Discourse Level Opinion Interpretation *

Swapna Somasundaran
Dept. of Computer Science
University of Pittsburgh
Pittsburgh, PA 15260
swapna@cs.pitt.edu

Janyce Wiebe
Dept. of Computer Science
University of Pittsburgh
Pittsburgh, PA 15260
wiebe@cs.pitt.edu

Josef Ruppenhofer
Intelligent Systems Program
University of Pittsburgh
Pittsburgh, PA 15260
josefr@cs.pitt.edu

Abstract

This work proposes *opinion frames* as a representation of discourse-level associations which arise from related opinion topics. We illustrate how *opinion frames* help gather more information and also assist disambiguation. Finally we present the results of our experiments to detect these associations.

1 Introduction

Opinions have been investigated at the phrase, sentence, and document levels. However, little work has been carried out regarding interpreting opinions at the level of the discourse.

Consider the following excerpt from a dialog about designing a remote control for a television (the opinion *targets* – what the opinions are about — are shown in *italics*).

(1) D :: And I thought **not too** *edgy* and *like a box*, **more** kind of *hand-held* **not** as *computery*, yeah, **more** *organic shape* I think. Simple designs, like the last one we just saw, **not too** *many buttons*.

Speaker *D* expresses an opinion in favor of a design that is simple and organic in shape, and against an alternative design which is not. Several individual opinions are expressed in this passage. The first is a negative opinion about the design being too edgy and box-like, the next is a positive opinion toward a hand-held design, followed by a negative opinion toward a computery shape, and so on. While recognizing individual expressions of opinions and their properties is important, discourse interpretation is needed as well. It is by understanding the passage as a discourse that we see *edgy*, *like a box*, *computery*, and *many buttons* as descriptions of the type of design *D* does not prefer, and *hand-held*, *organic shape*, and *simple designs* as descriptions of the type he does. These descriptions are not in general synonyms/antonyms of one another; for example, there are hand-held "computery" devices and simple designs that are edgy. The unison/opposition among the descriptions is due to how they are used in the discourse.

This paper focuses on such relations between the targets of opinions in discourse. Specifically, in this work, we propose a scheme of *opinion frames*, which consist of two opinions that are related by virtue of having united or opposed targets. We argue that recognizing opinion frames will provide more opinion information for NLP applications than recognizing individual opinions alone. Further, if there is uncertainty about any one of the components, we believe opinion frames are an effective representation incorporating discourse information to make an overall coherent interpretation (Hobbs et al., 1993). Finally, we also report the first results of experiments in recognizing the presence of these opinion frames.

We introduce our data in Section 2, present *opinion frames* in Section 3 and illustrate their utility in Section 4. Our experiments are in Section 5, related work is discussed in Section 6, and conclusions are in Section 7.

2 Data

The data used in this work is the AMI meeting corpus (Carletta et al., 2005) which contains multi-modal recordings of group meetings. Each meeting has rich transcription and seg-

*This research was supported in part by the Department of Homeland Security under grant N000140710152.

Proceedings of the 22nd International Conference on Computational Linguistics (Coling 2008), pages 801–808
Manchester, August 2008

ment (turn/utterance) information for each speaker. Each utterance consists of one or more sentences. We also use some of the accompanying manual annotations (like adjacency pairs) as features in our machine learning experiments.

3 Opinion Frames

In this section, we lay out definitions relating to opinion frames, illustrate with examples how these are manifested in our data, and consider them in the context of discourse relations.

3.1 Definitions

The components of opinion frames are individual opinions and the relationships between their targets. Following (Wilson and Wiebe, 2005; Somasundaran et al., 2007), we address two types of opinions, *sentiment* and *arguing*.

Sentiment includes positive and negative evaluations, emotions, and judgments. **Arguing** includes arguing *for* or *against* something, and arguing *that* something should or should not be done.

Opinions have a **polarity** that can be *positive* or *negative*. [1] The **target** of an opinion is the entity or proposition that the opinion is about. We establish relations between targets, in the process relating their respective opinions. We address two types of relations, *same* and *alternative*.

The **same** relation holds between targets that refer to the same entity, property, or proposition. Observing the relations marked by annotators, we found that *same* covers not only identity, but also part-whole, synonymy, generalization, specialization, entity-attribute, instantiation, cause-effect, epithets and implicit background topic, i.e., relations that have been studied by many researchers in the context of anaphora and co-reference (e.g. (Clark, 1975; Vieira and Poesio, 2000; Mueller and Strube, 2001)). Actually, *same* relations holding between entities often involve co-reference (where co-reference is broadly conceived to include relations such as part-whole listed above). However, there are no morpho-syntactic constraints on what targets may be. Thus, *same* relations may also hold between adjective phrases, verb phrases, and clauses. An instance of this is Example 1, where the *same* target relation holds between the adjectives *edgy* and *computery*.

[1] Polarity can also be *neutral* or *both* (Wilson and Wiebe, 2005), but these values are not significant for our opinion frames.

SPSPsame, SNSNsame, APAPsame, ANANsame, SPAPsame, APSPsame, SNANsame, ANSNsame, SPSNalt, SNSPalt, APANalt, ANAPalt, SPANalt, SNAPalt, APSNalt, ANSPalt
SPSNsame, SNSPsame, APANsame, ANAPsame, SPANsame, APSNsame, SNAPsame, ANSPsame, SPSPalt, SNSNalt, APAPalt, ANANalt, SPAPalt, SNANalt, APSPalt, ANSNalt

Table 1: Opinion Frames

The **alternative** relation holds between targets that are related by virtue of being opposing (mutually exclusive) options in the context of the discourse. For example, in the domain of TV remote controls, the set of all shapes are alternatives to one another, since a remote control may have only one shape at a time. In such scenarios, a positive opinion regarding one choice may imply a negative opinion toward competing choices, and vice versa. Objects appear as alternatives via world and domain knowledge (for example, shapes of a remote); the context of the discourse (for example, Hillary Clinton and Barak Obama are alternatives in discussions of the primaries, but not in discussions of the general election); and the way the objects are juxtaposed while expressing opinions (for instance *hand-held* and *computery* in Example 1).

While *same* and *alternative* are not the only possible relations between targets, they commonly occur in task-oriented dialogs such as those in the data we use.

Now that we have all the ingredients, we can define opinion frames. An **opinion frame** is defined as a structure composed of two opinions and their respective targets connected via their target relations. With four opinion type/polarity pairs (*SN*,*SP*,*AN*,*AP*), for each of two opinion slots, and two possible target relations, we have 4 * 4 * 2 = 32 types of frame, listed in Table 1.

3.2 Examples

We will now illustrate how the frames are applied with the following meeting snippets from the AMI meeting corpus. In our examples, the lexical anchors revealing the opinion type (as the words are interpreted in context) are indicated in **bold face**. The text span capturing the target of the opinion (as interpreted in context) is indicated in *italics*. To make it easier to understand the opinion frames, we separately list each opinion, followed by the major relation between the targets and, in parentheses, the relevant subtype of the major relation.

In the passage below, the speaker *D* expresses

his preferences about the material for the TV remote.

(2) D:: ... this kind of rubbery material, *it's* a **bit more bouncy**, like you said they get chucked around a lot. A **bit more durable** and *that* [2] can also be **ergonomic** and *it* kind of feels **a bit different from all the other remote controls**.

Opinion Span - target Span	Type
O1 **bit more bouncy** - *it's* [t1]	SP
O2 **bit more durable** - ellipsis [t2]	SP
O3 **ergonomic** - *that* [t3]	SP
O4 **a bit different from all the other remote** - *it* [t4]	SP

Target - target	Rel
t1 - t2	same (ellipsis)
t3 - t4	same (identity)
t1 - t3	same (identity)

The speaker's positive sentiment regarding the rubbery material is apparent from the text spans **bit more bouncy** (**S**entiment **P**ositive or SP), **bit more durable** (SP), **ergonomic** (SP) and **a bit different from all the other remote controls** (SP). As shown, the targets of these opinions (*it's* [t1], *that* [t3], and *it* [t4]) are related by the *same* relation. The ellipsis occurs with **bit more durable**. Target [t2] represents the (implicit) target of that opinion, and [t2] has a *same* relation to [t1], the target of the **bit more bouncy** opinion. The opinion frames occurring throughout this passage are all *SPSPsame* denoting that both the opinion components are sentiments with positive polarity with a *same* relation between their targets. One frame occurs between O1 and O2, another between O3 and O4, and so on.

Example 2 illustrates relatively simple *same* relations between targets. Now let us consider the more involved passage below, in which a meeting participant analyzes two leading remotes on the market.

(3) D:: These are two **leading** *remote controls* at the moment. You know *they*'re **grey**, *this one*'s got **loads of buttons**, it's **hard to tell** from here what *they* actually do, and *they* **don't look very exciting at all**.

Opinion Span - target Span	Rel
O1 **leading** - *remote controls* [t1]	SP
O2 **grey** - *they* [t2]	SN
O3 **loads of buttons** - *this one* [t3]	SN
O4 **hard to tell** - *they* [t4]	SN
O5 **don't look very exciting at all** - *they* [t5]	SN

Target - target	Rel
t1 - t2	same (identity)
t2 - t3	same (t3 subset of t2)
t3 - t4	same (t4 partof t3)
t5 - t1	same (identity)

Target [t2] is the set of two leading remotes; [t3], which is in a *same* relation with [t2], is one of those remotes. Target [t4], which is also in a *same* relation with [t3], is a part of that remote, namely its buttons. Thus, opinion O3 is directly about one of the remotes, and indirectly about the set of both remotes. Similarly, O4 is directly about the buttons of one of the remotes, and indirectly about that remote itself. The assessments at different levels accrue toward the analysis of the main topic under consideration.

Moving on to *alternative* (*alt*) relations, consider the passage below, where the speaker is arguing for the curved shape.

(4) C:: . . . shapes **should be** *curved*, so round shapes. **Nothing** *square-like*.
⋮
C:: . . . So we **shouldn't have too** *square corners* and that kind of thing.
B:: Yeah okay. **Not *the old box look***.

Opinion Span - target Span	Rel
O1 **should be** - *curved* [t1]	AP
O2 **Nothing** - *square-like* [t2]	AN
O3 **shouldn't have** - *square corners* [t3]	AN
O4 **too** - *square corners* [t3]	SN
O5 **Not** - *the old box look* [t4]	AN
O6 **the old box look** - *the old box look* [t4]	SN

Target - target	Rel
t1 -t2	alternatives
t2 - t3	same (specification)
t3 - t4	same (epithet)

Opinion O1 argues for a curved shape, O2 argues against a square shape, and O3 argues against square corners. Note that *square corners* is also the target of a negative sentiment, O4, expressed here by **too**. Opinion O5 argues against the old box look. In addition, the wording **old box look** implies a negative sentiment – O6 (we list the target span as "old box look," which refers to the look of having square corners).

There is an *alt* relation between [t1] and [t2]. Thus, we have an opinion frame of type *APANalt* between O1 and O2. From this frame, we are able to understand that a positive opinion is expressed toward something and a negative opinion is expressed toward its alternative.

3.3 Link Transitivity

When individual targets are linked, they form a chain-like structure. Due to this, a connecting path may exist between targets that were not directly

[2] Note that the "that" refers to the property of being durable; however, as our annotation scheme is not hierarchical, we connect it to the entity the opinion is about – in this case the rubbery material.

linked by the human annotators. This path can be traversed to create links between new pairs of targets, which in turn results in new opinion frame relations.

Let us illustrate this idea with Example 4. The frames with direct relations are *O1O2 APANalt*. By following the *alt* link from [t1] to [t2] and the *same* link from [t2] to [t3], we have an *alt* link between [t1] and [t3], and the additional frames *O1O3 APANalt* and *O1O4 APSNalt*. Repeating this process would finally link speaker *C's* opinion O1 with *B's* opinion O6 via a *APSNalt* frame.

Simple recipes such as this can be used by applications such as QA to gather more information from the discourse.

3.4 Frame Types

In our corpus, we found that the 32 frames of Table 1 can be categorized into two functional types: *reinforcing frames* and *non-reinforcing frames*.

The set of frames that occur in scenarios where the speaker intends to fortify or *reinforce* his opinion/stance are called ***reinforcing frames***. These are the ones in the top row of the Table 1. Note that these frames cover all opinion types, polarities and target relations. It is the particular combination of these frame components that bring about the reinforcement of the opinion in the discourse.

On the other hand, the frames at the bottom row of the table are ***non-reinforcing***. In our corpus, these frames occur when a speaker is ambivalent or weighing pros and cons.

Example 2 is characterized by opinion frames in which the opinions reinforce one another – that is, individual positive sentiments (SP) occurring throughout the passage fortify the positive regard for the rubbery material via the *same* target relations and the resulting *SPSPsame* frames.

Interestingly, interplays among different opinion types may show the same type of reinforcement. For instance, Example 4 is characterized by mixtures of opinion types, polarities, and target relations. However, the opinions are still unified in the intention to argue for a particular type of shape.

3.5 Discourse Relations and Opinion Frames

Opinion-frame recognition and discourse interpretation go hand in hand; together, they provide richer overall interpretations. For example, consider the opinion frames and the Penn Discourse Treebank relations (Prasad et al., 2007) for Example 2. PDTB would see a *list* or *conjunction* relation between the clauses containing opinions **bit more durable** (O2) and **ergonomic** (O3), as well as between the clauses containing opinions **ergonomic** (O3) and **a bit different from all the other remote controls** (O4). All of our opinion frames for this passage are of type *SPSPsame*, a reinforcing frame type. This passage illustrates the case in which discourse relations nicely correspond to opinion frames. The opinion frames flesh out the discourse relations: we have lists specifically of positive sentiments toward related objects.

However, opinion-frame and discourse-relation schemes are not redundant. Consider the following three passages.

(e1) **Non-reinforcing opinion frame (SNSPsame); Contrast discourse relation**
D:: . . . I draw for you this *schema* that can be maybe **too technical** for you but is **very important** for me . . ..

(e2) **Reinforcing opinion frame (SNAPalt); Contrast discourse relation**
D:: **not too** *edgy* and like *a box*, **more** kind of *hand-held*

(e3) **Reinforcing opinion frame (SPSPsame); no discourse relation**
. . . they **want** something that's *easier to use* straight away, *more intuitive* perhaps.

In both e1 and e2, the discourse relation between the two opinions is *contrast* ("too technical" is contrasted with "very important", and "not too edgy and like a box" is contrasted with "more kind of hand-held"). However, the opinion frame in e1 is *SNSPsame*, which is a non-reinforcing frame, while the opinion frame in e2 is *SNAPalt*, which is a reinforcing frame. In e3, the opinion frame holds between targets within a subordinated clause (*easier to use* and *more intuitive* are two desired targets); most discourse theories don't predict any discourse relation in this situation.

Generally speaking, we find that there are not definitive mappings between opinion frames and the relations of popular discourse theories. For example, Hobbs' (Hobbs et al., 1993) *contrast* covers at least four of our frames (*SPSPalt*, *APAPalt*, *APANsame*, *SPSNsame*), while, for instance, our *SPSPsame* frame can map to both the *elaboration* and *explanation* relations.

4 Benefits of Discourse Opinion Frames

This section argues for two motivations for opinion frames: they may unearth additional information over and above the individual opinions stated in the text, and they may contribute toward arriving

	Positive	Negative
Counting only individual opinions		
Accepted Items	120	20
Rejected Items	9	12
individual + opinions via Reinforcing Opinion frames		
Accepted Items	252	63
Rejected Items	22	26

Table 2: Opinion Polarity Distribution for Accepted/Rejected Items

at a coherent interpretation (Hobbs et al., 1993) of the opinions in the discourse.

4.1 Gathering More Information

Frame relations provide a mechanism to relate opinions expressed in non-local contexts - the opinion may occur elsewhere in the discourse, but will become relevant to a given target due to a relation between its target and the given target. For instance, in Example 3, there is one direct evaluation of the leading remotes (O1) and two evaluations via identity (O2, O5). Following frames constructed via t2-t3 and t3-t4, we get two more opinions (O3 and O4) for the leading remotes.

Furthermore, opinions regarding something not lexically or even anaphorically related can become relevant, providing more opinion information. This is particularly interesting when *alt* relations are involved, as opinions towards one alternative imply opinions of opposite polarity toward the competing options. For instance in Example 4, if we consider only the explicitly stated opinions, there is only one (positive) opinion, O1, about the curved shape. However, the speaker expresses several other opinions which reinforce his positivity toward the curved shape. Thus, by using the frame information, it is possible to gather more opinions regarding curved shapes for TV remotes.

As a simple proof of concept, we counted the number of positive and negative opinions towards the items that were accepted or rejected in the meetings (information about accepted and rejected items is obtained from the manual abstractive summaries provided by the AMI corpus). Counts are obtained, over opinions manually annotated in the data, for two conditions: with and without frame information. The items in our meeting data are mainly options for the new TV remote, which include attributes and features like different shapes, materials, designs, and functionalities. We observed that for the accepted items, the number of positive opinions is higher and, for rejected items, the number of negative opinions is higher. The top section of Table 2 shows a contingency table of counts of positive/negative opinions for accepted/rejected items for 5 AMI meetings.

Then we counted the number of reinforcing opinions that were expressed regarding these items. This meant also counting additional opinions that were related via reinforcing frames. The bottom section of Table 2 shows the counts when the reinforcing frames are considered. Compared to the counts of only individual opinions, we see that the numbers in each cell have increased, while maintaining the same pattern of distribution.

Thus, in effect we have procured more instances of opinions for the items. We believe this added information would help applications like meeting summarizers and QA systems to make more informed decisions.

4.2 Interdependent Interpretation

We believe that our opinion frames, anaphoric relations and discourse relations can symbiotically help disambiguate each other in the discourse. In particular, suppose that some aspect of an individual opinion, such as polarity, is unclear. If the discourse suggests certain opinion frames, this may in turn resolve the underlying ambiguity.

Revisiting Example 2 from above, we see that out of context, the polarities of **bouncy** and **different from other remotes** are unclear (bounciness and being different may be negative attributes for another type of object). However, the polarities of two of the opinions are clear (**durable** and **ergonomic**). There is evidence in this passage of discourse continuity and *same* relations such as the pronouns, the lack of contrastive cue phrases, and so on. This evidence suggests that the speaker expresses similar opinions throughout the passage, making the opinion frame *SPSPsame* more likely throughout. Recognizing the frames would resolve the polarity ambiguities of **bouncy** and **different**.

In the following example (5), the positive sentiment (SP) towards the *this* and the positive arguing (AP) for the *it* are clear. These two individual opinions can be related by a *same*/*alt* target relation, be unrelated, or have some other relation not covered by our scheme (in which case we would not have a relation between them). There is evidence in the discourse that makes one interpretation more likely than others. The “so” indicates that the two clauses are highly likely to be related by a *cause* discourse

relation (PDTB). This information confirms a discourse continuity, as well as makes a reinforcing scenario likely, which makes the reinforcing frame *SPAPsame* highly probable. This increase in likelihood will in turn help a coreference system to increase its confidence that the "that" and the "it" co-refer.

(5) B :: ... and *this* will **definitely enhance our market sales**, **so we should** take *it* into consideration also.

Opinion Span - target Span	Rel
O1 **definitely enhance our market sales** - *this* [t1]	SP
O2 **so we should** - *it* [t2]	AP

Target - target	Rel
t1 -t2	same (identity)

5 Experiments

There has been much work on recognizing individual aspects of opinions like extracting individual opinions from phrases or sentences and recognizing opinion type and polarity. Accordingly, in our machine learning experiments we assume oracle opinion and polarity information. Our experiments thus focus on the new question: "Given two opinion sentences, determine if they participate in *any* frame relation." Here, an opinion sentence is a sentence containing one or more sentiment or arguing expression. In this work, we consider frame detection only between sentence pairs belonging to the same speaker.

5.1 Annotation of Gold Standard

Creating gold-standard opinion-frame data is accomplished by annotating frame components and then building the frames from those underlying annotations.

We began with annotations created by Somasundaran et al. (2007), namely four meetings of the AMI meeting corpus annotated for *sentiment* and *arguing* opinions (text anchor and type). Following that annotation scheme, we annotated an additional meeting. This gave us a corpus of 4436 sentences or 2942 segments (utterances). We added attributes to the existing opinion annotations, namely *polarity* and *target-id*. The *target-id* attribute links the opinion to its local target span. Relations between targets were then annotated. When a newly annotated target is similar (or opposed) to a set of targets already participating in *same* relations, then the *same* (or *alt*) link is made only to one of them - the one that seems most natural. This is often the one that is physically closest.

Content Word overlap between the sentence pair
Focus space overlap between the sentence pair
Anaphoric indicator in the second sentence
Time difference between the sentence pair
Number of intervening sentences
Existence of adjacency pair between the sentence pair
Bag of words for each sentence

Table 3: Features for Opinion Frame detection

Link transitivity is then used to connect targets that are not explicitly linked by the annotators.

All annotations were performed by two of the co-authors of this paper by consensus labeling. The details of our annotation scheme and inter-annotator agreement studies are presented in (Somasundaran et al., 2008).

Once the individual frame components are annotated, conceptually, a frame exists for a pair of opinions if their polarities are either *positive* or *negative* and their targets are in a *same* or *alt* relation. For our experiments, if a path exists between two targets, then their opinions are considered to be participating in an opinion-frame relation.

The experimental data consists of pairs of opinion sentences and the gold-standard information whether there exists a frame between them. We approximate continuous discourse by only pairing sentences that are not more than 10 sentences apart. We also filter out sentences that are less than two words in length in order to handle data skewness. This filters out very small sentences (e.g., "Cool.") which rarely participate in frames. The experiments were performed on a total of 2539 sentence pairs, of which 551 are positive instances.

5.2 Features

The factor that determines if two opinions are related is primarily the target relations between them. Instead of first finding the target span for each opinion sentence and then inferring if they should be related, we directly try to encode target relation information in our features. By this approach, even in the absence of explicit target-span information, we are able to determine if the opinion sentence pairs are related.

We explored a number of features to incorporate this. The set that give the best performance are listed in Table 3. The *content word overlap* feature captures the degree of topic overlap between the sentence pair, and looks for target relations via identity. The *focus space overlap* feature is motivated by our observation that partici-

	Acc.	Prec.	Recall	F-measure
False	78.3%	-	0%	-
Distribution	66%	21.7%	21.7%	21.4%
Random	50.0%	21.5%	49.4%	29.8 %
True	21.7%	21.6%	100%	35.5 %
System	67.6%	36.8%	64.9%	46%

Table 4: Automatic Detection of Opinion Frames

pants refer to an established discourse topic without explicitly referring to it. Thus, we construct a focus space for each sentence containing recently used NP chunks. The feature is the percent overlap between the focus spaces of the two opinion sentences. The *anaphoric indicator feature* checks for the presence of pronouns such as *it* and *that* in the second sentence to account for target relations via anaphora. The *time difference* between the sentences and the *number of intervening sentences* are useful features to capture the idea that topics shift with time. The *existence of an adjacency pair* [3] between the sentences can clue the system that the opinions in the sentences are related too. Finally, standard *bag of words* features are included for each sentence.

5.3 Results

We performed 5-fold cross validation experiments, using the standard SVMperf package (Joachims, 2005), an implementation of SVMs designed for optimizing multivariate performance measures. We found that, on our skewed data, optimizing on F-measure obtains the best results.

Our system is compared to four baselines in Table 4. The majority class baseline which always guesses false (*False*) has good accuracy but zero recall. The baseline that always guesses true (*True*) has 100% recall and the best f-measure among the baselines, but poor accuracy. We also constructed a baseline that guesses true/false over the test set based on the distribution in the training data (*Distribution*). This baseline is smarter than the other baselines, as it does not indiscriminately guess any one of the class. The last baseline *Random* guesses true 50% of the time.

The bottom row of Table 4 shows the performance of our system (*System*). The skewness of the data affects the baselines as well as our system. Our system beats the best baseline f-measure by over 10 percentage points, and the best baseline precision by 14 percentage points. Comparing it to the baseline which has comparable accuracy, namely *Distribution*, we see that our system improves in f-measure by 24 percentage points.

Our results are encouraging - even using simple features to capture target relations achieves considerable improvement over the baselines. However, there is much room for improvement. Using more detailed target and discourse information promises to further improve system performance. These are avenues for future work.

6 Related work

Evidence from the surrounding context has been used previously to determine if the current sentence should be subjective/objective (Riloff et al., 2003; Pang and Lee, 2004) and adjacency pair information has been used to predict congressional votes (Thomas et al., 2006). However, these methods do not explicitly model the relations between opinions. An application of the idea of alternative targets can be seen in Kim and Hovy's (2007) work on election prediction. They assume that if a speaker expresses support for one party, all mentions of the competing parties have negative polarity, thus creating automatically labeled training data.

In the field of product review mining, sentiments and features (*aspects*) have been mined (Popescu and Etzioni, 2005), where the aspects correspond to our definition of targets. However, the aspects themselves are not related to each other in any fashion.

Polanyi and Zaenen (2006), in their discussion on contextual valence shifters, have also observed the phenomena described in this work - namely that a central topic may be divided into subtopics in order to perform evaluations, and that discourse structure can influence the overall interpretation of valence. Snyder and Barzilay (2007) combine an agreement model based on contrastive RST relations with a local aspect model to make a more informed overall decision for sentiment classification. In our scheme, their aspects would be related as *same* and their high contrast relations would correspond to the non-reinforcing frames *SPSNsame, SNSPsame*. Additionally, our frame relations would link the sentiments across non-adjacent clauses, and make connections via *alt* target relations.

With regard to meetings, the most closely related work includes the dialog-related annotation

[3] Adjacency Pairs are manual dialog annotations available in the AMI corpus.

schemes for various available corpora of conversation (e.g., Carletta et al. (2005) for AMI). As shown by Somasundaran et al. (2007), dialog structure information and opinions are in fact complementary. We believe that, like the discourse relations, the dialog information will additionally help in arriving at an overall coherent interpretation.

7 Conclusions

In this paper, we described the idea of *opinion frames* as a representation capturing discourse level relations that arise from related opinion targets and which are common in task-oriented dialogs. We introduced the *alternative* relations that hold between targets by virtue of being opposing in the discourse context. We discussed how our opinion-frame scheme and discourse relations go hand in hand to provide a richer overall interpretation. We also illustrated that such discourse level opinion associations have useful benefits, namely they help gather more opinion information and help interdependent interpretation. Finally, we showed via our machine learning experiments that the presence of opinion frames can be automatically detected.

References

Carletta, J., S. Ashby, S. Bourban, M. Flynn, M. Guillemot, T. Hain, J. Kadlec, V. Karaiskos, W. Kraaij, M. Kronenthal, G. Lathoud, M. Lincoln, A. Lisowska, I. McCowan, W. Post, D. Reidsma, and P. Wellner. 2005. The AMI Meetings Corpus. In *Proceedings of Measuring Behavior Symposium on "Annotating and measuring Meeting Behavior"*.

Clark, H. H. 1975. Bridging. *Theoretical issues in natural language processing . New York: ACM*.

Hobbs, J., M. Stickel, D. Appelt, and P. Martin. 1993. Interpretation as abduction. *AI*, 63.

Joachims, T. 2005. A support vector method for multivariate performance measures. In *ICML 2005*.

Kim, Soo-Min and Eduard Hovy. 2007. Crystal: Analyzing predictive opinions on the web. In *EMNLP-CoNLL 2007*.

Mueller, C. and M. Strube. 2001. Annotating anaphoric and bridging relations with mmax. In *2nd SIGdial Workshop on Discourse and Dialogue*.

Pang, B. and L. Lee. 2004. A sentimental education: Sentiment analysis using subjectivity summarization based on minimum cuts. In *ACl 2004*.

Polanyi, L. and A. Zaenen, 2006. *Contextual Valence Shifters*, chapter 1. Computing Attitude and Affect in Text: Theory and Applications. Springer.

Popescu, A.-M. and O. Etzioni. 2005. Extracting product features and opinions from reviews. In *HLT-EMNLP 2005*.

Prasad, R., E. Miltsakaki, N. Dinesh, A. Lee, A. Joshi, L. Robaldo, and B. Webber, 2007. *PDTB 2.0 Annotation Manual*.

Riloff, E., J. Wiebe, and T. Wilson. 2003. Learning subjective nouns using extraction pattern bootstrapping. In *CoNLL 2003*.

Snyder, B. and R. Barzilay. 2007. Multiple aspect ranking using the good grief algorithm. In *HLT 2007: NAACL*.

Somasundaran, S., J. Ruppenhofer, and J. Wiebe. 2007. Detecting arguing and sentiment in meetings. In *SIGdial Workshop on Discourse and Dialogue 2007*.

Somasundaran, S, J Ruppenhofer, and J Wiebe. 2008. Discourse level opinion relations: An annotation study. In *SIGdial Workshop on Discourse and Dialogue*. ACL.

Thomas, M., B. Pang, and L. Lee. 2006. Get out the vote: Determining support or opposition from congressional floor-debate transcripts. In *EMNLP 2006*.

Vieira, R. and M. Poesio. 2000. An empirically based system for processing definite descriptions. *Comput. Linguist.*, 26(4).

Wilson, T. and J. Wiebe. 2005. Annotating attributions and private states. In *Proceedings of ACL Workshop on Frontiers in Corpus Annotation II: Pie in the Sky*.