

Connotation Frames: A Data-Driven Investigation

Hannah Rashkin Sameer Singh Yejin Choi

Computer Science & Engineering
University of Washington

{hrashkin, sameer, yejin}@cs.washington.edu

Abstract

Through a particular choice of a predicate (e.g., “ x violated y ”), a writer can subtly connote a range of implied sentiments and presupposed facts about the entities x and y : (1) *writer’s perspective*: projecting x as an “antagonist” and y as a “victim”, (2) *entities’ perspective*: y probably dislikes x , (3) *effect*: something bad happened to y , (4) *value*: y is something valuable, and (5) *mental state*: y is distressed by the event.

We introduce *connotation frames* as a representation formalism to organize these rich dimensions of connotation using typed relations. First, we investigate the feasibility of obtaining connotative labels through crowdsourcing experiments. We then present models for predicting the connotation frames of verb predicates based on their distributional word representations and the interplay between different types of connotative relations. Empirical results confirm that connotation frames can be induced from various data sources that reflect how people use language and give rise to the connotative meanings. We conclude with analytical results that show the potential use of connotation frames for analyzing subtle biases in online news media.

1 Introduction

People commonly express their opinions through subtle and nuanced language (Thomas et al., 2006; Somasundaran and Wiebe, 2010). Often, through seemingly objective statements, the writer can influence the readers’ judgments toward an event and their participants. Even by choosing a particular predicate, the writer can indicate rich connotative information about the entities that interact through

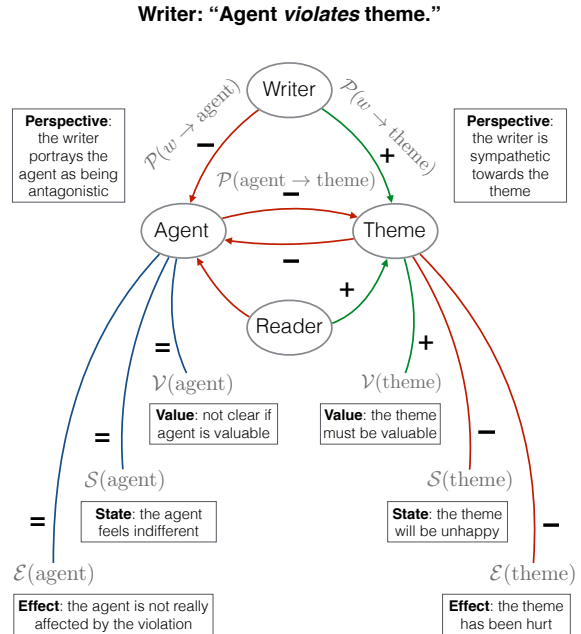


Figure 1: An example connotation frame of “violate” as a set of typed relations: perspective $\mathcal{P}(x \rightarrow y)$, effect $\mathcal{E}(x)$, value $\mathcal{V}(x)$, and mental state $\mathcal{S}(x)$.

- the predicate. Specifically, through a simple statement such as “ x violated y ”, the writer can convey:
- (1) **writer’s perspective**: the writer is projecting x as an “antagonist” and y as a “victim”, eliciting negative perspective from readers toward x (i.e., blaming x) and positive perspective toward y (supportive or sympathetic to y).
 - (2) **entities’ perspective**: y most likely feels negatively toward x as a result of being violated.
 - (3) **effect**: something bad happened to y .
 - (4) **value**: y is something valuable, since it does not make sense to violate something worthless. In other words, the writer is presupposing y ’s positive value as a fact.
 - (5) **mental state**: y is most likely unhappy about the outcome.¹

¹To be more precise, y is most likely in a negative mental state assuming it is an entity that can have a mental state.

Verb	Subset of Typed Relations		Example Sentences	L/R
suffer	$\mathcal{P}(w \rightarrow \text{agent}) = +$ $\mathcal{P}(w \rightarrow \text{theme}) = -$ $\mathcal{P}(\text{agent} \rightarrow \text{theme}) = -$	$\mathcal{E}(\text{agent}) = -$ $\mathcal{V}(\text{agent}) = +$ $\mathcal{S}(\text{agent}) = -$	The story begins in Illinois in 1987, when a 17-year-old girl suffered a botched abortion.	R
guard	$\mathcal{P}(w \rightarrow \text{agent}) = +$ $\mathcal{P}(w \rightarrow \text{theme}) = +$ $\mathcal{P}(\text{agent} \rightarrow \text{theme}) = +$	$\mathcal{E}(\text{theme}) = +$ $\mathcal{V}(\text{theme}) = +$ $\mathcal{S}(\text{theme}) = +$	In August, marshals guarded 25 clinics in 18 cities.	L
uphold	$\mathcal{P}(w \rightarrow \text{theme}) = +$ $\mathcal{P}(\text{agent} \rightarrow \text{theme}) = +$	$\mathcal{E}(\text{theme}) = +$ $\mathcal{V}(\text{theme}) = +$	A hearing is scheduled to make a decision on whether to uphold the clinic’s suspension.	R

Table 1: Example typed relations (perspective $\mathcal{P}(x \rightarrow y)$, effect $\mathcal{E}(x)$, value $\mathcal{V}(x)$, and mental state $\mathcal{S}(x)$), where w denotes the writer. Not all typed relations are shown due to space constraints. The example sentences demonstrate the usage of the predicates in left [L] or right [R] leaning news sources.

Even though the writer might not explicitly state any of the interpretations [1-5] above, the readers will be able interpret these intentions as a part of their comprehension. In this paper, we present an empirical study of how to represent and induce the connotative interpretations that can be drawn from a verb predicate, as illustrated above.

We introduce *connotation frames* as a representation framework to organize the rich dimensions of the implied sentiment and presupposed facts. Figure 1 shows an example of a connotation frame for the predicate *violate*. We define four different typed relations: $\mathcal{P}(x \rightarrow y)$ for perspective of x towards y , $\mathcal{E}(x)$ for effect on x , $\mathcal{V}(x)$ for value of x , and $\mathcal{S}(x)$ for mental state of x . These relationships can all be either positive (+), neutral (=), or negative (-).

Our work is the first study to investigate frames as a representation formalism for connotative meanings. This contrasts with previous computational studies and resource development for frame semantics, where the primary focus was almost exclusively on denotational meanings of language (Baker et al., 1998; Palmer et al., 2005). Our formalism draws inspirations from the earlier work of frame semantics, however, in that we investigate the connection between a word and the related world knowledge associated with the word (Fillmore, 1976), which is essential for the readers to interpret many layers of the implied sentiment and presupposed value judgments.

We also build upon the extensive amount of literature in sentiment analysis (Pang and Lee, 2008; Liu and Zhang, 2012), especially the recent emerging efforts on implied sentiment analysis (Feng et al., 2013; Greene and Resnik, 2009), entity-entity sentiment inference (Wiebe and Deng, 2014), opinion role induction (Wiegand and Ruppenhofer, 2015), and effect analysis (Choi and Wiebe, 2014).

However, our work is the first to organize aspects of the connotative information into coherent frames.

More concretely, our contributions are threefold: (1) a new formalism, model, and annotated dataset for studying the connotation frames from large-scale natural language data and statistics, (2) data-driven insights into the dynamics among different typed relations within each frame, and (3) an analytic study to show the potential use of connotation frames for analyzing subtle biases in journalism.

The rest of the paper is organized as follows: in §2, we provide the definitions and data-driven insights for connotation frames. In §3, we introduce models for inducing the connotation frames, followed by empirical results, annotation studies, and analysis on news media in §4. We discuss related work in §5 and conclude in §6.

2 Connotation Frame

Given a predicate v , we define a connotation frame $\mathcal{F}(v)$ as a collection of typed relations and their polarity assignments: (i) **perspective** $\mathcal{P}^v(x_i \rightarrow x_j)$: whether the predicate v implies directed sentiment from the entity x_i to the entity x_j , (ii) **value** $\mathcal{V}^v(x_i)$: whether x_i is presupposed to be valuable by the predicate v , (iii) **effect** $\mathcal{E}^v(x_i)$: whether the event denoted by the predicate v is good or bad for the entity x_i , and (iv) **mental state** $\mathcal{S}^v(x_i)$: the likely mental state of the entity x_i as a result of the event. We assume that each typed relation can have one of the three connotative polarities $\in \{+, -, =\}$, i.e., positive, negative, or neutral. Our goal in this paper is to focus on the general connotation of the predicate considered out of context. We leave contextual interpretation of connotation as future work.

Table 1 shows examples of connotation frame relations for the verbs *suffer*, *guard*, and *uphold*, along with example sentences. For instance, for the verb *suffer*, the writer is likely to have a positive

Verb	x 's role	$\mathcal{P}(w \rightarrow \cdot)$	Left-leaning Sources	Right-leaning Sources
<i>accuse</i>	agent	-	Putin, Progressives, Limbaugh, Gingrich	activist, U.S., protestor, Chavez
	theme	+	official, rival, administration, leader	Romney, Iran, Gingrich, regime
<i>attack</i>	agent	-	McCain, Trump, Limbaugh	Obama , campaign, Biden, Israel
	theme	+	Gingrich, Obama , policy	citizen, Zimmerman
<i>criticize</i>	agent	-	Ugandans, rival, Romney, Tyson	Britain, passage, Obama , Maddow
	theme	+	Obama , Allen, Cameron, Congress	Pelosi, Romey, GOP, Republicans

Table 2: Media Bias in Connotation Frames: Obama, for example, is portrayed as someone who *attacks* or *criticizes* others by the right-leaning sources, whereas the left-leaning sources portray Obama as the victim of harsh acts like “attack” and “criticize”.

perspective towards the agent (e.g., being supportive or sympathetic toward the “17-year-old girl” in the example shown on the right) and a negative perspective towards the theme (e.g., being negative towards ‘botched abortion”).

2.1 Data-driven Motivation

Since the meaning of language is ultimately contextual, the exact connotation will vary depending on the context of each utterance. Nonetheless, there still are common shifts or biases in the connotative polarities, as we found from two data-driven analyses.

First, we looked at words from the Subjectivity Lexicon (Wilson et al., 2005) that are used in the argument positions of a small selection of predicates in Google Syntactic N-grams (Goldberg and Orwant, 2013). For this analysis, we assumed that the agent is the word in the subject position while the theme is the word in the object position. We found 64% of the words in the agent role of *suffer* are positive, and 94% of the words in the theme role are negative, which is consistent with the polarities of the writer’s perspective towards these arguments, as shown in Table 1. For *guard*, 57% of the agents and 76% of the themes are positive, and in the case of *uphold*, 56% of the agents and 72% of the themes are positive.

We also investigated how media bias can potentially be analyzed through connotation frames. From the Stream Corpus 2014 dataset (KBA, 2014), we selected all articles from news outlets with known political biases,² and compared how they use polarised words such as “accuse”, “attack”, and “criticize” differently in light of $\mathcal{P}(w \rightarrow agent)$ and $\mathcal{P}(w \rightarrow theme)$ relations of the connotation frames. Table 2 shows interesting contrasts. Obama, for example, is frequently portrayed as

someone who *attacks* or *criticizes* others according to the right-leaning sources, whereas the left-leaning sources portray Obama as the victim of harsh acts like “attack” or “criticize”.³ Furthermore, by knowing the perspective relationships $\mathcal{P}(w \rightarrow x_i)$ associated with a predicate, we can make predictions about how the left-leaning and right-leaning sources feel about specific people or issues. For example, because left-leaning sources frequently use McCain, Trump, and Limbaugh in the agent position of attack, we might predict that these sources have a negative sentiment towards these entities.

2.2 Dynamics between Typed Relations

Given a predicate, the polarity assignments of typed relations are interdependent. For example, if the writer feels positively towards the agent but negatively towards the theme, then it is likely that the agent and the theme do not feel positively towards each other. This insight is related to that of Wiebe and Deng (2014), but differs in that the polarities are predicate-specific and do not rely on knowledge of prior sentiment towards the arguments, themselves. This and other possible interdependencies are summarized in Table 3. These interdependencies serve as general guidelines of what properties we expect to depend on one another, especially in the case where the polarities are non-neutral. We will promote these internal consistencies in our factor graph model (§3) as soft constraints.

There also exist other interdependencies that we will use to simplify our task. First, the directed sentiments between the agent and the theme are likely to be reciprocal, or at least do not directly conflict with + and - simultaneously. Therefore, we assume that $\mathcal{P}(x_i \rightarrow x_j) = \mathcal{P}(x_j \rightarrow x_i) = \mathcal{P}(x_i \leftrightarrow x_j)$, and we only measure for these binary

²The articles come from 30 news sources indicated by others as exhibiting liberal or conservative leanings (Mitchell et al., 2014; Center for Media and Democracy, 2013; Center for Media and Democracy, 2012; HWC Library, 2011)

³That is, even if someone truly deserves criticism from Obama, left-leaning sources would choose slightly different wordings to avoid a potentially harsh portrayal of Obama.

Perspective Triad: If argument x_i is positive towards x_j , and x_j is positive towards x_k , then we expect x_i is also positive towards x_k . Similar dynamics hold for the negative case.

$$\mathcal{P}_{w \rightarrow x_i} = \neg (\mathcal{P}_{w \rightarrow x_j} \oplus \mathcal{P}_{x_i \rightarrow x_j})$$

Perspective – Effect: If a predicate has a positive effect on one argument, then we expect that the interaction between the arguments was positive. Similar dynamics hold for the negative case.

$$\mathcal{E}_{x_i} = \mathcal{P}_{x_j \rightarrow x_i}$$

Perspective – Value: If argument x_i is presupposed as valuable, then we expect that the writer also views x_i positively. Similar dynamics hold for the negative case.

$$\mathcal{V}_{x_i} = \mathcal{P}_{w \rightarrow x_i}$$

Effect – Mental State: If the predicate has a positive effect on an argument x_i , then we expect that x_i will gain a positive mental state. Similar dynamics hold for the negative case.

$$\mathcal{S}_{x_i} = \mathcal{E}_{x_i}$$

Table 3: Potential Dynamics among Typed Relations: we propose models that parameterize these dynamics using log-linear models (frame-level model in §3).




Figure 2: A factor graph for predicting the polarities of the typed relations that define a connotation frame for a given verb predicate. The factor graph also includes unary factors (ψ_{emb}), which we left out for brevity.

relationships going in one direction. In addition, we assume the *predicted*⁴ perspective from the reader r to an argument, $\mathcal{P}(r \rightarrow x)$, is likely to be the same as the *implied* perspective from the writer w to the same argument, $\mathcal{P}(w \rightarrow x)$. So, we only try to learn the perspective of the writer. Lifting these assumptions will be future work.

For simplicity, our work only explores verb predicates and focuses on the polarities involving the agent and the theme roles, which we will refer to as a and t . We will assume that these roles are correlated to the subject and object positions.

3 Modeling Connotation Frames

Our task is essentially that of lexicon induction (Akkaya et al., 2009; Feng et al., 2013) in that we want to induce the connotation frames of previ-

⁴Surely different readers can and will form varying opinions after reading the same text. Here we concern with the most likely perspective of the general audience, as a result of reading the text.

ously unseen verbs. For each verb predicate, we infer a connotation frame composed of 9 relationship aspects that represent: *perspective* $\{\mathcal{P}(w \rightarrow t), \mathcal{P}(w \rightarrow a), \mathcal{P}(a \rightarrow t)\}$, *effect* $\{\mathcal{E}(t), \mathcal{E}(a)\}$, *value* $\{\mathcal{V}(t), \mathcal{V}(a)\}$, and *mental state* $\{\mathcal{S}(t), \mathcal{S}(a)\}$ polarities, where w, a, t denote the writer, the agent, and the theme, respectively.

We propose two models: an aspect-level model that makes the prediction for each typed relation independently based on the distributional representation of the context in which the predicate commonly appears (§3.1), and a frame-level model that makes the prediction over the connotation frame collectively in consideration of the dynamics between typed relations (§3.2).

3.1 Aspect-Level

Our aspect-level model predicts labels for each of these typed relations separately. As input, we use the 300-dimensional dependency-based word embeddings from Levy and Goldberg (2014). For each aspect, there is a separate MaxEnt (maximum entropy) classifier used to predict the label of that aspect on a given word-embedding, which is treated as a 300 dimensional input vector to the classifier. The MaxEnt classifiers learn their weights using LBFGS on the training data examples with re-weighting of samples to maximize for the best average F1 score.

3.2 Frame-Level

Next we present a factor graph model (Figure 2) of the connotation frames that parameterizes the dynamics between typed relations. Specifically, for each verb predicate, the factor graph contains 9 nodes representing the different aspects of the connotation frame involving the writer (w), the agent (a), and the theme (t). All these variables take polarity values from the set $\{-, =, +\}$.

We define $\mathbf{Y}_i := \{\mathcal{P}_{wt}, \mathcal{P}_{wa}, \mathcal{P}_{at}, \mathcal{E}_t, \mathcal{E}_a, \mathcal{V}_t, \mathcal{V}_a, \mathcal{S}_t, \mathcal{S}_a\}$ as the set of relational aspects for the i^{th} verb predicate. The factor graph for \mathbf{Y}_i , is illustrated in Figure 2, and we describe the factor potentials, ψ , in detail in the rest of this section. The probability of an assignment of polarities to the nodes in \mathbf{Y}_i is:

$$\begin{aligned} P(\mathbf{Y}_i) \propto & \psi_{PV}(\mathcal{P}_{wa}, \mathcal{V}_a) \psi_{PV}(\mathcal{P}_{wt}, \mathcal{V}_t) \\ & \psi_{PE}(\mathcal{P}_{at}, \mathcal{E}_a) \psi_{PE}(\mathcal{P}_{at}, \mathcal{E}_t) \\ & \psi_{ES}(\mathcal{E}_a, \mathcal{S}_a) \psi_{ES}(\mathcal{E}_t, \mathcal{S}_t) \\ & \psi_{PT}(\mathcal{P}_{wt}, \mathcal{P}_{wa}, \mathcal{P}_{at}) \prod_{y \in \mathbf{Y}_i} \psi_{emb}(y) \end{aligned}$$

Embedding Factors We include unary factors on all nodes to represent the results of the aspect-level classifier. Incorporating this knowledge as factors, as opposed to *fixing* the variables as observed, affords us the flexibility of representing noise in the labels as soft evidence. The potential function ψ_{emb} is a log-linear function of a feature vector \mathbf{f} , which is a one-hot feature vector representing the polarity of a node (+, -, or =). For example, with the node representing the value of the theme (\mathcal{V}_t):

$$\psi_{emb}(\mathcal{V}_t) = e^{\theta_{\mathcal{V}_t} \cdot \mathbf{f}(\mathcal{V}_t)}$$

The potential ψ_{emb} is defined similarly for the remaining eight nodes.

Weights θ are learned in a piecewise likelihood manner (Sutton and McCallum, 2009) for each factor independently using stochastic gradient descent (SGD) over the training data.

Interdependency Factors We include interdependency factors to promote the properties defined by the dynamics between relations (§2.2). The potentials for Perspective Triad, Perspective-Value, Perspective-Effect, and Effect-State Relationships (ψ_{PT} , ψ_{PV} , ψ_{PE} , ψ_{ES} respectively) are all defined using log-linear functions of one-hot feature vectors that encode the combination of polarities of the neighboring nodes. Thus the potential ψ_{PT} is:

$$\psi_{PT}(\mathcal{P}_{wt}, \mathcal{P}_{wa}, \mathcal{P}_{at}) = e^{\theta_{PT} \cdot \mathbf{f}(\mathcal{P}_{wt}, \mathcal{P}_{wa}, \mathcal{P}_{at})}$$

And we define the potentials for ψ_{PV} , ψ_{PE} , and ψ_{ES} for nodes pertaining to the agent as:

$$\begin{aligned} \psi_{PV}(\mathcal{P}_{wa}, \mathcal{V}_a) &= e^{\theta_{PV,a} \cdot \mathbf{f}(\mathcal{P}_{wa}, \mathcal{V}_a)} \\ \psi_{PE}(\mathcal{P}_{at}, \mathcal{E}_a) &= e^{\theta_{PE,a} \cdot \mathbf{f}(\mathcal{P}_{at}, \mathcal{E}_a)} \\ \psi_{ES}(\mathcal{E}_a, \mathcal{S}_a) &= e^{\theta_{ES,a} \cdot \mathbf{f}(\mathcal{E}_a, \mathcal{S}_a)} \end{aligned}$$

and we define the potentials for the theme nodes similarly. As with the unary seed factors, weights θ are learned using SGD over training data.

Belief Propagation We use belief propagation to induce the connotation frames of previously unseen verbs. In the belief propagation algorithm, messages are iteratively passed between the nodes to their neighboring factors. Each message μ , containing a scalar for each value $x \in \{-, =, +\}$, is defined from each node v to a neighboring factor f as follows:

$$\mu_{v \rightarrow f}(x) \propto \prod_{f^* \in N(v) \setminus a} \mu_{f^* \rightarrow v}(x)$$

and from each factor a to a neighboring node v as:

$$\mu_{f \rightarrow v} \propto \sum_{x', x'_v = x} \psi(x') \prod_{v^* \in N(f) \setminus v} \mu_{v^* \rightarrow f}(x'_{v^*})$$

Our factor graph does not contain any loops, so we are able to perform exact inference by choosing a root node and performing message passing from the leaves to the root and back to the leaves. At the conclusion of message passing, the probability of a specific polarity associated with node v being equal to x is proportional to $\prod_{f \in N(v)} \mu_{f \rightarrow v}(x)$.

4 Experiments

We first describe crowd-sourced annotations (§4.1), then present the empirical results of predicting connotation frames (§4.2), and conclude with qualitative analysis of a large corpus (§4.3).

4.1 Data and Crowdsourcing

In order to understand how humans interpret connotation frames, we designed an Amazon Mechanical Turk (AMT) annotation study. We gathered a set of transitive verbs commonly used in the New York Times corpus (Sandhaus, 2008), selecting the 2400 verbs that are used more than 200 times in the corpus. Of these, AMT workers annotated the 1000 most frequently used verbs.

Annotation Design In a pilot annotation experiment, we found that annotators have difficulty thinking about subtle connotative polarities when shown predicates without any context. Therefore, we designed the AMT task to provide a generic context as follows. We first split each verb predicate into 5 separate tasks that each gave workers a different generic sentence using the verb. To create generic sentences, we used Google Syntactic

N-grams (Goldberg and Orwant, 2013) to come up with a frequently seen Subject-Verb-Object tuple which served as a simple three-word sentence with generic arguments.⁵ For each of the 5 sentences, we asked 3 annotators to answer questions like “How do you think the agent feels about the event described in this sentence?” In total, each verb has 15 annotations aggregated over 5 different generic sentences containing the verb.

In order to help the annotators, some of the questions also allowed annotators to choose sentiment using additional classes for “positive or neutral” or “negative or neutral” for when they were less confident but still felt like a sentiment might exist. When taking inter-annotator agreement, we count “positive or neutral” as agreeing with either “positive” or “neutral” classes.

Annotator agreement Table 4 shows agreements and data statistics. The non-conflicting (NC) agreement only counts opposite polarities as disagreement.⁶ From this study, we can see that non-expert annotators are able to see these sort of relationships based on their understanding of how language is used. From the NC agreement, we see that annotators do not frequently choose completely opposite polarities, indicating that even when they disagree, their disagreements are based on the degree of connotations rather than the polarity itself. The average Krippendorff alpha for all of the questions posed to the workers is 0.25, indicating stronger than random agreement. Considering the subtlety of the implicit sentiments that we are asking them to annotate, it is reasonable that some annotators will pick up on more nuances than others. Overall, the percent agreement is encouraging that the connotative relationships are visible to human annotators.

Aggregating Annotations We aggregated over crowdsourced labels (fifteen annotations per verb) to create a polarity label for each aspect of a verb.⁷

Final distributions of the aggregated labels are included in the right-hand columns of Table 4. Notably, the distributions are skewed toward positive and neutral labels. The most skewed connotation

⁵Because Google Syntactic N-grams only provide dependency types and do not provide semantic roles, we approximate the agent and theme as the subject and the object respectively.

⁶Annotators were asked yes/no questions related to Value, so this does not have a corresponding NC agreement score.

⁷We take the average to obtain scalar value between $[-1, 1]$ for each aspect of a verb’s connotation frame. For simplicity, we cutoff the ranges of negative, neutral and positive polarities as $[-1, -0.25]$, $[-0.25, 0.25]$ and $(0.25, 1]$, respectively.

Aspect	% Agreement		Distribution	
	Strict	NC	% +	% -
$\mathcal{P}(w \rightarrow t)$	75.6	95.6	36.6	4.6
$\mathcal{P}(w \rightarrow a)$	76.1	95.5	47.1	7.9
$\mathcal{P}(a \rightarrow t)$	70.4	91.9	45.8	5.0
$\mathcal{E}(t)$	52.3	94.6	50.3	20.24
$\mathcal{E}(a)$	53.5	96.5	45.1	4.7
$\mathcal{V}(t)$	65.2	-	78.64	2.7
$\mathcal{V}(a)$	71.9	-	90.32	1.4
$\mathcal{S}(t)$	79.9	98.0	12.8	14.5
$\mathcal{S}(a)$	70.4	92.5	50.72	8.6

Table 4: Label Statistics: % Agreement refers to pairwise inter-annotator agreement. The strict agreement counts agreement over 3 classes (“positive or neutral” was counted as agreeing with either + or neutral), while non-conflicting (NC) agreement also allows agreements between neutral and -/+ (no direct conflicts). Distribution shows the final class distribution of -/+ labels created by averaging annotations.

frame aspect is the value $\mathcal{V}(x)$ which tends to be positive, especially for the agent argument. This makes some intuitive sense since, as the agent actively causes the predicate event to occur, they most likely have some intrinsic potential to be valuable. An example of a verb where the agent was labelled with negative value is “contaminate”. In the most generic case, the writer is using “contaminate” to frame the agent as being worthless (and even harmful) with regards to the other event participants. For example, in the sentence “his touch contaminated the food,” it is clear that the writer presupposes “his touch” to be of negative value in how it impacts the rest of the event.

4.2 Connotation Frame Prediction

Using the crowdsourced labels, we randomly divide the annotated verbs into training, dev, and held-out test sets of equal size (300 verbs each). For evaluation we measure average accuracy and F1 score of induced labels for the 9 different connotation frame relationship types for which we have annotations: $\mathcal{P}(w \rightarrow t)$, $\mathcal{P}(w \rightarrow a)$, $\mathcal{P}(a \rightarrow t)$, $\mathcal{V}(t)$, $\mathcal{V}(a)$, $\mathcal{E}(t)$, $\mathcal{E}(a)$, $\mathcal{S}(t)$, and $\mathcal{S}(a)$, where w refers to the writer, a to the agent, and t to the theme.

Baselines To show the non-trivial challenge of learning Connotation Frames, we include a simple majority-class baselines. The MAJORITY classifier assigns each of the 9 relationships the label of the majority of that relationship type found in the training data. Some of these relationships (in particular, the Value of agent/theme) have skewed distributions, so we expect this classifier to achieve

a much higher accuracy than random but a much lower overall F1 score.

Additionally, we add a GRAPH PROP baseline that is comparable to algorithms like graph propagation or label propagation which are often used for (sentiment) lexicon induction. We use a factor graph with nodes representing the polarity of each typed relation for each verb. Binary factors connect nodes representing a particular type of relation for two similar verbs (e.g. $\mathcal{P}(w \rightarrow t)$ for verbs *persuade* and *convince*). These binary factors have hand-tuned potentials that are proportional to the cosine similarity of the verbs’ embeddings, encouraging similar verbs to have the same polarity for the various relational aspects. We use words in the training data as the seed set and use loopy belief propagation to propagate polarities from known nodes to the unknown relationships.

Finally, we use a 3-NEAREST NEIGHBOR baseline that labels relationships for a verb based on the predicate’s 300-dimensional word embedding representation, using the same embeddings as in our aspect-level. 3-NEAREST NEIGHBOR labels each verb using the polarities of the three closest verbs found in the training set. The most similar verbs are determined using the cosine similarity between word embeddings.

Results As shown in Table 5, aspect-level and frame-level models consistently outperform all three baselines — MAJORITY, 3-NN, GRAPH PROP in the development set across the different types of relationships. In particular, the improved F1 scores show that these models are able to perform better across all three classes of labels even in the most skewed cases. The frame-level model also frequently improves the F1 scores of the labels from what they were in the aspect-level model. The summarized comparison of the classifiers’ performance test set is shown in Table 6. As with the development set, aspect-level and frame-level are both able to outperform the baselines. Furthermore, the frame-level formulation is able to make improvement over the results of the aspect-level classification, indicating that the modelling of inter-dependencies between relationships did help correct some of the mistakes made.

One point of interest about the frame-level results is whether the learned weights over the consistency factors match our initial intuitions about inter-dependencies between relationships. The weights learned in our algorithm do tell us something in-

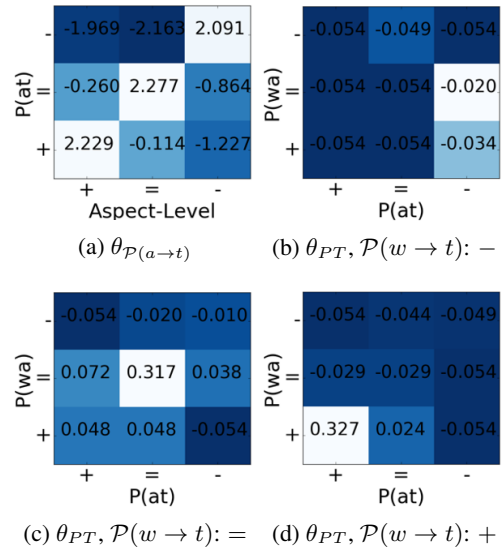


Figure 3: Learned weights, θ , of embedding factor for the perspective of agent to theme (3a) and the weights the perspective triad (PT) factor (3b-3d). Lighter shades are for weights that are more positive, whereas dark blue is more negative.

teresting about the degree to which these inter-dependencies are actually found in our data.

We show the heat maps for some of the learned weights in Figure 3. In 3a, we show the weights of one of the embedding factors, and how the node’s polarities are more strongly weighted when they match the aspect-level output. In the rest of the figure, we show the weights for the other perspective relationships when $\mathcal{P}(w \rightarrow t)$ is negative (3b), neutral (3c), and positive (3d), respectively. Based on the expected interdependencies, when $\mathcal{P}(w \rightarrow t) : -$, the model should favor $\mathcal{P}(w \rightarrow a) \neq \mathcal{P}(a \rightarrow t)$ and when $\mathcal{P}(w \rightarrow t) : +$, the model should favor $\mathcal{P}(w \rightarrow a) = \mathcal{P}(a \rightarrow t)$. Our model does, in fact, learn a similar trend, with slightly higher weights along these two diagonals in the maps 3b and 3d. Interestingly, when $\mathcal{P}(w \rightarrow t)$ is neutral, weights slightly prefer for the other two perspectives to resemble one another, but with highest weights being when other perspectives are also neutral.

4.3 Analysis of a Large News Corpus

Using the connotation frame, we present measured implied sentiment in online journalism.

Data From the Stream Corpus (KBA, 2014), we select 70 million news articles. We extract subject-verb-object relations for this subset using the direct dependencies between noun phrases and verbs as identified by the BBN Serif system, obtaining 1.2 billion unique tuples of the form

Aspect	Algorithm	Acc.	Avg F ₁
$\mathcal{P}(w \rightarrow t)$	Majority	56.52	24.07
	Graph Prop	59.53	50.20
	3-nn	62.88	47.93
	Aspect-Level	67.56	56.18
	Frame-Level	67.56	56.18
$\mathcal{P}(w \rightarrow a)$	Majority	49.83	22.17
	Graph Prop	52.84	42.93
	3-nn	55.18	45.88
	Aspect-Level	60.54	60.72
	Frame-Level	61.87	63.07
$\mathcal{P}(a \rightarrow t)$	Majority	49.83	22.17
	Graph Prop	52.17	46.57
	3-nn	56.52	52.94
	Aspect-Level	63.21	61.70
	Frame-Level	63.88	62.56
$\mathcal{E}(t)$	Majority	48.83	21.87
	Graph Prop	54.85	51.40
	3-nn	55.18	51.53
	Aspect-Level	64.21	63.63
	Frame-Level	65.22	64.67
$\mathcal{E}(a)$	Majority	49.83	22.17
	Graph Prop	52.17	35.56
	3-nn	54.85	42.63
	Aspect-Level	62.54	53.82
	Frame-Level	63.88	56.81
$\mathcal{V}(t)$	Majority	79.60	29.55
	Graph Prop	71.91	35.10
	3-nn	76.25	39.09
	Aspect-Level	75.92	45.45
	Frame-Level	76.25	48.13
$\mathcal{V}(a)$	Majority	89.30	31.45
	Graph Prop	84.62	38.82
	3-nn	85.62	38.45
	Aspect-Level	87.96	48.06
	Frame-Level	87.96	48.06
$\mathcal{S}(t)$	Majority	71.91	27.89
	Graph Prop	69.90	55.57
	3-nn	72.91	59.26
	Aspect-Level	81.61	72.85
	Frame-Level	81.61	72.85
$\mathcal{S}(a)$	Majority	50.84	22.47
	Graph Prop	48.83	35.40
	3-nn	54.85	45.51
	Aspect-Level	61.54	53.88
	Frame-Level	61.54	53.88

Table 5: Detailed breakdown of results on the development set using accuracy and average F1 over the three class labels (+,-,=).

Algorithm	Acc.	Avg F ₁
Graph Prop	58.81	41.46
3-nn	63.71	47.30
Aspect-Level	67.93	53.17
Frame-Level	68.26	53.50

Table 6: Performance on the test set. Results are averaged over the different aspects.




Figure 4: Average sentiment of Democrats and Republicans (as agents) to selected nouns (as their themes), aggregated over a large corpus using the learned lexicon (§4.2). The line indicates identical sentiments, i.e. Republicans are more positive towards the nouns that are above the line.

(*url,subject,verb,object,count*). We also extract tuples from news articles from the Annotated English Gigaword Corpus (Napoles et al., 2012), which contains nearly 10 million articles, resulting in an additional 120 million unique tuples.

Estimating Entity Polarities Using connotation frames, we can also measure entity-to-entity sentiment at a large scale. Figure 4, for example, presents the polarity of entities “Democrats” and “Republicans” towards a selected set of nouns, by computing the average estimated $\mathcal{P}(a \rightarrow t)$ polarity (using our frame-level output) over triples where one of these entities appears as part of the phrase in the agent role (e.g. “Democrats” or “Republican party”). Apart from nouns that both entities are positive (“business”, “constitution”) or negative (“the allegations”, “veto threat”) towards, we can also see interesting examples in which Democrats feel more positively (below the line: “nancy pelosi”, “unions”, “gun control”, etc.) and ones where Republicans feel more positive (“the pipeline”, “gop leadership”, “budget cuts”, etc.). Also, both entities are neutral towards “idea” and “the proposal”, which probably owes to the fact that ideas or proposals can be good or bad for either entity depending on the context.

5 Related Work

Most prior work on sentiment lexicons focused on the overall polarity of words without taking

into account their semantic arguments (Wilson et al., 2005; Baccianella et al., 2010; Wiebe et al., 2005; Velikovich et al., 2010; Kaji and Kitsuregawa, 2007; Kamps et al., 2004; Takamura et al., 2005; Adreevskaia and Bergler, 2006). Several recent studies began exploring more specific and nuanced aspects of sentiment such as connotation (Feng et al., 2013), good and bad effects (Choi and Wiebe, 2014), and evoked sentiment (Mohammad and Turney, 2010). Drawing inspirations from them, we present connotation frames as a unifying representation framework to encode the rich dimensions of implied sentiment, presupposed value judgements, and effect evaluation, and propose a factor graph formulation that captures the interplay among different types of connotation relations.

Goyal et al. (2010a; 2010b) investigated how characters (protagonists, villains, victims) in children’s stories are affected by certain predicates, which is related to the effect relations studied in this work. While Klenner et al. (2014) similarly investigated the relation between the polarity of the verbs and arguments, our work introduces new perspective types and proposes a unified representation and inference model. Wiegand and Ruppenhofer (2015) also looked at perspective-based relationships induced by verb predicates with a focus on opinion roles. Building on this concept, our framework also incorporates information about the perspectives’ polarities as well as information about other typed relations. There have been growing interests for modeling framing (Greene and Resnik, 2009; Hasan and Ng, 2013), biased language (Recasens et al., 2013) and ideology detection (Yano et al., 2010). All these tasks are relatively less studied, and we hope our connotation frame lexicon will be useful for them.

Sentiment inference rules have been explored by the recent work of Wiebe and Deng (2014) and Deng and Wiebe (2014). The focus of their work was on general inference rules that are not predicate-specific. In contrast, our work focuses on the notion that connotative polarities can be determined directly from the predicate, rather than partial knowledge of the arguments or the context in which it is being used. In brief, we make a novel conceptual connection between inferred sentiments and frame semantics, organized as connotation frames, and present a unified model that integrates different aspects of the connotation frames.

Finally, in a broader sense, what we study as connotation frames draws a connection to schema

and script theory (Schank and Abelson, 1975). Unlike prior work that focused on directly observable actions (Chambers and Jurafsky, 2009; Frermann et al., 2014; Bethard et al., 2008), we focus on implied sentiments that are framed by predicate verbs.

6 Conclusion

In this paper, we presented a novel system of connotative frames that define a set of implied sentiment and presupposed facts for a predicate. Our work also empirically explores different methods of inducing and modelling these connotation frames, incorporating the interplay between relations within frames. Our work suggests new research avenues on learning connotation frames, and their applications to deeper understanding of social and political discourse. All the learned connotation frames and annotations are available at <http://homes.cs.washington.edu/~hrashkin/connframe.html>.

Acknowledgements

We thank the anonymous reviewers for many insightful comments. We also thank members of UW NLP for discussions and support. This material is based upon work supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE-1256082, in part by NSF grants IIS-1408287, IIS-1524371, and gifts by Google and Facebook.

References

- Alina Adreevskaia and Sabine Bergler. 2006. Mining wordnet for fuzzy sentiment: Sentiment tag extraction from wordnet glosses. In *11th Conference of the European Chapter of the Association for Computational Linguistics*, pages 209–216.
- Cem Akkaya, Janyce Wiebe, and Rada Mihalcea. 2009. Subjectivity word sense disambiguation. In *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing*, volume 2, pages 190–199.
- Stefano Baccianella, Andrea Esuli, and Fabrizio Sebastiani. 2010. Sentiwordnet 3.0: An enhanced lexical resource for sentiment analysis and opinion mining. In *Proceedings of the Seventh conference on International Language Resources and Evaluation (LREC’10)*.
- Collin F Baker, Charles J Fillmore, and John B Lowe. 1998. The berkeley framenet project. In *Proceed-*

- ings of the 17th international conference on Computational linguistics, volume 1, pages 86–90.
- Steven Bethard, William J Corvey, Sara Klingsenstein, and James H Martin. 2008. Building a corpus of temporal-causal structure. In *Proceedings of the Sixth International Conference on Language Resources and Evaluation (LREC'08)*.
- Center for Media and Democracy. 2012. Sourcewatch: Conservative news outlets. http://www.sourcewatch.org/index.php/Conservative_news_outlets.
- Center for Media and Democracy. 2013. Sourcewatch: Liberal news outlets. http://www.sourcewatch.org/index.php/Liberal_news_outlets.
- Nathanael Chambers and Dan Jurafsky. 2009. Unsupervised learning of narrative schemas and their participants. In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP*, volume 2 of *ACL '09*, pages 602–610.
- Yoonjung Choi and Janyce Wiebe. 2014. +/-effectwordnet: Sense-level lexicon acquisition for opinion inference. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1181–1191. Association for Computational Linguistics, October.
- Lingjia Deng and Janyce Wiebe. 2014. Sentiment propagation via implicature constraints. In *Proceedings of the Conference of the European Chapter of the Association for Computational Linguistics (EACL)*.
- Song Feng, Jun Seok Kang, Polina Kuznetsova, and Yejin Choi. 2013. Connotation lexicon: A dash of sentiment beneath the surface meaning. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (ACL)*, volume 1, pages 1774–1784. Association for Computational Linguistics.
- Charles J. Fillmore. 1976. Frame semantics and the nature of language. In *In Annals of the New York Academy of Sciences: Conference on the Origin and Development of Language and Speech*, volume 280, pages 2032.
- Lea Frermann, Ivan Titov, and Manfred Pinkal. 2014. A hierarchical bayesian model for unsupervised induction of script knowledge. In *Proceedings of the Conference of the European Chapter of the Association for Computational Linguistics*.
- Yoav Goldberg and Jon Orwant. 2013. A dataset of syntactic-ngrams over time from a very large corpus of english books. In *Second Joint Conference on Lexical and Computational Semantics (*SEM)*, volume 1, pages 241–247, June.
- Amit Goyal, Ellen Riloff, and Hal Daumé, III. 2010a. Automatically producing plot unit representations for narrative text. In *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing*, Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing (EMNLP), pages 77–86.
- Amit Goyal, Ellen Riloff, Hal Daumé III, and Nathan Gilbert. 2010b. Toward plot units: Automatic affect state analysis. In *Proceedings of HLT/NAACL Workshop on Computational Approaches to Analysis and Generation of Emotion in Text (CAET)*.
- Stephan Greene and Philip Resnik. 2009. More than words: Syntactic packaging and implicit sentiment. In *Proceedings of Human Language Technologies: The 2009 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pages 503–511.
- Kazi Saidul Hasan and Vincent Ng. 2013. Frame semantics for stance classification. *Proceedings of the Seventeenth Conference on Computational Natural Language Learning (CONLL)*, pages 124–132.
- HWC Library. 2011. Consider the Source: A Resource Guide to Liberal, Conservative, and Nonpartisan Periodicals. www.ccc.edu/colleges/washington/departments/Documents/PeriodicalsPov.pdf. Compiled by HWC Librarians in January 2011.
- Nobuhiro Kaji and Masaru Kitsuregawa. 2007. Building lexicon for sentiment analysis from massive collection of html documents. In *Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL)*, pages 1075–1083.
- Jaap Kamps, Maarten Marx, Robert J Mokken, and Maarten De Rijke. 2004. Using wordnet to measure semantic orientations of adjectives. In *Proceedings of the Fourth International Conference on Language Resources and Evaluation (LREC'04)*, volume 4, pages 1115–1118.
- TREC KBA. 2014. Knowledge Base Acceleration Stream Corpus. <http://trec-kba.org/kba-stream-corpus-2014.shtml>.
- Manfred Klenner, Michael Amsler, and Nora Hollenstein. 2014. Verb polarity frames: a new resource and its application in target-specific polarity classification. In *Proceedings of KONVENS 2014*, pages 106–115.
- Omer Levy and Yoav Goldberg. 2014. Dependency-based word embeddings. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 302–308.
- Bing Liu and Lei Zhang. 2012. A survey of opinion mining and sentiment analysis. In *Mining text data*, pages 415–463. Springer.

- Amy Mitchell, Jeffrey Gottfried, Jocelyn Kiley, and Katerina Eva Matsa. 2014. Political Polarization & Media Habits. www.journalism.org/2014/10/21/political-polarization-media-habits/. Produced by Pew Research Center in October, 2014.
- Saif M Mohammad and Peter D Turney. 2010. Emotions evoked by common words and phrases: Using mechanical turk to create an emotion lexicon. In *Proceedings of the NAACL HLT 2010 Workshop on Computational Approaches to Analysis and Generation of Emotion in Text*, pages 26–34. Association for Computational Linguistics.
- Courtney Napoles, Matthew Gormley, and Benjamin Van Durme. 2012. Annotated gigaword. In *Proceedings of the Joint Workshop on Automatic Knowledge Base Construction and Web-scale Knowledge Extraction*, pages 95–100. Association for Computational Linguistics.
- Martha Palmer, Daniel Gildea, and Paul Kingsbury. 2005. The proposition bank: An annotated corpus of semantic roles. *Computational linguistics*, 31(1):71–106.
- Bo Pang and Lillian Lee. 2008. Opinion mining and sentiment analysis. *Foundations and trends in information retrieval*, 2(1-2):1–135.
- Marta Recasens, Cristian Danescu-Niculescu-Mizil, and Dan Jurafsky. 2013. Linguistic models for analyzing and detecting biased language. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 1650–1659.
- Evan Sandhaus. 2008. The new york times annotated corpus. *Linguistic Data Consortium, Philadelphia*, 6(12):e26752.
- Roger C Schank and Robert P Abelson. 1975. *Scripts, plans, and knowledge*. Yale University.
- Swapna Somasundaran and Janyce Wiebe. 2010. Recognizing stances in ideological on-line debates. In *Proceedings of the NAACL HLT 2010 Workshop on Computational Approaches to Analysis and Generation of Emotion in Text*, pages 116–124. Association for Computational Linguistics.
- Charles Sutton and Andrew McCallum. 2009. Piecewise training for structured prediction. *Machine Learning*, 77(2):165–194.
- Hiroya Takamura, Takashi Inui, and Manabu Okumura. 2005. Extracting semantic orientations of words using spin model. In *Proceedings of 43rd Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Matt Thomas, Bo Pang, and Lillian Lee. 2006. Get out the vote: Determining support or opposition from Congressional floor-debate transcripts. In *Proceedings of the 2006 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 327–335.
- Leonid Velikovich, Sasha Blair-Goldensohn, Kerry Hannan, and Ryan McDonald. 2010. The viability of web-derived polarity lexicons. In *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, HLT '10, pages 777–785.
- Janyce Wiebe and Lingjia Deng. 2014. An account of opinion implicatures. *CoRR*, abs/1404.6491.
- Janyce Wiebe, Theresa Wilson, and Claire Cardie. 2005. Annotating expressions of opinions and emotions in language. *Language resources and evaluation*, 39(2-3):165–210.
- Michael Wiegand and Josef Ruppenhofer. 2015. Opinion holder and target extraction based on the induction of verbal categories. *Proceedings of the 2015 Conference on Computational Natural Language Learning (CoNLL)*, page 215.
- Theresa Wilson, Janyce Wiebe, and Paul Hoffmann. 2005. Recognizing contextual polarity in phrase-level sentiment analysis. In *Proceedings of the conference on human language technology and empirical methods in natural language processing*, pages 347–354.
- Tae Yano, Philip Resnik, and Noah A. Smith. 2010. Shedding (a thousand points of) light on biased language. In *Proceedings of the NAACL HLT 2010 Workshop on Creating Speech and Language Data with Amazon's Mechanical Turk*, CSLDAMT '10, pages 152–158.