BERT** is logistic regression where our features are replaced with the BERT embedding of the input sentence (Devlin et al., 2018). Contextualized BERT embeddings have achieved

<sup>11</sup>We tried the number of topics  $\in \{10, 50, 100\}$ , and 50 has the best AUC on the val set for both prediction settings.

<sup>12</sup>We also tried a multilayer perceptron to model feature interactions, but it consistently performed worse than LR.

	Attacked			Successful		
	P@1	A@3	AUC	P@1	A@3	AUC
Random	35.9	66.0	50.1	18.9	45.0	50.1
Length	42.9	73.7	54.5	22.3	52.1	55.7
LR	47.1	76.2	61.7	24.2	54.5	59.3
(×) Content	45.2	74.4	58.1	24.0	52.6	57.0
(×) Knowledge	47.0	76.0	61.7	24.1	54.3	59.0
(×) Prop Type	46.7	75.9	61.5	24.4	53.6	59.0
(×) Tone	47.0	76.0	61.9	25.2	56.2	59.4
BERT	49.6	77.8	64.4	28.3	57.2	62.0
Humans <sup>†</sup>	51.7	80.1	–	27.8	54.2	–

Table 4: Prediction accuracy. All LR/BERT scores (rows 3–8) have standard deviations between 0.1 and 1.1, significantly outperforming “Length”. <sup>†</sup>The average bootstrap accuracy after resampling 100K times with sample size 200—the standard deviations of P@1 and A@3 range between 2.1 and 3.5.

state-of-the-art performance in many NLP tasks. We use the pretrained, uncased base model from Hugging Face (Wolf et al., 2019) and fine-tune it during training.<sup>13</sup>

We explore two baseline models. **Random** is to rank sentences randomly. **Length** is to rank sentences from longest to shortest, with the intuition that longer sentences may contain more information and thus more content to attack as well.

Lastly, we estimate laypeople’s performance on this task. Three undergraduate students each read 100 posts and rank three sentences to attack for each post. Posts that have at least one positive instance are randomly selected from the test set.<sup>14</sup>

### 6.3 Results

All computational models were run 10 times, and their average accuracy is reported in Table 4. Both the LR and BERT models significantly outperform the baselines, while the BERT model performs best. For predicting attacked sentences, the BERT model’s top 1 decisions match the gold standard 50% of the time; its decisions match 78% of the time when three sentences are chosen. Predicting successfully attacked sentences is harder, but the performance gap between our models and the baselines gets larger. The BERT model’s top 1 decisions match the gold standard 28% of the time—a 27% and 10% boost from random and length-based performance, respectively.

<sup>13</sup>Details for reproducibility are in Appendix F.

<sup>14</sup>We were interested in the performance of young adults who are academically active and have a moderate level of life experience. Their performance may not represent the general population, though.

To examine the contribution of each feature category, we did ablation tests based on the best performing LR model (Table 4 rows 4–7). The two prediction settings show similar tendencies. Regarding P@1 for successful attack, content has the highest contribution, followed by knowledge, proposition types, and tone. This result reaffirms the importance of content for a sentence’s attackability. But the other features still have significant contribution, yielding higher P@1 and AUC (Table 4 row 4) than the baselines.

It is worth noting that our features, despite the lower accuracy than the BERT model, are clearly informative of attackability prediction as Table 4 row 3 shows. Moreover, since they directly operationalize the sentence characteristics we compiled, it is pretty transparent that they capture relevant information that contributes to sentence attackability and help us better understand what characteristics have positive and negative signals for sentence attackability. We speculate that transformer models like BERT are capable of encoding these characteristics more sophisticatedly and may include some additional information, e.g., lexical patterns, leading to higher accuracy. But at the same time, it is less clear exactly what they capture and whether they capture relevant information or irrelevant statistics, as is often the case in computational argumentation (Niven and Kao, 2019).

Figure 2 illustrates how LR allows us to interpret the contribution of different features to attackability, by visualizing a post with important features highlighted. For instance, external knowledge plays a crucial role in this post; all successfully attacked sentences match substantially more Kialo statements than other sentences. The attackability scores of these sentences are also increased by the use of hypotheticals and certain *n*-grams like *could*. These features align well with the actual attacks by successful challengers. For instance, they pointed out that the expulsion of Russian diplomats (sentence 2) is not an aggressive reaction because the diplomats can be simply replaced with new ones. Kialo has a discussion on the relationship between the U.S. and Russia, and one statement puts forward exactly the same point that the expulsion was a forceful-looking but indeed a nice gesture. Similarly, a successful challenger pointed out the consistent attitude of the U.S. toward regime change in North Korea (sentence 3), and the North Korean regime is a controversial topic in Kialo. Lastly,



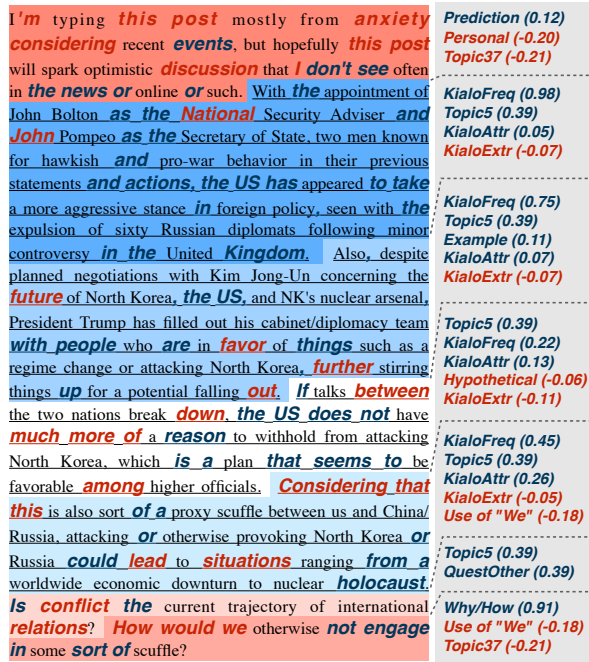


Figure 2: Prediction visualization. Background color indicates predicted attackability (blue: high, red: low). Successfully attacked sentences are underlined. Features with high/low weights are indicated with blue/red.

one successful challenger attacked the hypothetical outcomes in sentences 4 and 5, pointing out that those outcomes are not plausible, and the LR model also captures the use of hypothetical and the word *could* as highly indicative of attackability. More successful and erroneous cases are in Appendix H.

Laypeople perform significantly better than the BERT model for predicting attacked sentences, but only comparably well for successfully attacked sentences (Table 4 row 9). Persuasive argumentation in CMV requires substantial domain knowledge, but laypeople do not have such expertise for many domains. The BERT model, however, seems to take advantage of the large data and encodes useful linguistic patterns that are predictive of attackability. A similar tendency has been observed in predicting persuasive refutation (Guo et al., 2020), where a machine-learned model outperformed laypeople. Nevertheless, in our task, the humans and the BERT model seem to make similar decisions; the association between their choices of sentences is high, with odds ratios ranging between 3.43 (top 1) and 3.33 (top 3). Interestingly, the LR model has a low association with the human decisions for top 1 (OR=2.65), but the association exceeds the BERT model for top 3 (OR=3.69). It would be interesting to further examine the similarities and differences

in how humans and machines choose sentences to attack.

## 7 Conclusion

We studied how to detect attackable sentences in arguments for successful persuasion. Using on-line arguments, we demonstrated that a sentence’s attackability is associated with many of its characteristics regarding its content, proposition types, and tone, and that Kialo provides useful information about attackability. Based on these findings we demonstrated that machine learning models can automatically detect attackable sentences, comparably well to laypeople.

Our work contributes a new application to the growing literature on causal inference from text (Egami et al., 2018), in the setting of “text as a treatment”. Specifically, our findings in Section 5 pave the way towards answering the causal question: would attacking a certain type of sentence (e.g., questions or expressions of confusion) in an argument increase the probability of persuading the opinion holder? While our findings suggest initial hypotheses about the characteristics of sentences that can be successfully attacked, establishing causality in a credible manner would require addressing confounders, such as the challenger’s reputation (Manzoor et al., 2020) and persuasive skill reflected in their attack (Tan et al., 2014). We leave this analysis to future work.

Our work could be improved also by including discourse properties (coherence, cohesiveness). Further, argumentation structure (support relations between sentences or lack thereof) might provide useful information about each sentence’s attackability.

## Acknowledgments

This research was supported by the Kwanjeong Educational Foundation.

## References

Aristotle. 2007. *On Rhetoric*. A Theory of Civic Discourse (Translated by George Alexander Kennedy). Oxford University Press, USA.

Marc Brysbaert, Amy Beth Warriner, and Victor Kuperman. 2014. Concreteness ratings for 40 thousand generally known English word lemmas. *Behavior Research Methods*, 46(3):904–911.

- Amparo Elizabeth Cano-Basave and Yulan He. 2016. A Study of the Impact of Persuasive Argumentation in Political Debates. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1405–1413, San Diego, California. Association for Computational Linguistics.
- Tuhin Chakrabarty, Christopher Hidey, Smaranda Muresan, Kathy McKeown, and Alyssa Hwang. 2019. AMPERSAND: Argument Mining for PERSuasive oNline Discussions. 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 2926–2936, Hong Kong, China. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. [BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding](#). *arXiv*.
- Amanda M Durik, M Anne Britt, Rebecca Reynolds, and Jennifer Storey. 2008. The Effects of Hedges in Persuasive Arguments: A Nuanced Analysis of Language. *Journal of Language and Social Psychology*, 27(3):217–234.
- Naoki Egami, Christian J Fong, Justin Grimmer, Margaret E Roberts, and Brandon M Stewart. 2018. How to make causal inferences using texts. *arXiv preprint arXiv:1802.02163*.
- Yunfan Gu, Zhongyu Wei, Maoran Xu, Hao Fu, Yang Liu, and Xuan-Jing Huang. 2018. Incorporating Topic Aspects for Online Comment Convincingness Evaluation. *Proceedings of the 5th Workshop on Argument Mining*, pages 97–104.
- Zhen Guo, Zhe Zhang, and Munindar Singh. 2020. [In Opinion Holders’ Shoes: Modeling Cumulative Influence for View Change in Online Argumentation](#). In *Proceedings of The Web Conference 2020*, WWW, page 2388–2399, New York, NY, USA. Association for Computing Machinery.
- Ivan Habernal and Iryna Gurevych. 2016a. 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.
- Ivan Habernal and Iryna Gurevych. 2016b. Which argument is more convincing? Analyzing and predicting convincingness of Web arguments using bidirectional LSTM. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1589–1599, Berlin, Germany. Association for Computational Linguistics.
- Xinyu Hua, Zhe Hu, and Lu Wang. 2019. [Argument Generation with Retrieval, Planning, and Realization](#). *arXiv*.
- Ken Hyland. 2005. *Metadiscourse : Exploring Interaction in Writing*. Continuum Discourse Series. Continuum.
- Lu Ji, Zhongyu Wei, Xiangkun Hu, Yang Liu, Qi Zhang, and Xuanjing Huang. 2018. Incorporating Argument-Level Interactions for Persuasion Comments Evaluation using Co-attention Model. In *Proceedings of COLING , the th International Conference on Computational Linguistics*, pages 1–12.
- Yohan Jo and Alice H Oh. 2011. Aspect and Sentiment Unification Model for Online Review Analysis. In *Proceedings of the fourth ACM international conference on Web search and data mining*, pages 815–824.
- Yohan Jo, Shivani Poddar, Byungsoo Jeon, Qinlan Shen, Carolyn Penstein Rosé, and Graham Neubig. 2018. Attentive Interaction Model: Modeling Changes in View in Argumentation. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 103–116, New Orleans, Louisiana. Association for Computational Linguistics.
- Zixuan Ke, Hrishikesh Inamdar, Hui Lin, and Vincent Ng. 2019. Give Me More Feedback II: Annotating Thesis Strength and Related Attributes in Student Essays. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3994–4004, Florence, Italy. Association for Computational Linguistics.
- Emaad Manzoor, George H. Chen, Dokyun Lee, and Michael D. Smith. 2020. Influence via ethos: On the persuasive power of reputation in deliberation online. *arXiv preprint arXiv:2006.00707*.
- Gaku Morio, Ryo Egawa, and Katsuhide Fujita. 2019. Revealing and Predicting Online Persuasion Strategy with Elementary Units. 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 6273–6278, Hong Kong, China. Association for Computational Linguistics.
- Elena Musi. 2017. How did you change my view? A corpus-based study of concessions’ argumentative role. *Discourse Studies*, 20(2):270–288.
- Timothy Niven and Hung-Yu Kao. 2019. [Probing Neural Network Comprehension of Natural Language Arguments](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, ACL, pages 4658 – 4664.

- E Michael Nussbaum, Ian J Dove, Nathan Slife, CarolAnne M Kardash, Refika Turgut, and David Vallett. 2018. Using critical questions to evaluate written and oral arguments in an undergraduate general education seminar: a quasi-experimental study. *Reading and Writing*, 19(2):1–22.
- Nils Reimers, Benjamin Schiller, Tilman Beck, Johannes Daxenberger, Christian Stab, and Iryna Gurevych. 2019. Classification and Clustering of Arguments with Contextualized Word Embeddings. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 567–578, Florence, Italy. Association for Computational Linguistics.
- Michael Ringgaard, Rahul Gupta, and Fernando C. N. Pereira. 2017. *SLING: A framework for frame semantic parsing*. *CoRR*, abs/1710.07032.
- Sara Rosenthal, Noura Farra, and Preslav Nakov. 2017. SemEval-2017 Task 4: Sentiment Analysis in Twitter. In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, pages 502–518, Vancouver, Canada. Association for Computational Linguistics.
- Yi Song, Paul Deane, and Beata Beigman Klebanov. 2017. Toward the Automated Scoring of Written Arguments: Developing an Innovative Approach for Annotation. *ETS Research Report . . .*, 152(3):157.
- Chenhao Tan, Lillian Lee, and Bo Pang. 2014. The effect of wording on message propagation: Topic- and author-controlled natural experiments on Twitter. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 175–185, Baltimore, Maryland. Association for Computational Linguistics.
- Chenhao Tan, Vlad Niculae, Cristian Danescu-Niculescu-Mizil, and Lillian Lee. 2016. Winning Arguments: Interaction Dynamics and Persuasion Strategies in Good-faith Online Discussions. In *Proceedings of the 25th International Conference on World Wide Web*, pages 613–624. International World Wide Web Conferences Steering Committee.
- URL1. [http://www-di.inf.puc-rio.br/~endler/students/Hedging\\_Handout.pdf](http://www-di.inf.puc-rio.br/~endler/students/Hedging_Handout.pdf).
- URL2. <https://github.com/words/hedges/blob/master/data.txt>.
- URL3. <https://www.janefriedman.com/hedge-word-inflation-words-prune/>.
- Henning Wachsmuth, Nona Naderi, Ivan Habernal, Yufang Hou, Graeme Hirst, Iryna Gurevych, and Benno Stein. 2017a. Argumentation Quality Assessment: Theory vs. Practice. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 250–255, Vancouver, Canada. Association for Computational Linguistics.
- Henning Wachsmuth, Nona Naderi, Yufang Hou, Yonatan Bilu, Vinodkumar Prabhakaran, Tim Alberdingk Thijm, Graeme Hirst, and Benno Stein. 2017b. Computational Argumentation Quality Assessment in Natural Language. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pages 176–187, Valencia, Spain. Association for Computational Linguistics.
- 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. Association for Computational Linguistics.
- Douglas Walton, Chris Reed, and Fabrizio Macagno. 2008. *Argumentation Schemes*. Cambridge University Press.
- 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 Association of Computational Linguistics*, 5:219–232.
- Amy Beth Warriner, Victor Kuperman, and Marc Brysbaert. 2013. Norms of valence, arousal, and dominance for 13,915 English lemmas. *Behavior Research Methods*, 45(4):1191–1207.
- Zhongyu Wei, Yang Liu, and Yi Li. 2016. Is This Post Persuasive? Ranking Argumentative Comments in Online Forum. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 195–200, Berlin, Germany. Association for Computational Linguistics.
- Theresa Wilson, Janyce Wiebe, and Paul Hoffmann. 2005. Recognizing Contextual Polarity in Phrase-Level Sentiment Analysis. In *Proceedings of Human Language Technology Conference and Conference on Empirical Methods in Natural Language Processing*, pages 347–354, Vancouver, British Columbia, Canada. Association for Computational Linguistics.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2019. *HuggingFace’s Transformers: State-of-the-art Natural Language Processing*. *arXiv*.



## Appendices

### A Annotating Attacked Sentences

We tried capturing sentences in posts that are addressed by comments but not directly quoted. To see its feasibility, we randomly sampled 100 post-comment pairs that do not contain direct quotes and then asked an undergraduate native speaker of English (who has no knowledge about this work) to mark attacked sentences in each post, if any. This revealed two challenges. First, human annotation is subjective when compared to a co-author’s result and very time-consuming (2.5 min/comment). Second, we tried several methods to automatically identify attacked sentences. We compared the similarity between each post sentence with the comment (first sentence of the comment, first sentence of each paragraph, or all comment text) based on word overlap with/without synonym expansion and the GloVe embeddings. But it turned out to be difficult to get similar results to human annotations. Therefore, we decided to use only those sentences that are direct quoted or have at least 4 common words with a comment’s sentence as the most reliable labels.

### B External Knowledge

In this section, we describe the methods that we explored to use Kialo as a knowledge base but that were not successful.

#### B.1 UKP Sentence Embedding-Based Retrieval

We measured the similarity between CMV sentences and Kialo statements using the UKP sentence embedding—BERT embeddings fine-tuned to measure argument similarity (Reimers et al., 2019). Specifically, the authors provide pretrained embeddings constructed by appending a final softmax layer to BERT to predict a numerical *dissimilarity* score between 0 and 1 for each sentence pair in the UKP ASPECT corpus. The 3,595 sentence pairs in this corpus were drawn from 28 controversial topics and annotated via crowd workers to be “unrelated” or of “no”, “some” or “high” similarity. They report a mean F1-score of 65.39% on a held-out subset of this corpus, which was closest to human performance (F1=78.34%) among all competing methods that were not provided with additional information about the argument topic.

We used this fine-tuned model to measure the dis-

similarity between each CMV sentence and Kialo statements. Based on this information, we extracted the feature **UKP Avg Distance 10**, which is the average dissimilarity score of the 10 Kialo statements that are closest to the sentence. This score is expected to be low if a sentence has many similar statements in Kialo. In addition, we extracted the same **frequency**, **attractiveness**, and **extremeness** features as in §4.2.2. Here, we determine whether a CMV sentence and a Kialo statement are “matched” based on several dissimilarity thresholds (0.1, 0.2, 0.3, 0.4); A Kialo statement is considered matched with a CMV sentence if the dissimilarity is below the selected threshold.

#### B.2 Semantic Frame-Based Knowledge

We extracted semantic frames from CMV sentences and Kialo statements, using Google SLING (Ringgaard et al., 2017). For each frame in a sentence or statement, a “knowledge piece” is defined as the concatenation of the predicate and arguments (except negation); the predicate is lemmatized and the arguments are stemmed to remove differences in verb/noun forms. We also mark each knowledge piece as negated if the frame contains negation. Example knowledge pieces include:

- ARG0:peopl-ARG1:right-ARGM-MOD:should-PRED:have (Negation: true)
- ARG1:person-ARG2:abl-ARGM-MOD:should-PRED:be (Negation: false)

For each CMV sentence, we extracted two features: the count of knowledge pieces in Kialo that are **consistent** with those in the sentence, and the count of knowledge pieces in Kialo that are **conflicting** with those in the sentence. Two knowledge pieces are considered consistent if they are identical, and conflicting if they are identical but negated. Attackable sentences are expected to have many consistent and conflicting knowledge pieces in Kialo. If we assume that most statements in Kialo are truthful, attackable sentences may have more conflicting knowledge pieces than consistent knowledge pieces.

#### B.3 Word Sequence-Based Knowledge

Treating each frame as a separate knowledge piece does not capture the dependencies between multiple predicates within a sentence. Hence, we tried a simple method to capture this information, where a knowledge pieces is defined as the concatenation of



Knowledge Feature	Attacked	Successful
Word Overlap Frequency (log2)	1.18 (***)	1.07 (***)
Word Overlap Attractiveness (log2)	1.30 (***)	1.18 (***)
Word Overlap Extremeness	1.51 (***)	1.19 (***)
UKP Avg Distance 10 <sup>‡</sup>	0.93 (***)	0.98 ( * )
UKP 0.1 Frequency <sup>†</sup>	1.08 ( * )	0.99 ( )
UKP 0.1 Attractiveness <sup>†</sup>	1.11 ( * )	1.08 ( )
UKP 0.1 Extremeness	3.49 ( * )	6.77 ( )
UKP 0.2 Frequency <sup>†</sup>	1.02 ( ** )	1.01 ( )
UKP 0.2 Attractiveness <sup>†</sup>	1.05 (***)	1.06 ( )
UKP 0.2 Extremeness	1.69 (***)	1.76 ( )
UKP 0.3 Frequency <sup>†</sup>	1.04 (***)	1.01 ( )
UKP 0.3 Attractiveness <sup>†</sup>	1.09 (***)	1.02 ( )
UKP 0.3 Extremeness	2.44 (***)	1.40 ( )
UKP 0.4 Frequency <sup>†</sup>	1.04 (***)	1.01 ( ** )
UKP 0.4 Attractiveness <sup>†</sup>	1.12 (***)	1.01 ( )
UKP 0.4 Extremeness	2.35 (***)	1.02 ( )
Frame Knowledge Consistent	1.28 (***)	1.01 ( )
Frame Knowledge Conflict	1.37 (***)	1.08 ( )
Word Sequence Knowledge Consistent	1.05 ( )	0.98 ( )
Word Sequence Knowledge Conflict	1.18 ( )	1.49 ( )

Table 5: Odds ratio (OR) and statistical significance of features. An effect is positive (blue) if OR > 1 and negative (red) if OR < 1. (†: log2, ‡: standardized / \*:  $p < 0.05$ , \*\*:  $p < 0.01$ , \*\*\*:  $p < 0.001$ )

verbs, nouns, adjectives, modal, prepositions, subordinating conjunctions, numbers, and existential *there* within a sentence; but independent clauses (e.g., a *because* clause) were separated off. All words were lemmatized. Each knowledge piece is negated if the source text has negation words. Example knowledge pieces include:

- gender-be-social-construct (Negation: true)
- congress-shall-make-law-respect-establishment-of-religion-prohibit-free-exercise (Negation: false)

For each CMV sentence, we extracted the same two features as in semantic frame-based knowledge pieces: the count of knowledge pieces in Kialo that are **consistent** with those in the sentence, and the count of knowledge pieces in Kialo that are **conflicting** with those in the sentence.

#### B.4 Effects and Statistical Significance

The effects and statistical significance of the above features were estimated in the same way as §5 and are shown in Table 5. Word sequence-based knowledge has no effect, probably because not many knowledge pieces are matched. Most of the other features have significant effects only for “Attacked”. We speculate that a difficulty comes from the fact that both vector embedding-based matching and frame-based matching are inaccurate in many cases. UKP sentence embeddings often retrieve

Kialo statements that are only topically related to a CMV sentence. Similarly, frame-based knowledge pieces often cannot capture complex information conveyed in a CMV sentence. In contrast, word overlap-based matching seems to be more reliable and better retrieve Kialo statements that have similar content to a CMV sentence.

## C Lexicons

Table 6 shows the lexicons and regular expressions used in feature extraction. `r"pattern"` represents a regular expression.

Feature	Pattern
Question - Confusion	r" (^  ) i (\S + ){,2} (not n't never) (understand know) ", r" (not n't) make sense", r" (^  ) i (\S + ){,2} (curious confused) ", r" (^  ) i (\S + ){,2} wonder", r" (me myself) wonder "
Question - Why/How	r" (^  ) (why how) .*\?"
Question - Other	?
Normative	should, must, "(have has) to", "have got to", "'ve got to", gotta, need, needs
Prediction	r" (am\$  \$'m\$  \$are\$  \$'re\$  \$is\$  \$'s) (not )? (going to\$  \$gonna) ", will, won't, would, shall
Hypothetical	r" (^  , ) if unless"
Citation	r" {PATTERN} that [^.,!]" (PATTERN: said, reported, mentioned, declared, claimed, admitted, explained, insisted, promised, suggested, recommended, denied, blamed, apologized, agreed, answered, argued, complained, confirmed, proposed, replied, stated, told, warned, revealed), according to, r"https?:"
Comparison	than, compared to
Examples	r" (^  ) (for example for instance such as e.g\.) (  \$) "
Definition	define, definition
Personal Story	<b>Epistemic verbs:</b> think, believe, see, know, feel, say, understand, mean, sure, agree, argue, consider, guess, realize, hope, support, aware, disagree, post, mention, admit, accept, assume, convince, wish, appreciate, speak, suppose, doubt, explain, wonder, discuss, view, suggest, recognize, respond, acknowledge, clarify, state, sorry, advocate, propose, define, apologize, curious, figure, claim, concede, debate, list, oppose, describe, suspect, reply, bet, realise, defend, convinced, offend, concern, intend, certain, conclude, reject, challenge, thank, condone, value, skeptical, contend, anticipate, maintain, justify, recommend, confident, promise, guarantee, comment, unsure, elaborate, posit, swear, dispute, imply, misunderstand. <b>Epistemic nouns:</b> view, opinion, mind, point, argument, belief, post, head, position, reasoning, understanding, thought, reason, question, knowledge, perspective, idea, way, stance, vote, best, cmv, response, definition, viewpoint, example, claim, logic, conclusion, thinking, comment, statement, theory, bias, assumption, answer, perception, intention, contention, word, proposal, thesis, interpretation, reply, guess, evidence, explanation, hypothesis, assertion, objection, criticism, worldview, impression, apology, philosophy
Use of <i>You</i>	you, your, yours
Use of <i>We</i>	r" (^  ) we   (?<!the) (us our ours) (  \$) "
Subjectivity	Wilson et al. (2005)
Concreteness	Brysbaert et al. (2014)
Hedges	<b>Downtoners (score=1):</b> allegedly, apparently, appear to, conceivably, could be, doubtful, fairly, hopefully, i assume, i believe, i do not believe, i doubt, i feel, i do not feel, i guess, i speculate, i think, i do not think, if anything, imo, imply, in my mind, in my opinion, in my understanding, in my view, it be possible, it look like, it do not look like, kind of, mainly, may, maybe, might, my impression be, my thinking be, my understanding be, perhaps, possibly, potentially, presumably, probably, quite, rather, relatively, seem, somehow, somewhat, sort of, supposedly, to my knowledge, virtually, would. <b>Boosters (score=-1):</b> be definite, definitely, directly, enormously, entirely, evidently, exactly, explicitly, extremely, fundamentally, greatly, highly, in fact, incredibly, indeed, inevitably, intrinsically, invariably, literally, necessarily, no way, be obvious, obviously, perfectly, precisely, really, be self-evident, be sure, surely, totally, truly, be unambiguous, unambiguously, be undeniable, undeniably, undoubtedly, be unquestionable, unquestionably, very, wholly (Hyland, 2005; URL1; URL2)
Qualification	<b>Qualifiers (score=1):</b> a bit, a few, a large amount of, a little, a lot of, a number of, almost, approximately, except, generally, if, in general, largely, likely, lots of, majority of, many, more or less, most, mostly, much, nearly, normally, occasionally, often, overall, partly, plenty of, rarely, roughly, several, some, sometimes, tend, ton of, tons of, typically, unless, unlikely, usually. <b>Generality words (score=-1):</b> all, always, every, everybody, everyone, everything, never, no, no one, nobody, none, neither, not any, ever, forever (Hyland, 2005; URL2; URL3)
Arousal	Warriner et al. (2013)
Dominance	Warriner et al. (2013)

Table 6: Lexicons and regular expressions used in feature extraction.

## D Statistical Model for Feature Effects

For each feature, we use the following logistic regression model:

$$\log \frac{\mathbb{P}(\mathbf{Y} = 1)}{1 - \mathbb{P}(\mathbf{Y} = 1)} = \beta_0 + \beta_X \mathbf{X} + \alpha_1 \mathbf{D}_1 + \cdots + \alpha_{|D|} \mathbf{D}_{|D|},$$

where  $\mathbf{X}$  is a continuous or binary explanatory variable that takes the value of a characteristic that we are interested in.  $\mathbf{D}_d$  ( $d = 1, \dots, |D|$ ) is a binary variable that takes 1 if the sentence belongs to the  $d$ -th domain.  $\mathbf{Y}$  is a binary response variable that takes 1 if the sentence is attacked or if the sentence is attacked successfully.  $\beta_X$  is the regression coefficient of the characteristic  $\mathbf{X}$ , which is the main value of our interest for examining the association between the characteristic and the response;  $\exp(\beta_X)$  is the *odds ratio* (OR) that is interpreted as the change of odds (i.e., the ratio of the probability that a sentence is (successfully) attacked to the probability that a sentence is not (successfully) attacked) when the value of the characteristic increases by one unit. If  $\beta_X$  is significant, we can infer that  $\mathbf{X}$  has an effect on  $\mathbf{Y}$ . If  $\beta_X$  is positive (and significant), we can infer that the characteristic and the response have positive association, and vice versa.



## **E Important $n$ -gram Features**

Table 7 shows the top 100  $n$ -grams that have the highest or lowest weights for attacked sentences (vs. unattacked sentences) and for successfully attacked sentences (vs. unsuccessfully attacked).

	Attacked (vs. Unattacked)	Attacked Successfully (vs. Unsuccessfully)
High	<p>is are no - ? life why women should to society men a  nothing 1_) would money if_i they n't people if *  someone 2_ human never believe 2_) 3_ your i_belong  and 5_ americans tax 4 , being :- :_* feel because *_the  than could republicans do be government ) sex 3_)  nobody why_should the_government "i seems religion  their ca_n't less 4_ pay world war an )_the 6_  without ,_why science 4_) reason humans animals racism  military selfish racist of when social 3 gun makes you  speech climate get kids have can white should_i ,_is *_**  proven how_can</p>	<p>is without are ? would public life women weapons data  how_can usa no should if sex of_ ,_would n't why  money % someone the_us customers coffee since 1_ :  skills are_a end 3_ available ,_they technology 2_ - ,_if  people_with cost need a car the pretty_much racist  so_many to know third such_as white dog could be  towards_the americans song actions seems formal ,_he  gender_is nothing this_: power see teams job years  videos rates why_would cream expectations ca god  people feet global i_belong sounds n't_the 100  think_that_it crime to_pay firstly because ,_why immoral  and_not can_also scooby "i issues %_of ca_n't  marriages ability in_many</p>
Low	<p>edit cmv i /_? / thanks ( edit_: [ ! post ] )_( this thank  thank_you comments please view — &amp;t discussion  here topic sorry changed my_view some cmv_ posts _”  my_delta comment i_will points responses :_1_ of_you  /_ article title i_'ll 'll = thanks_for now 'm &amp;t got  i_'m was *_edit above recently reddit view_ lot i_was  below change_my hi 's a few edit_2 on_this again “ )_  my_view_ this_post discuss arguments you_all deltas  few there_are 1_ i_'ve /_)— i_have currently edit_2 :  comments_ let_me a_lot hello let i_still here_  background course )— context you_guys appreciate  thread perspective and_i posted</p>	<p>edit cmv i thanks / edit_: view this thank ! 1_ definitely  ] post discussion thank_you some 's_a changed that_this  here i_have tv points today responses above ,_it's ]_(  perspective both thought i_was to_any do_this ( there_are  &amp;t continue_to currently :i_delta comments certainly  taxes my_you_can discuss matters person_a please let_me  got ,_that not_all 'm i_'m more_of n't_want to obvious  posts friends has_been honest true_ background great  hypocritical case_ work_ account not_the results article  bit all_the that_would_be grow whose thread fine_  point_ do_you remember still hope now standard  thanks_for asking try_to go started wealth = bitcoin  series arguments super does_n't</p>

Table 7:  $n$ -grams ( $n = 1, 2, 3$ ) with the highest/lowest weights. Different  $n$ -grams are split by a space, and words within an  $n$ -gram are split by “\_”.

## F Reproducibility Checklist

Criterion	LR	BERT
Computing infrastructure	Intel(R) Core(TM) i7-3770K CPU @ 3.50GHz / 19GiB System memory	Intel(R) Core(TM) i7-8700K CPU @ 3.70GHz / 31GiB System memory / NVIDIA GP102 [TITAN Xp]
Average runtime	Attacked: 225.9 mins / Successful: 31.5 mins	Attacked: 279.5 mins / Successful: 43.4 mins
Number of parameters	20,105	108M
Validation performance	Attacked: P@1=47.4, A@3=75.8, AUC=61.8 / Successful: P@1=26.5, A@3=54.6, AUC=60.1	Attacked: P@1=50.3, A@3=77.6, AUC=64.6 / Successful: P@1=28.3, A@3=57.2, AUC=62.0
Bounds for hyperparameters	Norm: {L1, L2} / Regularization weight: {1e-4, 1e-3, 1e-2, 1e-1}	Learning rate: 1e-5 / Adam $\epsilon$ : 1e-8
Hyperparameter configurations for best-performing models	Norm: L2 / Regularization weight: 1e-1	Learning rate: 1e-5 / Adam $\epsilon$ : 1e-8
Number of hyperparameter search trials	8	(No hyperparameter search)
Method of choosing hyperparameter values	Grid search	(No hyperparameter search)
Criterion for selecting optimal hyperparameter values	AUC	(No hyperparameter search)

Table 8: Reproducibility checklist.

## G Prediction Results

Table 9 shows the prediction accuracy with an additional metric mean average precision (MAP).

	Attacked				Successfully Attacked			
	P@1	Any@3	MAP	AUC	P@1	Any@3	MAP	AUC
Random	35.9	66.0	48.0	50.1	18.9	45.0	34.0	50.1
Length	42.9	73.7	53.7	54.5	22.3	52.1	38.8	55.7
Logistic Regression	47.1	76.2	56.5	61.7	24.2	54.5	41.0	59.3
(×) Content	45.2	74.4	54.7	58.1	24.0	52.6	39.9	57.0
(×) Knowledge	47.0	76.0	56.4	61.7	24.1	54.3	40.5	59.0
(×) Prop Types	46.7	75.9	56.2	61.5	24.4	53.6	40.7	59.0
(×) Tone	47.0	76.0	56.4	61.9	25.2	56.2	41.4	59.4
BERT	49.6	77.8	57.9	64.4	28.3	57.2	43.1	62.0
Human	51.7	80.1	–	–	27.8	54.2	–	–

Table 9: Prediction accuracy.



## H Visualization Examples

For the successful example in Figure 3a, the model finds evidence for the successfully attacked sentences 3 and 5 from the external knowledge source (Kialo). Although some of the other sentences (7–8) also match Kialo statements, the degree of match is relatively low, and the model determines that their  $n$ -grams reduce attackability (*many, think, needs*). Sentence 4 is properly found to have high attackability, since it makes a comparison and contains many  $n$ -grams predictive of attackability (*because, Democrats, Republicans, opposing*).

For the successful example in Figure 3b, topics play important roles for determining attackability. The topics of the successfully attacked sentences 2–4 all increase attackability, whereas the topics of other sentences 5–9 reduce attackability.

For the erroneous example in Figure 4a, all sentences have relatively little evidence for attackability/unattackability. The model determines sentence 5 to have relatively high attackability because of many  $n$ -grams that increase attackability (*know, absolutely, nothing*). On the other hand, the successfully attacked sentence 6 is assigned a low attackability score despite its match with Kialo statements, because its use of *we*, personal stories, and certain  $n$ -grams (*many, times, and friends*).

For the erroneous example in Figure 4b, the model finds sentence 4 to have high attackability because it matches with Kialo statements, makes a comparison and prediction, and certain  $n$ -grams (*believe, presence, society, market*). Sentence 5 is also assigned a relatively high attackability score due to its use of examples and certain  $n$ -grams (*know, committed, weapons*). However, these sentences were not successfully attacked. In contrast, the successfully attacked sentences 2–4 do not have strong enough evidence for attackability compared to their negatively signals, such as personal stories and  $n$ -grams *own* and *I*.

The last **Presidential** election (2016) and most succeeding **elections** have **proven** that **elections are more** about **party** affiliations than **actual** views or **the character** of the individual being elected. In one of the most extreme **examples**, Roy Moore **was backed** by the **Republican** Party **even** though **he was** accused of sexual misconduct **and sexual** assault of minors **simply because he was** a **Republican**. **This** also allows voters to be **lazy, as many** will **simply** vote for **their party** **without** researching **the values** and **character** of the person **they are voting** for. Our Congress **is slow** an inefficient **because Democrats and Republicans are more** **focused** on **opposing one** another than **they are on** developing **actual** solutions to **issues** like gun control **and abortion**. **It is the job** of elected officials to represent **ALL** of **the people** of their district/state/country, **not just the people** that voted for **them** or agree with **them**, and following **the ideals** of a **political** party **does not** allow for this. Political **parties** force us **to think in** terms of **black and white**, and this **is both** inefficient **and** inappropriate for **issues** that affect the **entire** country. **Also, many** young voters do not **think this** way--**many Americans are** becoming disenfranchised with the **entire** political system. **This** is an outdated system, and either **needs to** adapt or change **completely to better** fit **the needs** of **the people**.

- Comparison (0.20)
- Topic35 (-0.07)
- KialoFreq (0.72)
- KialoFreq (0.14)
- Prediction (0.12)
- KialoExtr (-0.08)
- Comparison (0.20)
- Topic39 (-0.12)
- KialoFreq (0.86)
- KialoAttr (0.20)
- Use of "We" (-0.19)
- KialoFreq (0.23)
- KialoAttr (0.14)
- KialoExtr (-0.12)
- KialoFreq (0.23)
- KialoAttr (0.22)
- Normative (0.18)
- Topic33 (0.08)
- KialoExtr (-0.12)

(a) Successful example 2.

I **believe** that **socialism is an obvious** and humanitarian next **step for the U.S**. **It should be the** responsibility of vastly **successful** people to **provide a** tiny fraction of their income to provide services **for people who** were not **given** the same opportunities. Everyone has **the right to** safety, universal health **care**, social security, **education (affordable collage)**, **a livable minimum wage (\$15 per hour)**, **and not to get** screwed over by businesses more **interested** in capital than people. Businesses **don't give** their fair share back to the community **they** leach off of (**wages or taxes**), and it **should be the** responsibility of the **government** to make sure **they do**. When **many** people speak about socialism **they** quote **nations like the** U.S.S.R. (**Soviet** Union). I **believe** that the **problems** with these nations **are a** weak constitution that stems **from a violent revolution instead of a political one**. Socialism **is** an economic policy **and** can be used **in** cooperation with **the** current governing body. I **believe that many** **European** country's sudo-**socialist ideas** (like universal healthcare) **are a** perfect example of how socialism can be beneficial **to people**. B.T.W. **this is** my **first post**, I **can't wait** to have reasonable debates with **you all!**

- Topic5 (0.39)
- Normative (0.18)
- Topic28 (0.16)
- KialoFreq (0.93)
- KialoAttr (0.09)
- Topic46 (0.06)
- KialoExtr (-0.11)
- Topic30 (-0.17)
- Topic30 (-0.17)
- KialoAttr (0.44)
- KialoFreq (0.23)
- Topic30 (-0.17)
- KialoFreq (0.36)
- KialoAttr (0.18)
- Topic2 (-0.54)
- Use of "You" (-0.15)
- Personal (-0.20)
- Topic9 (-1.04)

(b) Successful example 3.

Figure 3: Successful examples.

<p>So let's say I'm good friends with Amanda and Bailey. I'm compatible with both of them on a platonic level, but I only take a romantic interest in Bailey because she's (physically) my type. Not to say that Amanda is ugly, just that I'm not really into her body structure. Another piece of evidence to support this is when you feel attracted to a complete stranger, because of their physical appearance. You know absolutely nothing about them yet, you could envision a happy relationship with them just from their looks. I feel this way because many times when I'm hanging out with my friends (of both genders) I think to myself "wow we'd make such a good couple" but even so don't feel the desire to enter a relationship with them.</p>	<p>Topic15 (-0.18) Personal (-0.20)</p> <p>Topic15 (-0.18) Personal (-0.20)</p> <p>Topic15 (-0.18) Personal (-0.20)</p> <p>Use of "You" (-0.15) Topic15 (-0.18)</p> <p>Use of "You" (-0.15) Topic15 (-0.18)</p> <p>KialoFreq (0.23) Topic15 (-0.18) Use of "We" (-0.19) Personal (-0.20)</p>
--	--

(a) Erroneous example 1.

<p>I realize I have a bias because I grew up in a big city in Canada and not a single person I knew owned a gun and most law enforcement officers I saw on the street also didn't carry guns and I perceive Canada to generally be safer than the open carry US state that I now live in. I see zero reason to own a gun, not even for hunting. I think hunters should use bows and arrows. I admit I've never been hunting myself. I believe the presence of guns in society makes society less safe and we would all be safer if there were fewer of them and they were far more difficult and expensive to buy on the black market rather than being able to pick one up easily at a gun show parking lot using cash and with no background check. I know that violence can be committed with other weapons such as knives or running someone over with a car. But we have laws about who can drive a car and it's actually more difficult to kill people with such things and less efficient.</p>	<p>KialoFreq (0.78) Comparison (0.20) Topic43 (0.19) KialoAttr (0.13) Use of "We" (-0.19) Personal (-0.20)</p> <p>Topic43 (0.19)</p> <p>Topic43 (0.19) Normative (0.18)</p> <p>Topic43 (0.19) Personal (-0.20)</p> <p>KialoFreq (0.75) Comparison (0.20) Topic43 (0.19) Prediction (0.12) KialoAttr (0.06) KialoExtr (-0.12) Use of "We" (-0.19)</p> <p>Topic43 (0.19) Example (0.11)</p> <p>KialoFreq (0.36) Topic32 (0.11) Use of "We" (-0.19)</p>
---	--

(b) Erroneous example 2.

Figure 4: Erroneous examples.