

# Detecting type of Persuasion : Is there structure in persuasion tactics?

Rahul R Iyer  
Carnegie Mellon University  
5000 Forbes Avenue  
Pittsburgh, PA 15213  
rahuli@andrew.cmu.edu

Katia Sycara  
Carnegie Mellon University  
5000 Forbes Avenue  
Pittsburgh, PA 15213  
katia@cs.cmu.edu

Yuezhang Li  
Carnegie Mellon University  
5000 Forbes Avenue  
Pittsburgh, PA 15213  
yuezhanl@andrew.cmu.edu

## ABSTRACT

Existing works on detecting persuasion in text make use of lexical features for detecting persuasive tactics, without taking advantage of the possible structures inherent in the tactics used. In this paper, we propose a multi-class classification, unsupervised domain-independent model for detecting the type of persuasion used in text, that makes use of the sentence structure inherent in the different persuasion tactics. Our work shows promising results as compared to existing work, and vector-embedding models.

## KEYWORDS

persuasion detection, multi-class classification, text mining, unsupervised learning

## 1 INTRODUCTION

Persuasion is being used at every type of forum these days, from politics and military to social media. Detecting persuasion in text helps address many challenging problems: Analyzing chat forums to find grooming attempts of sexual predators; training salespeople and negotiators; and developing automated sales-support systems. Furthermore, ability to detect persuasion tactics in social flows such as SMS and chat forums can enable targeted and relevant advertising. Additionally, persuasion detection is very useful in detecting spam campaigns and promotions on social media: especially those that relate to terrorism. Persuasion identification is also potentially applicable to broader analyses of interaction, such as the discovery of those who shape opinion or the cohesiveness and/or openness of a social group.

Existing work on detecting persuasion in text, focuses mainly on lexical features without taking advantage of the inherent structure present in persuasion tactics. In this work, we attempt to build an unsupervised, domain-independent model for detecting persuasion tactics, that relies on the sentence structures of the tactics. Our contributions to the literature are: 1) we show that persuasive tactics have inherent sentential structures that can be exploited, 2) we propose an unsupervised approach that does not require annotated data, 3) we propose a way to synthesize prototype strings for the different persuasion tactics, 4) our approach takes much less time to execute as compared to models that require training; for example, our approach is faster than Doc2Vec by a factor of almost 1.5, 5) our approach is domain-independent, in that it is independent of the vocabulary and can be applied

to various domains such as politics, blogs, supreme court arguments etc., with very minimal changes (unlike vector-embedding models and other models that make use of lexical features, because they are dependent on the vocabulary).

We compare our proposed approach with existing methods that use lexical features, and also some vector embedding models, such as Doc2Vec [10].

We had an intuition about arguments in the persuasive space having similar sentential structures because we had seen and observed a few examples. Consider two examples in the *Reason* category: 1a) *Are we to stoop to their level just because of this argument?*, 1b) *I am angry at myself because I did nothing to prevent this*, and two examples in the *Scarcity* category: 2a) *Their relationship is not something you see everyday*, 2b) *It is only going to go downhill from here*. As we can see, there is a pattern in the structure between arguments in the same category, and there is a structural difference across these two categories. This led us to investigate the problem further and hypothesize our claim.

The rest of the paper is organized as follows: Section 1.1 talks about the related work that has been done in the area, section 2 explains the problem that we are trying to tackle, section 3 gives descriptions about the different datasets that have been used in the paper, section 4 explains the proposed model and the baselines, section 5 discusses the experimental results obtained, section 6 goes over some brief applications of the model, and section 7 concludes the paper with a discussion and future work.

### 1.1 Related Work

There has been some work in the literature on detection of persuasion in texts. In [22], Young et al. present a corpus for persuasion detection, which is derived from hostage negotiation transcripts. The corpus is called “NPS Persuasion Corpus”, consists of 37 transcripts from four sets of hostage negotiation transcriptions. Cialdini’s model [5] was used to hand-annotate each utterance in the corpus. There were nine categories of persuasion used: reciprocity, commitment, consistency, liking, authority, social proof, scarcity, other, and non-persuasive. Then algorithms like Naive Bayes, SVM, Maximum Entropy were used for the classification.

Gilbert [8] presented an annotation scheme for a persuasion corpus. A pilot application of this scheme showed some agreement between annotators, but not a very strong one. After revising the annotation scheme, a more extensive study

showed significant agreement between annotators. The authors in [14], determined that it is possible to automatically detect persuasion in conversations using three traditional machine learning techniques, naive bayes, maximum entropy, and support vector machine. Anand et al. [1] describe the development of a corpus of blog posts that are annotated for the presence of attempts to persuade and corresponding tactics employed in persuasive messages. The authors make use of lexical features like unigrams, topic features from LDA, and List count features from the Linguistic Inquiry and Word Count [7], and also the tactics themselves, which are provided by human annotators. Tactics represent the type of persuasion being employed: social generalization, threat/promise, moral appeal etc. Carlo et al. [18] analyze political speeches and use it in a machine learning framework to classify the transcripts of the political discourses, according to their persuasive power, and predicting the sentences that trigger applause in the audience. In [19], Tan et al. look at the wordings of a tweet to determine its popularity, as opposed to the general notion of author/topic popularity. The computational methods they propose perform better than an average human. In [12], Lukin et al. determine the effectiveness of a persuasive argument based on the audience reaction. They report a set of experiments testing at large scale how audience variables interact with argument style to affect the persuasiveness of an argument.

In addition to text, there has been some work on persuasion in the multimedia domain. In [17], Siddiquie et al. work on the task of automatically classifying politically persuasive videos, and propose a multi-modal approach for the task. They extract audio, visual and textual features that attempt to capture affect and semantics in the audio-visual content and sentiment in the viewers’ comments. They work on each of these modalities separately and show that combining all of them works best. For the experiments, they use *Rallying a Crowd* (RAC) dataset, which consists of over 230 videos from YouTube, comprising over 27 hours of content. Chatterjee et al. [4] aim to detect persuasiveness in videos by analysis of the speaker’s verbal behavior, specifically based on his lexical usage and paraverbal markers of hesitation (a speaker’s stuttering or breaking his/her speech with filled pauses, such as *um* and *uh*). Paraverbal markers of hesitation have been found to influence how other people perceive the speaker’s persuasiveness. The analysis is performed on a multimedia corpus of 1000 movie review videos annotated for persuasiveness. Park et al. collected and annotated a corpus of movie review videos in [15]. From this data, they demonstrate that the verbal and non-verbal behavior of the presenter is predictive of how persuasive they are as well as predictive of the cooperative nature of a dyadic interaction.

Tasks similar to persuasion detection have been explored, such as sentiment detection and perspective detection. Lin et al. investigated the idea of perspective identification at the sentence and document level [11]. Using the articles from the bitterlemons website<sup>1</sup>, they were able to discriminate

<sup>1</sup><http://www.bitterlemons.org>

between Palestinian authors and Israeli authors who had written about the same topic. Bikel and Soren used machine learning techniques to differentiate between differing opinions [2]. They report an accuracy of 89% when distinguishing between 1-star and 5-star consumer reviews, using only lexical features.

## 2 PROBLEM FORMULATION

### 2.1 Preliminaries

Persuasion is an attempt to influence a person’s beliefs, attitudes, intentions, motivations, or behaviors. In persuasion, one party (the ‘persuader’) induces a particular kind of mental state in another party (the ‘persuadee’), like flattery or threats, but unlike expressions of sentiment, persuasion also involves the potential change in the mental state of the other party. Contemporary psychology and communication science further require the persuader to be acting intentionally. Correspondingly, any instance of (successful) persuasion is composed of two events: (a) an attempt by the persuader, which we term the persuasive act, and (b) subsequent uptake by the persuadee. In this work, we consider (a) only, the different persuasive acts, and how to detect them. Working with (b) is a whole other problem. Throughout the rest of the paper, when we say *persuasive arguments*, we mean the former without taking the effectiveness of the persuasion into account. We are only interested in whether the arguments contain persuasion, and if so, the type.

### 2.2 Problem Statement

The main objective of our work is to detect whether a given piece of text contains persuasion or not. If it does, then we can look into the type of persuasion strategy being used, such as threat/promise, outcome, reciprocity etc. In this paper we look at 14 different persuasion strategies. These are listed in Table 1. These are the common tactics for persuasive acts contributed by Marwell and Scmitt [13], Cialdini [5], as well as argumentative patterns inspired by Walton et al. [21]. The intuition behind this investigation along with some examples was discussed in section 1.

It has to be noted that an entire text is deemed to contain persuasion, if it includes a few arguments that use some of these tactics to persuade. So, it is important to extract such arguments from the text, before applying the persuasion model to them. So, our approach has two steps: 1) a very simple argument extractor model, to extract arguments from a given piece of text, and 2) the output of the extractor is fed into the persuasion detection model, which classifies the arguments into the different tactic classes. This process is represented in the flowchart shown in Figure 1. It has to be noted that, in this work, we are not concerned with the effectiveness of the persuasion.

## 3 DATASETS

We have used datasets from different domains for the experiments to test the robustness of our model. Below are the datasets used for persuasion detection:




Figure 1: The outline of the problem considered

Outcomes	Generalizations	Broad Categories		Other
		External	Interpersonal	
<p><b>Outcome.</b> Mentions some particular consequences from uptake or failure to uptake</p> <p><b>Social Esteem.</b> States that people the persuadee values will think more highly of them</p> <p><b>Threat/Promise.</b> Poses a direct threat or promise to the persuadee</p> <p><b>Self-Feeling.</b> States that uptake will result in a better self-valuation by the persuadee</p>	<p><b>Good/Bad Traits.</b> Associates the intended mental state with a “good” or “bad” person’s traits.</p> <p><b>Deontic/Moral Appeal.</b> Mentions duties or obligations, goodness, badness</p>	<p><b>VIP.</b> Appeals to authority (bosses, experts, trend-setters)</p> <p><b>Popularity.</b> Invokes popular opinion as support for uptake</p>	<p><b>Favors/Debts.</b> Mentions returning a favor or injury</p> <p><b>Consistency.</b> Mentions keeping promises or commitments</p> <p><b>Empathy.</b> Attempts to make the persuadee connect with someone else’s emotional perspective</p> <p><b>Scarcity.</b> Mentions rarity, urgency, or opportunity of some outcome</p>	<p><b>Recharacterization.</b> Reframes an issue by analogy or metaphor.</p> <p><b>Reasoning.</b> Provides a justification for an argumentative point based upon additional argumentation schemes e.g., causal reasoning, arguments from absurdity</p>

Table 1: List of Persuasion Tactics Used, Taken from [1]. We do not consider the broad categories in our experiments, and only work with the finer categories.

1. **ChangeMyView**, an active community on Reddit, provides a platform where users present their own opinions and reasoning, invite others to contest them, and acknowledge when the ensuing discussions change their original views. The training data period is: 2013/01/01 - 2015/05/07, and the test data period is: 2015/05/08 - 2015/09/01. The training dataset contains 3456 posts and the holdout dataset contains 807 posts. The dataset is organized as follows: each post that is written by a user who wants his views changed, has two argumentative threads – one that is successful and one that is not. This has been used to determine the persuasion strategies employed by the successful thread. This dataset is used in [20].

2. **Supreme Court Dialogs Corpus**: This corpus contains a collection of conversations from the U.S. Supreme Court Oral Arguments. 1) 51,498 utterances making up 50,389 conversational exchanges, 2) from 204 cases involving 11 Justices and 311 other participants 3) metadata like case-outcome, vote of the Justice, gender annotation etc. This dataset is used in [6].

3. **Blog Authorship Corpora**: This dataset, contributed by Pranav Anand et al. and used in [1], is a subset of the blog authorship corpus. Each directory corresponds to a blog. Each blog has sub-directories corresponding to a day. Inside the day sub-directory, there may be multiple posts; posts are identified with an underscore followed by a number. Out of around 25048 posts, only around 457 were annotated with persuasive acts. Each blog post has been broken down into different “text tiles”, which are a few sentences long, and

each of these tiles are annotated with a persuasive tactic (if present).

4. **Political Speeches**: We collected a number of speeches of Donald Trump and Hillary Clinton, to analyze the kinds of persuasive tactics they use.

To train the argumentation extraction model, we use an **Argumentation Essay Dataset**<sup>2</sup>: This consists of about 402 essays. There are two files for each essay - the original essay and the annotated file. Annotations include breaking down the essay into different components: claim, premise, stance. This can be used to train a simple classifier, to identify arguments from text passages.

The Blog Authorship Corpus is already annotated, as noted above. In order to have the ground truth, i.e. the annotations, for the other datasets, we needed to annotate the arguments of the corpora with the persuasion tactics mentioned in Table 1. For this, we used Amazon Mechanical Turk<sup>3</sup>. Using the argument extraction model, we extracted arguments from all of the corpora combined (excluding the blog authorship corpus)<sup>4</sup>. We had each argument annotated by two different turkers, and the turkers were given the freedom to classify a piece of text as either a non-argument or as one of the tactics from Table 1. There was about 65% inter-annotator agreement, between the turkers and the

<sup>2</sup><https://www.ukp.tu-darmstadt.de/data/argumentation-mining/argument-annotated-essays-version-2/>

<sup>3</sup><https://www.mturk.com/mturk/>

<sup>4</sup>The whole dataset, along with the annotation guidelines, classification criteria and the prototype strings (both median and synthetic) for all the persuasion tactics, can be found at <https://github.com/rrahul15/Persuasion-Dataset>

conflicts were resolved manually. After this, we had a total of 1457 persuasive arguments from all the datasets combined. The distribution of arguments from the different datasets is given in Table 2. The guidelines for annotation were built on the ones provided in [1], with some changes<sup>4</sup>.

Dataset	# Arguments
ChangeMyView	362
Supreme Court	440
Political Speeches	198
Blog	457

**Table 2: Distribution of arguments from the different datasets**

## 4 TECHNICAL APPROACH

In this section, we describe our proposed models, along with a couple of baselines for comparison. We describe the baselines in the following section. It has to be kept in mind, as noted in section 3, that the dataset used consists solely of persuasive arguments.

### 4.1 Baselines

Here, we discuss the different baselines that we use for comparison. We describe a simple supervised approach that makes use of lexical features and then move onto more complicated models involving vector-embedding. In all the supervised approaches, we use a 80 : 20 split for training and testing.

- (1) **Simple Supervised:** Here, the learning phase involved extracting simple textual features from the training set: unigrams, bigrams, without punctuation, and then training an SVM (Support Vector Machine) model, using Sequential Minimal Optimization (SMO) [16], to learn a model from these features that could be applied to the holdout set. This model was then used to test the remaining posts.
- (2) **Supervised Document Vectors:** This method uses the Doc2Vec model proposed by Quoc and Mikolov [10]. First the arguments were separated into different categories based on the persuasion tactic. Then, the Doc2Vec model was applied, to each such cluster, to embed all the arguments into vectors. The prototype vector for each category was then chosen as the mean of all the vectors in that category. To classify the holdout set, one would compute the vector of the argument in consideration and then compute the similarity (cosine) to the prototype vectors. The category which has the highest similarity is the one that is chosen.

The cosine similarity between two vectors  $\mathbf{a}$  and  $\mathbf{b}$  is defined as follows:

$$\text{similarity} = \frac{\mathbf{a} \cdot \mathbf{b}}{\|\mathbf{a}\| \|\mathbf{b}\|} \quad (1)$$

- (3) We also compare our approach with that proposed by [1]. Here, the authors make use of different features to account for fewer word-dependent features a) 25 topic features, which were extracted using Latent Dirichlet Allocation (LDA) [3], with a symmetric Dirichlet prior b) 14 Tactic count features, i.e., a vector consisting of the count of the tactics. Naive Bayes was used for the classification, to assess the degree to which these feature sets complement each other.

### 4.2 Proposed Approach

In this section, we describe the proposed unsupervised, domain-independent approach to identify the persuasion tactics in a given set of arguments. By domain-independence, we mean that our proposed model is robust across different genres, be it political speeches or blogs, and this is important because different domains might have their own vocabulary. Before heading into the details of the algorithm, we present a few useful definitions.

**4.2.1 Preliminaries.** In this subsection, we describe a few preliminary concepts

**Parse Tree:** A parse tree is an ordered, rooted tree that represents the syntactic structure of a sentence, according to some context-free grammar. It captures the sentence structure: multiple types of sentences can have a similar sentence structure, even if their vocabularies are not the same. This is the essence of the approach.

**Edit Distance:** Edit distance is a way of quantifying how dissimilar two strings are to one another by counting the minimum number of operations required to transform one string into the other. Given two strings  $a$ , and  $b$  on an alphabet  $\Sigma$ , the edit distance  $d(a, b)$  is the minimum number of edit operations that transforms  $a$  into  $b$ . The different edit operations are: 1) Insertion of a single symbol, 2) Deletion of a single symbol, 3) Substitution of a single symbol, for another.

**Median String:** The median string of a set of strings is defined to be that element in the set which has the smallest sum of distances from all the other elements [9]. In our case, the distance between strings is the edit distance.

**4.2.2 Parse-Tree Model.** We are proposing a domain-independent classification. The idea is that persuasive arguments may have certain characteristic sentence structures, which we might be able to exploit. The training and testing phase are given below:

#### Training Phase

- (1) As mentioned earlier, we have 14 different categories for persuasive tactics. We obtain one representative prototype argument for each category. We obtain

these in two different ways, which we discuss after detailing the algorithm.

- (2) We then perform phrase-structure parsing on each of these prototype arguments to obtain their parse-trees, which gives the structure of the argument.
- (3) These parse trees are then converted into *parse strings*, keeping the structure intact, and the leaf nodes (the terminal symbols, namely words) are removed, to get a domain independent representation of the structure of the argument.
- (4) By now, we have representative prototype parse strings for each persuasive tactic, i.e. 14 different prototype parse strings. We use these strings to classify a new argument into one of the persuasive categories. As mentioned earlier, every instance in the dataset is a persuasive argument. This can be construed as the “training phase”.

### Testing Phase

- (1) For the testing phase, we have to classify a new argument into one of the categories. Since, each new argument in the dataset is persuasive, we don’t have to worry about non-arguments. We build a model to account for non-arguments in section 6.
- (2) Given a new argument, compute its parse string, similar to the procedure used in obtaining the parse strings for the prototype arguments.
- (3) To then classify this argument, we compute the normalized edit distances (normalized by the lengths of the strings) between its parse string and the prototype parse strings of each category.
- (4) The persuasive category with the least edit distance is logically the most structurally similar to the given argument, and hence the argument is classified into that category.
- (5) This process is explained in the flowchart, given in Figure 2.

**Choosing the Prototype Strings** We propose two methods to obtain the prototype argument strings.

- (1) *Median as the Prototype*: Take a set of arguments from each persuasion category and obtain the prototype string for that category as the median string of the set. We now have to determine the ideal size of the set in question. For obvious reasons, we get the best representation if we consider all arguments of that category, but this would require a completely annotated dataset (making the model supervised).

In order to determine the ideal set size, we conduct additional experiments on a particular dataset, the supreme court dataset, with different parameter values to observe the trend of the performance. We choose different set sizes: 2%, 5%, 10%, 20%, 30% and All (the set sizes chosen is a percentage of the total number of arguments in that category). For each set size, we conduct 5 different trials, choosing a random sample each time, to see the average performance. This trend is shown in Figure 3. As we can see, the

performance across successive trials stabilizes as we increase the set size. The performance is best when we consider all the arguments, and is quite stable and close to the best when we consider 30% of all the arguments. So, we settle for 30% as the ideal set size because we get a stable performance with a very small loss in accuracy and much fewer arguments. Examples of the prototype parse strings, using the median method, for two different persuasion tactics: Reasoning and Scarcity, are given in Figure 7. We only display prototypes for two of the tactics for purposes of brevity<sup>4</sup>.

Now, although different arguments in the same category are structurally similar, they may each have certain parts in their structure that capture the essence of that category much better. We rule out taking advantage of these individual segments, when we pick one median argument out of the set. This led us to the second method of obtaining prototype strings.

- (2) *Synthetic Prototype*: We noted earlier that there could be certain segments in different arguments of the same category, that capture the essence of the category better. To accommodate this, we chop up the different arguments in a set into a number of segments and choose different segments to synthesize an artificial prototype string. To obtain the best  $i^{th}$  segment for the synthetic string, we choose the median of the  $i^{th}$  segments for all strings in the set. It has to be noted that we chop the strings uniformly. This process of synthesizing the prototype string is illustrated in Figures 4 and 5. As before, we need some parameters to tune here. In addition to the optimal set size, we also need to determine the optimal number of segments.

In order to determine the optimal number of segments and set size, we conduct additional experiments on the supreme court dataset, with different parameter values to observe the trend as before. We choose different set sizes: 2%, 5%, 10%, 20%, 30% and All, as before, and different number of segments: 2, 3, 5, 7, and 9. For each (set size, number of segment) pair, we conduct 5 trials, choosing a random sample for the sets each time, and compute the average performance. This trend is shown in Figure 6. We do not show the performance for different trials as before, rather just the average across the trials. We see that the trend stabilizes as we increase the set size, as before, and the accuracy improves as we consider more number of segments. But here, it is a tradeoff between accuracy and speed because having a large number of segments will require us to compute the median for every segment. We settle for 30% as the ideal set size and 9 as the ideal number of segments. The prototype strings using this method

are not parse strings of meaningful sentences and so we do not display them here.




Figure 3: Trend Graph for Median as the Prototype string




Figure 5: Synthesis of the prototype string. The coloured segments are the medians of those string segments. In this case, the median of each segment is coming from a different string. This is just for illustrative purposes, it need not always be the case.




Figure 6: Trend Graph for Synthetic Prototype string

## 5 EXPERIMENTAL RESULTS

In this section, we present the results of our proposed model and compare it with the different approaches described earlier. In the following section, we discuss the metrics that we have used for evaluation.

### 5.1 Evaluation Metrics

The metrics used for evaluation are listed below.

- (1) Precision: The percentage of arguments the system identified as having a particular tactic that in fact had that tactic

$$\text{Precision}_t = \frac{|\{\text{retrieved documents}\} \cap |\{\text{relevant documents}\}|}{|\{\text{retrieved documents}\}|} \quad (2)$$

where  $\text{Precision}_t$  is the precision for tactic  $t$ .

- (2) Recall: the percentage of arguments of a particular tactic that the system classified into that category

$$\text{Recall}_t = \frac{|\{\text{retrieved documents}\} \cap |\{\text{relevant documents}\}|}{|\{\text{relevant documents}\}|} \quad (3)$$

where  $\text{Recall}$  is the recall for tactic  $t$ .

- (3)  $F_1$ -measure: the harmonic mean of precision and recall

$$F_{1_t} = \frac{2 \times \text{precision}_t \times \text{recall}_t}{\text{precision}_t + \text{recall}_t} \quad (4)$$

where  $F_{1_t}$  is the  $F_1$  measure for tactic  $t$ .

It is important to note that, the precision, recall and  $F_1$  measure are computed for each persuasion tactic separately, akin to a binary classifier. We report the mean of these measures, over all the tactics, in our experiments.

### 5.2 Results

First, we run the proposed parse-tree model on the arguments extracted from the datasets and obtain the *average per-category accuracy*. The per-category accuracy is defined as the percentage of accurate classifications for a specific category. The categories are the persuasion tactics in our case. For this task, we combined the arguments from each of the 4 datasets to form a combined set, in order to get an average performance estimate (refer to Table 2 for the distribution of arguments in each dataset). We classified the arguments in the combined set and calculated the fraction of correct classifications for each category. The results are given in Table 3. We do not consider the broad categories in our experiments, and only work with the finer categories.

We also compute the distribution of the tactics in the different datasets. We do this by classifying the arguments in each of the 4 datasets, and calculating the frequency of appearance of each tactic in the corpus as a percentage over all the arguments in that corpus. These are listed in Tables 4-7. The ranking of the tactics in these tables, with respect to the percentages, aligns closely with manual evaluations. These distributions are shown just to give an idea of the ranking of the tactics, as predicted by the algorithm (which makes sense intuitively).

**ChangeMyView:** Each user posts his/her stance on a particular topic and challenges others to change his opinion. For example, one of the posts was about a man who did not believe in essential-oils and believed that they were destructive, whereas his wife believed the oils were beneficial. He requested the other users to make him change his mind about essential oils by giving him sufficient evidence. If a person is




Figure 2: The parse-tree model explained





Figure 4: Outline for Synthesis of the Prototype strings



(a) Parse Tree for Scarcity

Sentence: Their relationship is not something you see everyday

Parse String: (NP+SBAR+S (S (NP (PRP\$) (NN)) (VP (VBZ) (RB) (NP (NN)))) (S (NP (PRP)) (VP (VB) (NP (DT) (NN)))))



(b) Parse Tree for Reasoning

Sentence: I'm angry because of this, I did NOTHING

Parse String: (SBAR+S (NP (PRP)) (VP (VBP) (VB) (SBAR (IN) (S (PP (IN) (NP (DT))) (,) (NP (PRP)) (VP (VBD) (ADJP (JJ)))))))

Figure 7: Examples of parse trees with the parse-strings for 2 persuasion tactics: Scarcity and Reason. This is just an illustration to show how the parse strings look like.

Category	Accuracy
Reasoning	79.8
Deontic/Moral Appeal	69.6
Outcome	65.7
Empathy	61.3
Threat/Promise	58.2
Popularity	56.4
Recharacterization	54.9
VIP	53.5
Social Esteem	50.3
Consistency	45.6
Favors/Debts	41.1
Self-Feeling	37.7
Good/Bad Traits	35.5
Scarcity	29.6

Table 3: Per-category Accuracy for the parse-tree model, when 14 categories are used

successful in changing the mind of the OP (Original Poster), the OP gives that person a *delta* in their comments. All the conversations are monitored by Reddit and hence the quality is high.

In our dataset, for each post, there are two threads of comments – one successful in changing the mind of the OP

and one that is unsuccessful. We analyzed the persuasive strategies that are used by the successful threads because these would be examples of good uses of the different persuasion tactics. For our purposes of classifying tactics, we could have also used the unsuccessful threads (we are not concerned about the uptake of the persuasion by the persuadee) but we chose not to. Firstly, we extracted the positive comments from the threads (those which were given a delta by the OP). We then applied the parse-tree persuasion model that we developed earlier, to these texts, to perform the classification. Many of the comments had links to other credible sources which listed facts that were opposed to the OP's view. We did not venture into these links. After determining the persuasion strategies used in the comments, we observed that *Reasoning* and *Outcomes* were the most frequently used strategies. A more detailed distribution of tactics is given in Table 4.

Tactic	Percentage
Reasoning	40.7
Outcomes	41.2
Good/Bad traits	10.0
Social	8.1

Table 4: Distribution of Persuasion Tactics used in the ChangeMyView dataset.

**Supreme Court Dataset:** This dataset includes the transcript of the conversation exchanges over 204 cases, along with the outcome of the cases. The outcome could either be *Respondent* or *Petitioner*. Petitioner is the person who files the petition/case against a particular party requesting action on a certain matter, and the respondent is the person against whom the said relief is sought. We have collected all the cases where the petitioner has won and analyzed the argument structure.

Tactic	Percentage
Deontic Appeal	33.3
Reasoning	35.5
Recharacterization	12.6
Outcome	8.6
Empathy	5.2
VIP	4.8

**Table 5: Distribution of Persuasion Tactics used in the Supreme Court dataset.**

We have taken these cases and analyzed the arguments. Using the argumentation-model, we were able to identify the key arguments and then using the parse-tree model, we were able to classify the type of argument that was used. It was found that most of the presented arguments were *Deontic Appeal* and *Reasoning*. The distribution of arguments is given in Table 5.

**Political Speeches:** We analyze the persuasive tactics present in the speeches of political candidates, specifically those of Donald Trump and Hillary Clinton. These distributions are given in Tables 6 and 7. It was observed that the most frequently used tactic by Trump was Outcome (“Make America Great Again”), while for Hillary, the most frequently used tactic is Empathy

Tactic	Percentage
Outcome	39.1
Principles	31.2
VIP	18.5
Reasoning	11.2

**Table 6: Distribution of Persuasion Tactics used in Trump’s Speeches.**

Finally, we present the results of the performance of the different algorithms, described earlier, in Table 8. We also performed these experiments in a binary setting: whether a given argument contains persuasion or not. These results are presented in Table 9. We have run these experiments on the arguments extracted from each dataset (refer to Table 2 for the distribution of arguments in each dataset). The

Tactic	Percentage
Empathy	35.2
Consistency	33.8
Favors/Debts	18.2
Social	12.8

**Table 7: Distribution of Persuasion Tactics used in Hillary’s Speeches.**

performances are measured by the precision (P), recall (R) and the  $F_1$  measure (F), as described earlier.

As can be seen, the domain-independent parse-tree model with synthetic prototype strings performs the best, almost 7–8% better than Doc2Vec. Thus, our intuition that different segments of arguments in the same category capture the essence of the category better than others, is validated. It has to be noted that in a multi-class classification setting, the  $F_1$  scores, obtained in Table 8, are reasonable. It also has to be noted that our model performs faster than Doc2Vec, by a factor of almost 1.5.

### 5.3 Sensitivity Analysis

We also performed a sensitivity analysis on the parse-tree model to observe its robustness. For this, we combined the arguments from all the 4 datasets and ran the model on the combined set: 1) using only 10 instances in the dataset, 2) using only 100 instances in the dataset, 3) using 1000 instances in the dataset, and 4) using all the instances in the dataset. The prototype argument strings, for this dataset, were synthesized according to the method mentioned earlier. The results are given in Table 10. For the cases which did not involve the whole dataset, we randomly sampled 5 times from the whole corpus and averaged the results. As we can see, the results show that the proposed model is relatively robust and invariant with the amount of data.

Another aspect of sensitivity analysis, involving variation of the parameters of the models was discussed earlier in section 4.2.2.

Data	Metrics		
	P	R	F
10	0.370	0.375	0.372
100	0.402	0.393	0.397
1000	0.491	0.387	0.432
ALL	0.442	0.412	0.426

**Table 10: Sensitivity Analysis of the parse-tree model using 10, 100, 1000 and all instances of data.**



## 6 APPLICATIONS

The persuasion-detection model that is proposed here, is very versatile and can be applied in many scenarios. A few applications are detailed below

### Basic Argument vs. Non-argument Classifier

The parse-tree model proposed can be used as an argument classifier. To test its applicability here, we collected a set of simple sentences<sup>5</sup> with less than 10 words. We collected a total of 1000 such simple sentences. It was our intuition that such simple phrases should have a very low similarity with the different persuasion categories (as they are structurally different). So, we needed to establish a threshold to classify a particular piece of text as an argument versus a non-argument. If the normalized edit distance similarities between the given string and the prototypes of the different categories are less than that threshold (all of them should be less than the threshold), classify it as a non-argument. Else, classify it into its correct persuasion category. Choosing this threshold is not such an easy task for these reasons:

- 1) Higher threshold: Lower chance of classifying a non-argument as an argument and Higher chance of classifying an argument as a non-argument
- 2) Lower Threshold: Higher chance of classifying a non-argument as an argument and Lower chance of classifying an argument as a non-argument

So, we needed to choose a threshold that is neither too high nor too low. We tried different thresholds, and we show a graph of threshold vs. accuracy in Figure 8. The accuracy here is the mean  $F_1$  score. As we can see, the threshold of 0.1 seems to work best. For this set threshold, the  $F_1$  score is observed to be 0.412. It can be seen that the performance of this system is not as good as with just persuasive arguments. It has to be noted here that this is not a binary problem: argument vs. non-argument. The  $F_1$  score presented here is for the problem of 14 persuasion tactic categories vs. non-arguments.

This model cannot be used, in this form, as a robust argument classifier yet because some non-arguments have structures similar to some of the persuasion tactics described earlier. For example, consider the sentence: “The men smoked and most of the women knitted while they talked”. Although this is a non-argument, the model could confuse it with one of the persuasion categories like Reason/Promise. It is for this reason that we consider very simple, straight-forward non-argument sentences of a few words. In order to build a classifier with such capabilities, we would be required to incorporate domain-independent lexical features to the parse tree model (more discussion in section 7.1).

<sup>5</sup><http://www.cs.pomona.edu/~dkauchak/simplification/>




Figure 8: Graph of Threshold vs. Accuracy ( $F_1$  Score)

### Political Speech Analysis

The parse-tree model can also be used to detect spam campaigns on social media, and to detect terrorist campaigns. We analyzed some speeches of Osama Bin Laden to see what kind of tactics he used to influence the people in his speeches. From the analysis, mostly *Empathy* was used. The detailed distribution is given in Table 11.

Tactic	Percentage
Empathy	29.7
Recharacterization	20.3
Reasoning	12.9
Good/Bad Traits	20.2
Scarcity	8.0
Outcome	8.9

Table 11: Distribution of Persuasion Tactics used in Osama Bin Laden’s Speeches.

## 7 DISCUSSION

We have proposed a fairly simple, domain-independent unsupervised model for detecting types of persuasion used in text. This model can be used in any context/domain because it only uses the inherent structure of the persuasion tactics. This versatility gives it a variety of applications. Almost all persuasive arguments can be classified under the categories mentioned earlier. It has to be noted that the reason we did not include lexical features to our model is because that would make the model slightly domain-dependent. It is for this reason, that we just focused on the structural aspects. Of course, as we mention in section 7.1, it might be possible to include a few lexical terms like *because*, *if*, *while* etc., which are domain-independent and tactic-dependent, to further strengthen our model.

From the obtained results, we see that our model’s accuracy is highest for the following persuasion tactics (refer to Table 3): reasoning, deontic/moral appeal, outcome, empathy. This is in agreement with the observations made by the authors in [1]. This further validates our model.

From the metrics computed, we see that the performance of the proposed model, with synthesized prototype strings, is better than that of vector-embedding models, such as Doc2Vec, which uses deep learning. There is almost a 7 – 8% improvement in performance (compared to Doc2Vec), and this over 14 categories in total is a reasonable amount. This is very interesting because: 1) The Doc2Vec model that we have used as baseline uses all the data, whereas our model just uses a very small subset of the dataset to compute the initial set of 14 prototype argument strings, one for each category, 2) The Doc2Vec model requires a training phase which can take considerable amount of time, considering its complex structure. In addition, our model runs faster than Doc2Vec by a factor of almost 1.5, as already mentioned earlier. The reason that our model beats lexical-features based methods could be attributed to the fact that domain-words somehow restrict the performance.

We see that the baseline SVM, on lexical features had considerably high precision but it fell short on recall. That is why we use the  $F_1$  measure, as a combination of both aspects of the model: precision and recall. The most suitable model should be the one with a high  $F_1$  score, which can be achieved with high values for both precision and recall. We also see that our model performs better, at both binary and multi-class classification, than the approach used in [1].

Sensitivity analysis was also done to test the robustness of the method. As can be seen from the results, the method is fairly stable and robust with respect to the size of the data.

The main takeaway here is that, complex methods like neural networks may not be the best for all tasks. We have shown that with very simple methods like the one we have proposed, we are able to achieve a performance better than the methods discussed in the paper, while avoiding high computational costs and the opacity of results of neural computation based methods.

## 7.1 Future Work

There is scope for improving the proposed model further. As of now, we just use the sentence structures of the different persuasion tactics. We have not made use of the fact that there could be some domain-independent words for each tactic, like *because*, *if*, *while* etc. Incorporation of such keywords into the model could result in improved performance. We will investigate this in the future. Additionally, we will use our approach for different applications, such as detecting spam campaigns, measuring how effective a spam campaign can be (combination of persuasiveness and connectivity in the network, which can be measured by PageRank), identifying terrorist campaigns etc.

## 8 ACKNOWLEDGEMENTS

This work has been funded by ARO award #W911NF-13-1-0416.

## REFERENCES

- [1] Pranav Anand, Joseph King, Jordan L Boyd-Graber, Earl Wagner, Craig H Martell, Douglas W Oard, and Philip Resnik. 2011.

- Believe Me-We Can Do This! Annotating Persuasive Acts in Blog Text.. In *Computational Models of Natural Argument*.
- [2] Daniel M Bikel and Jeffrey Sorensen. 2007. If we want your opinion. In *International Conference on Semantic Computing (ICSC 2007)*. IEEE, 493–500.
- [3] David M Blei, Andrew Y Ng, and Michael I Jordan. 2003. Latent dirichlet allocation. *Journal of machine Learning research* 3, Jan (2003), 993–1022.
- [4] Moitrey Chatterjee, Sunghyun Park, Han Suk Shim, Kenji Sagae, and Louis-Philippe Morency. 2014. Verbal behaviors and persuasiveness in online multimedia content. *SocialNLP 2014* (2014), 50.
- [5] Robert B Cialdini. 2001. Influence: Science and practice. *Boston: Allyn & Bacon* (2001).
- [6] Cristian Danescu-Niculescu-Mizil, Lillian Lee, Bo Pang, and Jon Kleinberg. 2012. Echoes of power: Language effects and power differences in social interaction. In *Proceedings of the 21st international conference on World Wide Web*. ACM, 699–708.
- [7] James W Pennebaker Martha E Francis and Roger J Booth. 1993. *Linguistic Inquiry and Word Count*. Technical Report. Technical Report, Dallas, TX: Southern Methodist University.
- [8] Henry T Gilbert. 2010. *Persuasion detection in conversation*. Ph.D. Dissertation. Monterey, California. Naval Postgraduate School.
- [9] Teuvo Kohonen. 1985. Median strings. *Pattern Recognition Letters* 3, 5 (1985), 309–313.
- [10] Quoc V Le and Tomas Mikolov. 2014. Distributed Representations of Sentences and Documents.. In *ICML*, Vol. 14. 1188–1196.
- [11] Wei-Hao Lin, Theresa Wilson, Janyce Wiebe, and Alexander Hauptmann. 2006. Which side are you on?: identifying perspectives at the document and sentence levels. In *Proceedings of the Tenth Conference on Computational Natural Language Learning*. Association for Computational Linguistics, 109–116.
- [12] Stephanie Lukin, Pranav Anand, Marilyn Walker, and Steve Whittaker. 2017. Argument Strength is in the Eye of the Beholder: Audience Effects in Persuasion. (2017).
- [13] Gerald Marwell and David R Schmitt. 1967. Dimensions of compliance-gaining behavior: An empirical analysis. *Sociometry* (1967), 350–364.
- [14] Pedro Ortiz. 2010. *Machine learning techniques for persuasion detection in conversation*. Ph.D. Dissertation. Monterey, California. Naval Postgraduate School.
- [15] Sunghyun Park, Han Suk Shim, Moitrey Chatterjee, Kenji Sagae, and Louis-Philippe Morency. 2014. Computational analysis of persuasiveness in social multimedia: A novel dataset and multimodal prediction approach. In *Proceedings of the 16th International Conference on Multimodal Interaction*. ACM, 50–57.
- [16] John Platt. 1998. Sequential minimal optimization: A fast algorithm for training support vector machines. (1998).
- [17] Behjat Siddique, Dave Chisholm, and Ajay Divakaran. 2015. Exploiting Multimodal Affect and Semantics to Identify Politically Persuasive Web Videos. In *Proceedings of the 2015 ACM on International Conference on Multimodal Interaction*. ACM, 203–210.
- [18] Carlo Strapparava, Marco Guerini, and Oliviero Stock. 2010. Predicting Persuasiveness in Political Discourses.. In *LREC*.
- [19] Chenhao Tan, Lillian Lee, and Bo Pang. 2014. The effect of wording on message propagation: Topic-and author-controlled natural experiments on Twitter. *arXiv preprint arXiv:1405.1438* (2014).
- [20] 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*. International World Wide Web Conferences Steering Committee, 613–624.
- [21] Douglas Walton, Christopher Reed, and Fabrizio Macagno. 2008. *Argumentation schemes*. Cambridge University Press.
- [22] Joel Young, Craig H Martell, Pranav Anand, Pedro Ortiz, Henry Tucker Gilbert IV, et al. 2011. A Microtext Corpus for Persuasion Detection in Dialog.. In *Analyzing Microtext*.

Method	Blogs			ChangeMyView			Supreme Court			Political Speeches		
	P	R	F	P	R	F	P	R	F	P	R	F
SVM Baseline	<b>0.594</b>	0.132	0.216	0.511	0.107	0.176	<b>0.605</b>	0.051	0.094	0.454	0.038	0.070
NB+Tactic	0.361	<b>0.483</b>	0.413	0.309	<b>0.465</b>	0.371	0.319	<b>0.511</b>	0.393	0.267	<b>0.417</b>	0.325
NB+LDA	0.098	0.132	0.112	0.071	0.116	0.088	0.083	0.151	0.107	0.032	0.045	0.037
NB+Tactic+LDA	0.114	0.229	0.152	0.099	0.212	0.135	0.138	0.242	0.176	0.041	0.145	0.064
S-Doc2Vec	0.493	0.439	0.464	0.472	0.427	0.448	0.496	0.442	0.467	0.411	0.392	0.401
ParseTree	0.498	0.443	0.468	0.491	0.419	0.452	0.477	0.462	0.468	0.418	0.371	0.393
ParseTree+SP	0.531	0.470	<b>0.498</b>	<b>0.539</b>	0.448	<b>0.489</b>	0.521	0.483	<b>0.501</b>	<b>0.464</b>	0.405	<b>0.432</b>

**Table 8:** Comparison of results for the different methods, considering all the 14 tactics. Here, ParseTree is the first model proposed, ParseTree+SP stands for the parse-tree model with the synthetic prototype strings, NB stands for Naive Bayes, LDA stands for Latent Dirichlet Allocation, and S-Doc2Vec stands for the supervised version of the Doc2Vec method. NB+Tactic, NB+LDA, NB+Tactic+LDA are the features used by the authors in [1]

Method	Blogs			ChangeMyView			Supreme Court			Political Speeches		
	P	R	F	P	R	F	P	R	F	P	R	F
SVM Baseline	<b>0.741</b>	0.179	0.288	<b>0.721</b>	0.159	0.261	<b>0.763</b>	0.171	0.279	0.651	0.107	0.184
NB+Tactic	0.537	<b>0.672</b>	0.597	0.515	<b>0.645</b>	0.573	0.533	<b>0.656</b>	0.588	0.467	<b>0.599</b>	0.525
NB+LDA	0.133	0.285	0.181	0.111	0.262	0.156	0.137	0.282	0.184	0.051	0.203	0.082
NB+Tactic+LDA	0.169	0.437	0.244	0.144	0.416	0.214	0.154	0.441	0.228	0.097	0.354	0.152
S-Doc2Vec	0.732	0.599	0.659	0.705	0.589	0.642	0.723	0.627	0.672	0.648	0.533	0.585
ParseTree	0.701	0.603	0.648	0.692	0.589	0.636	0.716	0.614	0.661	0.633	0.539	0.582
ParseTree+SP	0.737	0.629	<b>0.679</b>	<b>0.721</b>	0.623	<b>0.668</b>	0.751	0.642	<b>0.692</b>	<b>0.661</b>	0.562	<b>0.607</b>

**Table 9:** Comparison of results for the different methods in a *binary* setting. Here, ParseTree is the first model proposed, ParseTree+SP stands for the parse-tree model with the synthetic prototype strings, NB stands for Naive Bayes, LDA stands for Latent Dirichlet Allocation, and S-Doc2Vec stands for the supervised version of the Doc2Vec method. NB+Tactic, NB+LDA, NB+Tactic+LDA are the features used by the authors in [1]