NAACL-HLT 2010 Workshop on Computational Approaches to Analysis and Generation of Emotion in Text

Toward Plot Units: Automatic Affect State Analysis

Amit Goyal and **Ellen Riloff** and **Hal Daume III** and **Nathan Gilbert**
School of Computing
University of Utah
Salt Lake City, UT 84112
{amitg,riloff,hal,ngilbert}@cs.utah.edu

Abstract

We present a system called AESOP that automatically produces affect states associated with characters in a story. This research represents a first step toward the automatic generation of plot unit structures from text. AESOP incorporates several existing sentiment analysis tools and lexicons to evaluate the effectiveness of current sentiment technology on this task. AESOP also includes two novel components: a method for acquiring *patient polarity verbs*, which impart negative affect on their patients, and *affect projection rules* to propagate affect tags from surrounding words onto the characters in the story. We evaluate AESOP on a small collection of fables.

1 Introduction

In the 1980s, plot units (Lehnert, 1981) were proposed as a knowledge structure for representing narrative stories and generating summaries. Plot units are fundamentally different from the story representations that preceded them because they focus on the emotional states and tensions between characters as the driving force behind interesting plots and cohesive stories. Plot units were used in narrative summarization studies, both in computer science and psychology (Lehnert et al., 1981), but the computational models of plot units relied on tremendous amounts of manual knowledge engineering.

Given the recent swell of activity in automated methods for sentiment analysis, we embarked on a project to see whether current techniques could automatically detect the affect states needed for plot unit analysis. Plot units are complex structures that include affect states, causal links, and cross-character links, and generating complete plot unit structures is beyond the scope of this work. As an initial step toward the long-term goal of automatically generating plot units, we began by creating a system to automatically identify the affect states associated with characters. An *affect state* represents the emotional state of a character, based on their perspective of events in the story. Plots units include three types of affect states: positive (+) states, negative (-) states, and mental (M) states that have neutral emotion (these are often associated with plans and goals).

Our system, called AESOP, pulls together a variety of existing technologies in sentiment analysis to automatically identify words and phrases that have positive/negative polarity or that correspond to speech acts (for mental states). However, we needed to develop a method to automatically map these affect tags onto characters in the story.[1] To address this issue, we created *affect projection rules* that propagate affect tags from words and phrases to characters in the story via syntactic relations.

During the course of our research, we came to appreciate that affect states, of the type required for plot units, can represent much more than just direct expressions of emotion. A common phenomena are affect states that result from a character being acted upon in a positive or negative way. For example, *"the cat ate the mouse"* produces a positive affect state for the cat and a negative affect

[1]This is somewhat analogous to, but not exactly the same as, associating opinion words with their targets or topics (Kim and Hovy, 2006; Stoyanov and Cardie, 2008).

The Father and His Sons

(s1) A father had a family of sons who were perpetually quarreling among themselves. (s2) When he failed to heal their disputes by his exhortations, he determined to give them a practical illustration of the evils of disunion; and for this purpose he one day told them to bring him a bundle of sticks. (s3) When they had done so, he placed the faggot into the hands of each of them in succession, and ordered them to break it in pieces. (s4) They tried with all their strength, and were not able to do it. (s5) He next opened the faggot, took the sticks separately, one by one, and again put them into his sons' hands, upon which they broke them easily. (s6) He then addressed them in these words: "My sons, if you are of one mind, and unite to assist each other, you will be as this faggot, uninjured by all the attempts of your enemies; but if you are divided among yourselves, you will be broken as easily as these sticks."

(a) "Father and Sons" Fable

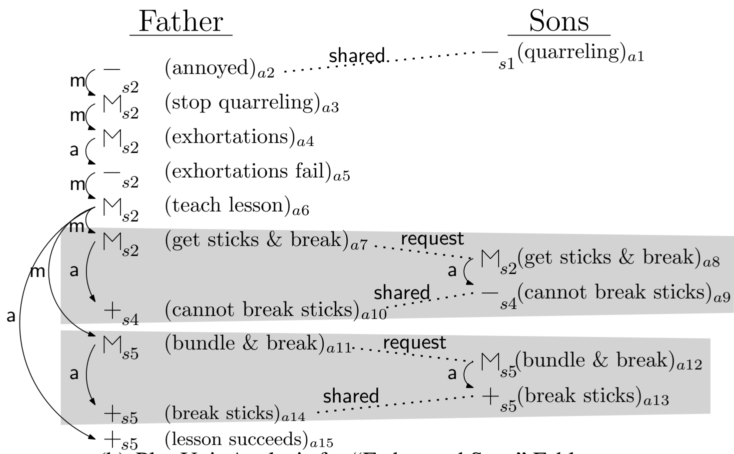

(b) Plot Unit Analysis for "Father and Sons" Fable

state for the mouse because obtaining food is good but being eaten is bad. This type of world knowledge is difficult to obtain, yet essential for plot unit analysis. In AESOP, we use corpus statistics to automatically learn a set of *negative patient polarity verbs* which impart a negative polarity on their patient (e.g., *eaten, killed, injured, fired*). To acquire these verbs, we queried a large corpus with patterns to identify verbs that frequently occur with agents who stereotypically have evil intent.

We evaulate our complete system on a set of AESOP's fables. In this paper, we also explain and categorize different types of situations that can produce affect states, several of which cannot be automatically recognized by existing sentiment analysis technology. We hope that one contribution of our work will be to create a better awareness of, and appreciation for, the different types of language understanding mechanisms that will ultimately be necessary for comprehensive affect state analysis.

2 Overview of Plot Units

Narratives can often be understood in terms of the emotional reactions and affect states of the characters therein. The plot unit formalism (Lehnert, 1981) provides a representational mechanism for affect states and the relationships between them. Plot unit structures can be used for tasks such as narrative summarization and question answering.

Plot unit structures consist of *affect states* for each character in a narrative, and links explaining the relationships between these affect states. The affect states themselves each have a type: (+) for positive states, (-) for negative states, and (M) for mental states (with neutral affect). Although affect states are *not* events per se, events often trigger affect states. If an event affects multiple characters, it can trigger multiple affect states, one for each character.

Affect states are further connected by causal links, which explain how the narrative hangs together. These include motivations (m), actualizations (a), terminations (t) and equivalences (e). Causal links exist between affect states for the same character. Cross-character links explain how single events affect two characters. For instance, if one character *requests* something of the other, this is an M-to-M link, since it spans a shared mental affect for both characters. Other speech acts can be represented as M to + (promise) or M to - (threat).

To get a better feeling of the plot unit representation, a short fable, "The Father and His Sons," is shown in Figure 1(a) and our annotation of its plot unit structure is shown in Figure 1(b). In this fable, there are two characters (the "Father" and the "Sons") who go through a series of affect states, depicted chronologically in the two columns.

In this example, the first affect state is a negative state for the sons, who are quarreling ($a1$). This state is *shared* by the father (via a cross-character link) who has a negative annoyance state ($a2$). The father then decides that he wants to stop the sons from quarreling, which is a mental event ($a3$). The causal link from $a2$ to $a3$ with an m label indicates a "motivation." His first attempt is by exhortations ($a4$).

This produces an M ($a3$) linked to an M ($a4$) with a m (motivation) link, which represents subgoaling. The father's overall goal is to stop the quarreling ($a3$) and in order to do so, he creates a subgoal of exhorting the sons to stop ($a4$). The exhortations fail, which produces a negative state ($a5$) for the father. The a causal link indicates an "actualization", representing the failure of the plan ($a4$).

The failure of the father's exhortations leads to a new subgoal: to teach the sons a lesson ($a6$). The m link from $a5$ to $a6$ is an example of "enablement." At a high level, this subgoal has two parts, indicated by the two gray regions ($a7 - a10$ and $a11 - a14$). The first gray region begins with a cross-character link (M to M), which indicates a request (in this case, to break a bundle of sticks). The sons fail at this, which upsets them ($a9$) but pleases the father ($a10$). The second gray region depicts the second part of the father's subgoal; he makes a second request ($a11$ to $a12$) to separate the bundle and break the sticks, which the sons successfully do, making them happy ($a13$) and the father happy ($a14$). This latter structure (the second gray region) is an HONORED REQUEST plot unit. At the end, the father's plan succeeds ($a15$) which is an actualization (a link) of his goal to teach the sons a lesson ($a6$).

In this example, as well as the others that we annotated in our gold standard, (see Section 5.1), we annotated *conservatively.* In particular, in reading the story, we may *assume* that the father's original plan of stopping the son's quarrelling also succeeded. However, this is not mentioned in the story and therefore we chose not to represent it. It is also important to note that plot unit representations can have t (termination) and e (equivalence) links that point backwards in time, but they do not occur in the Father and Sons fable.

3 Where Do Affect States Come From?

We began this research with the hope that recent research in sentiment analysis would supply us with effective tools to recognize affect states. However, we soon realized that affect states, as required for plot unit analysis, go well beyond the notions of positive/negative polarity and private states that have been studied in recent sentiment analysis work. In this section, we explain the wide variety of situations that can produce an affect state, based on our observations in working with fables. Most likely, an even wider variety of situations could produce affect states in other text genres.

3.1 Direct Expressions of Emotion

Plot units can include affect states that correspond to explicit expressions of positive/negative emotional states, as has been studied in the realm of sentiment analysis. For example, *"Max was disappointed"* produces a negative affect state for Max, and *"Max was pleased"* produces a positive affect state for Max. However, the affect must relate to an event that occurs in the story's plot. For example, a hypothetical expression of emotion would not yield an affect state (e.g., *"if the rain stops, she will be pleased"*).

3.2 Situational Affect States

Positive and negative affect states also frequently represent good and bad situational states that characters find themselves in. These states do not represent emotion, but indicate whether a situation is good or bad for a character based on world knowledge. For example, *"Wolf, who had a bone stuck in his throat, ..."* produces a negative affect state for the wolf. Similarly, *"The Old Woman recovered her sight..."* produces a positive affect state. Sentiment analysis is not sufficient to generate these affect states. Sometimes, however, a direct expression of emotion will also be present (e.g., *"Wolf was unhappy because he had a bone stuck..."*), providing redundancy and multiple opportunities to recognize the correct affect state for a character.

Situational affect states are common and often motivate plans and goals that are central to the plot.

3.3 Plans and Goals

Plans and goals are another common reason for affect states. The existence of a plan or goal is usually represented as a mental state (M). Plans and goals can be difficult to detect automatically. A story may reveal that a character has a plan or goal in a variety of ways, such as:

Direct expressions of plans/goals: a plan or goal may be explicitly stated (e.g., *"the lion wanted to find food"*). In this case, a mental state (M) should be generated.

Speech acts: a plan or goal may be revealed through a speech act between characters. For example, *"the wolf asked an eagle to extract the bone"* is a directive speech act that indicates the wolf's plan to resolve its negative state (having a bone stuck). This example illustrates how a negative state (bone stuck) can motivate a mental state (plan). When a speech act involves multiple characters, it produces multiple mental states. For example, a mental state should also be produced for the eagle, because it now has a plan to help the wolf (by virtue of being asked).

Inferred plans/goals: plans and goals sometimes must be inferred from actions. For example, *"the lion hunted deer"* reveals the lion's plan to obtain food. Similarly, *the serpent spat poison into the man's water"* implies that the serpent had a plan to kill the man.

Plans and goals also produce positive/negative affect states when they succeed/fail. For example, if the eagle successfully extracts the bone from the wolf's throat, then both the wolf and the eagle will have positive affect states, because both were successful in their respective goals. A directive speech act between two characters coupled with positive affect states for both characters is a common plot unit structure called an HONORED REQUEST, depicted by the second gray block shown in Fig.1(b).

The affect state for a character is always with respect to *its* view of the situation. For example, consider: *"The owl besought a grasshopper to stop chirping. The grasshopper refused to desist, and chirped louder and louder."* Both the owl and the grasshopper have M affect states representing the request from the owl to the grasshopper (i.e., the owl's plan to stop the chirping is to ask the grasshopper to knock it off). The grasshopper refuses the request, so a negative affect state is produced for the owl, indicating that its plan failed. However, a positive affect state is produced for the grasshopper, because its goal was to continue chirping which was accomplished by refusing the request. This scenario is also a common plot unit structure called a DENIED REQUEST.

3.4 Patient Role Affect States

Many affect states come directly from events. In particular, when a character is acted upon (the *theme* or *patient* of an event), a positive or negative affect state often results for the character. These affect states reflect world knowledge about what situations are good and bad. For example:

Negative patient roles: *killed X*, *ate X*, *chased X*, *captured X*, *fired X*, *tortured X*

Positive patient roles: *rescued X*, *fed X*, *adopted X*, *housed X*, *protected X*, *rewarded X*

For example, *"a man captured a bear"* indicates a negative state for the bear. Overall, this sentence would generate a SUCCESS plot unit consisting of an M state and a + state for the man (with an actualization a causal link between them representing the plan's success) and a - state for the bear (as a cross-character link indicating that what was good for the man was bad for the bear). A tremendous amount of world knowledge is needed to generate these states from such a seemingly simple sentence. Similarly, if a character is rescued, fed, or adopted, then a + affect state should be produced for the character based on knowledge that these events are desirable. We are not aware of existing resources that can automatically identify affect polarity with respect to event roles. In Section 4.1.2, we explain how we automatically acquire *Patient Polarity Verbs* from a corpus to identify some of these affect states.

4 AESOP: Automatic Affect State Analysis

We created a system, called AESOP, to try to automatically identify the types of affect states that are required for plot unit analysis. AESOP incorporates existing resources for sentiment analysis and speech act recognition, and includes two novel components: *patient polarity verbs*, which we automatically generate using corpus statistics, and *affect projection rules*, which automatically project and infer affect labels via syntactic relations.

AESOP produces affect states in a 3-step process. First, AESOP labels individual words and phrases with an M, +, or - affect tag. Second, it identifies all references to the two main characters of the story. Third, AESOP applies affect projection rules to propagate affect states onto the characters, and in some cases, to infer new affect states.

4.1 Step 1: Assigning Affect Tags to Words

4.1.1 Sentiment Analysis Resources

AESOP incorporates several existing sentiment analysis resources to recognize affect states associated with emotions and speech acts.

• OpinionFinder[2] (Wilson et al., 2005) (Version 1.4) is used to identify all three types of states. We use the +/- labels assigned by its contextual polarity classifier (Wilson, 2005) to create +/- affect tags. The MPQASD tags produced by its Direct Subjective and Speech Event Identifier (Choi et al., 2006) are used as M affect tags.

• Subjectivity Lexicon[3] (Wilson, 2005): The positive/negative words in this list are assigned +/- affect tags, when they occur with the designated part-of-speech (POS).

• Semantic Orientation Lexicon[4] (Takamura et al., 2005): The positive/negative words in this list are assigned +/- affect tags, when they occur with the designated part-of-speech.

• A list of 228 speech act verbs compiled from (Wierzbicka, 1987)[5], which are used for M states.

4.1.2 Patient Polarity Verbs

As we discussed in Section 3.4, existing resources are not sufficient to identify affect states that arise from a character being acted upon. Sentiment lexicons, for example, assign polarity to verbs irrespective of their agents or patients. To fill this gap, we tried to automatically acquire verbs that have a strong patient polarity (i.e., the patient will be in a good or bad state by virtue of being acted upon).

We used corpus statistics to identify verbs that frequently occur with agents who typically have evil (negative) or charitable (positive) intent. First, we identified 40 words that are stereotypically evil agents, such as *monster*, *villain*, *terrorist*, and *murderer*, and 40 words that are stereotypically charitable agents, such as *hero*, *angel*, *benefactor*, and *rescuer*. Next, we searched the google Web $1T$ 5-gram corpus[6] using patterns designed to identify verbs that co-occur with these words as agents. For each agent term, we applied the pattern "*ed by [a,an,the] AGENT" and extracted the list of matching verbs.[7]

Next, we rank the extracted verbs by computing the ratio between the frequency of the verb with a negative agent versus a positive agent. If this ratio is > 1, then we save the verb as a *negative patient polarity verb* (i.e., it imparts negative polarity to its patient). This process produced 408 negative patient polarity verbs, most of which seemed clearly negative for the patient. Table 1 shows the top 20 extracted verbs. We also tried to identify positive patient polarity verbs using a positive-to-negative ratio, but the extracted verbs were often neutral for the patient, so we did not use them.

scammed	damaged	disrupted	ripped
raided	corrupted	hindered	crippled
slammed	chased	undermined	possesed
dogged	tainted	grounded	levied
patched	victimized	posessed	bothered

Table 1: Top 20 negative patient polarity verbs

4.2 Step 2: Identifying the Characters

The problem of coreference resolution in fables is somewhat different than for other genres, primarily because characters are often animals (e.g., *"he"*=*"owl"*). So we hand-crafted a simple rule-based coreference system. For the sake of this task, we made two assumptions: (1) There are only two characters per fable, and (2) Both characters are mentioned in the fable's title.

We then apply heuristics to determine number and gender for the characters based on word lists, WordNet (Miller, 1990) and POS tags. If no determination of a character's gender or number can be made from these resources, a process of elimination is employed. Given the two character assumption, if one character is known to be male, but there are female pronouns in the fable, then the other character is assumed to be female. The same is done for number agreement. Finally, if there is only one character between a pronoun and the beginning of a document,

[2]http://www.cs.pitt.edu/mpqa/opinionfinderrelease/

[3]http://www.cs.pitt.edu/mpqa/lexiconrelease/collectinfo1.html

[4]http://www.lr.pi.titech.ac.jp/∼takamura/pndic_en.html

[5]http://openlibrary.org/b/OL2413134M/English_speech_act_verbs

[6]http://www.ldc.upenn.edu/Catalog/CatalogEntry.jsp?catalogId=LDC2006T13

[7]The corpus is not POS tagged so there is no guarantee these will be verbs, but they usually are in this construction.

the pronoun is assumed to corefer with that character. The character then assumes the gender and number of that pronoun. Lastly, WordNet is used to obtain a small set of non-pronominal, non-string-match resolutions by exploiting hypernym relations, for instance, linking *Peasant* with *the man*.

4.3 Step 3: Affect Projection

Our goal is to produce affect states for each character in the story. Therefore every affect tag needs to be attributed to a character, or discarded. Since plots typically revolve around actions, we used the verbs as the basis for projecting affect tags onto the characters. In some cases, we also spawn new affect tags associated with mental states to indicate that an action is likely the manifestation of a plan.

We developed 6 types of *affect projection rules* that orchestrate how affect tags are assigned to the characters based on verb argument structure. We use the Sundance shallow parsing toolkit (Riloff and Phillips, 2004) to generate a syntactic analysis of each sentence, including syntactic chunking, clause segmentation, and active/passive voice recognition. We normalize the verb phrases (VPs) with respect to voice (i.e., we transform the passive voice constructions into an active voice equivalent) to simplify our rules. We then make the assumption that the Subject of the VP is its AGENT and the Direct Object of the VP is its PATIENT.[8] The affect projection rules only project affect states onto AGENTS and PATIENTS that correspond to a character in the story. The five types of rules are described below.

1. AGENT **VP** : This case applies when the VP has no PATIENT, or a PATIENT that is not a character in the story, or the PATIENT corefers with the AGENT. All affect tags associated with the VP are projected onto the AGENT. For example, *"Mary **laughed** (+)"* projects a positive affect state onto Mary.

2. **VP** PATIENT[9]: All affect tags associated with the VP are projected onto the PATIENT, unless both M and +/- tags exist, in which case only the +/- tags are projected. For example, *"**loved** (+) the cat"*, projects a positive affect state onto the cat.

[8] We are not actually doing thematic role recognition, so this will not always be correct, but it is a reasonable approximation.

[9] Agent is missing or not a character.

3. AGENT **VP** PATIENT: This case applies when the AGENT and PATIENT refer to different characters. All affect tags associated with the VP are projected onto the PATIENT, unless both M and +/- tags exist, in which case only the +/- tags are projected (as in Rule #2). If the VP has an M tag, then we also project an M tag onto the AGENT (representing a shared, cross-character mental state). If the VP has a +/- tag, then we project a + tag onto the agent (as an inference that the AGENT accomplished some action).

4. AGENT **VERB1** to **VERB2** PATIENT. We divide this into two cases: (a) If the agent and patient refer to the same character, then Rule #1 is applied (e.g., *"Bo decided to teach himself..."*). (b) If the agent and patient are different, we apply Rule #1 to **VERB1** to agent and Rule #2 to **VERB2**. If no affect tags are assigned to either verb, then we create an M affect state for the agent (assuming that the VP represents some sort of plan).

5. If a noun phrase refers to a character and includes a modifying adjective with an affect tag, then the affect is mapped onto the character. For example, *"the **happy** (+) fox"*.

Finally, if an adverb or adjectival phrase (e.g., predicate adjective) has an affect tag, then that affect tag is mapped onto the preceding VP and the projection rules above are applied. For all of the rules, if a clause contains a negation word, then we flip the polarity of all words in that clause. Our negation list contains: *no, not, never, fail, failed, fails, don't, and didn't.*

5 Evaluation

5.1 Data Set

Plot unit analysis of ordinary text is enormously complex – even the idea of *manually* creating gold standard annotations seemed like a monumental task. So we began our exploration with simpler and more constrained texts that seemed particularly appropriate for plot unit analysis: fables. Fables have two desirable attributes: (1) they have a small cast of characters, and (2) they typically revolve around a moral, which is exemplified by a short and concise plot. Even so, fables are challenging for NLP due to anthropomorphic characters, flowery language, and sometimes archaic vocabulary.

State	**M** (66)			**+** (52)			**-** (39)			**All** (157)		
System	*R*	*P*	*F*	*R*	*P*	*F*	*R*	*P*	*F*	*R*	*P*	*F*
Bsent baseline	.65	.10	.17	.52	.08	.14	.74	.06	.11	.63	.08	.14
Bclause baseline	.48	.28	.35	.44	.22	.29	.69	.17	.27	.52	.22	.31
All 4 resources (w/proj. rules)	.48	.43	.45	.23	.39	.29	.23	.41	.29	.34	.41	.37
OpinionFinder	.36	.42	.39	.00	.00	.00	.00	.00	.00	.15	.35	.21
Subjectivity Lexicon	.45	.43	.44	.23	.35	.28	.21	.44	.28	.32	.41	.36
Semantic Dictionary	.42	.45	.43	.00	.00	.00	.00	.00	.00	.18	.45	.26
Semantic Orientation Lexicon	.41	.43	.42	.17	.53	.26	.08	.43	.13	.25	.45	.32
PPV Lexicon	.41	.42	.41	.02	.17	.04	.21	.73	.33	.23	.44	.30
AESOP (All 4 + PPV)	.48	.40	.44	.25	.36	.30	.33	.46	.38	.37	.40	.38

Table 2: Evaluation results for 2 baselines, 4 sentiment analysis resources with projection rules, and our PPV lexicon with projection rules. (The # in parentheses is the number of occurrences of that state in the gold standard).

We collected 34 fables from an Aesop's Fables web site[10], choosing fables that have a true plot (some only contain quotes) and exactly two characters. We divided them into a development set of 11 stories, a tuning set of 8 stories, and a test set of 15 stories. The Father and Sons story from Figure 1(a) is an example from our set.

Creating a gold standard was itself a substantial undertaking. Plot units are complex structures, and training non-experts to produce them did not seem feasible in the short term. So three of the authors discussed and iteratively refined manual annotations for the development and tuning set stories until we became comfortable that we had a common understanding for the annotation task. Then to create our gold standard test set, two authors independently created annotations for the test set, and a third author adjudicated the differences. The gold standard contains complete plot unit annotations, including affect states, causal links, and cross-character links. For the experiments in this paper, however, only the affect state annotations were used.

5.2 Baselines

We created two baselines to measure what would happen if we use all 4 sentiment analysis resources *without* any projection rules. The first one (Bsent) operates at the sentence level. It naively projects every affect tag that occurs in a sentence onto every character in the same sentence. The second baseline (Bclause) operates identically, but at the clause level.

[10]http://www.pacificnet.net/~johnr/aesop/

5.3 Evaluation

As our evaluation metrics we used recall (R), precision (P), and F-measure (F). We evaluate each system on individual affect states (+, - and M) as well as across all affect states. The evaluation is done at the sentence level. Meaning, if a system produces the same affect state as present in the gold standard for a sentence, we count it as a correct affect state. Our main evaluation also requires each affect state to be associated with the correct character.

Table 2 shows the coverage of our two baseline systems as well as the four Sentiment Analysis Resources used with our projection rules. We can make several observations:

- As expected, the baselines achieve relatively high recall, but low precision.
- Each of the sentiment analysis resources alone is useful, and using them with the projection rules leads to improved performance over the baselines (10 points in F score for M and 6 points overall). This shows that the projection rules are helpful in identifying the characters associated with each affect state.
- The PPV Lexicon, alone, is quite good at capturing negative affect states. Together with the projection rules, this leads to good performance on identifying mental states as well.

To better assess our projection rules, we evaluated the systems both with respect to characters and without respect to characters. In this evaluation, system-produced states are correct even if they are assigned to the wrong character. Table 3 reveals several results: (1) For the baseline: there is a large drop when

State	**M** (66)			**+** (52)			**-** (39)			**All** (157)		
System	*R*	*P*	*F*	*R*	*P*	*F*	*R*	*P*	*F*	*R*	*P*	*F*
Bclause w/o char	.65	.37	.47	.50	.25	.33	.77	.19	.30	.63	.26	.37
AESOP w/o char	.55	.44	.49	.33	.47	.39	.36	.50	.42	.43	.46	.44
Bclause w/ char	.48	.28	.35	.44	.22	.29	.69	.17	.27	.52	.22	.31
AESOP w/ char	.48	.40	.44	.25	.36	.30	.33	.46	.38	.37	.40	.38

Table 3: Evaluating affect states with and without respect to character.

State	**M** (66)			**+** (52)			**-** (39)			**All** (157)		
System	*R*	*P*	*F*	*R*	*P*	*F*	*R*	*P*	*F*	*R*	*P*	*F*
Bclause PCoref	.48	.28	.35	.44	.22	.29	.69	.17	.27	.52	.22	.31
AESOP PCoref	.48	.40	.44	.25	.36	.30	.33	.46	.38	.37	.40	.38
Bclause ACoref	.42	.45	.43	.25	.34	.29	.54	.24	.33	.39	.33	.36
AESOP ACoref	.41	.54	.47	.12	.40	.18	.26	.45	.33	.27	.49	.35

Table 4: Final results of Bclause and AESOP systems with perfect and automated coreference

evaluated with respect to the correct character. (2) For AESOP: there is a smaller drop in both precision and recall for M and -, suggesting that our projection rules are doing well for these affect states. (3) For AESOP: there is a large drop in both precision and recall for +, suggesting that there is room for improvement of our projection rules for positive affect.

Finally, we wish to understand the role that coreference plays. Table 4 summarizes the results with perfect coreference and with automated coreference. AESOP is better than both baselines when we use perfect coreference (PCoref), which indicates that the affect projection rules are useful. However, when we use automated coreference (ACoref), recall goes down and precision goes up. Recall goes down because our automated coreference system is precision oriented: it only says "coreferent" if it is sure.

The increase in precision when moving to automated coreference is bizarre. We suspect it is primarily due to the handling of quotations. Our perfect coreference system resolves first and second person pronouns in quotations, but the automated system does not. Thus, with automated coreference, we almost never produce affect states from quotations. This is a double-edged sword: sometimes quotes contain important affect states, sometimes they do not. For example, from the Father and Sons fable, "if you are **divided** among yourselves, you will be **broken** as easily as these sticks." Automated coreference does not produce any character resolutions and therefore AESOP produces no affect states. In this case this is the right thing to do. However, in another well-known fable, a tortoise says to a hare: "although you be as **swift** as the wind, I have **beaten** you in the race." Here, perfect coreference produces multiple affect states, which *are* related to the plot: the hare recieves a negative affect state for having been beaten in the race.

6 Conclusions

AESOP demonstrates that sentiment analysis tools can successfully recognize many affect states when coupled with syntax-based projection rules to map the affect states onto characters. We also showed that *negative patient polarity verbs* can be harvested from a corpus to identify characters that are in a negative state due to an action. However, performance is still modest, revealing that much work remains to be done. In future work, new methods will be needed to represent affect states associated with plans/goals, events, and inferences.

7 Acknowledgments

The authors thank the anonymous reviewers for many helpful comments. This work was supported in part by the Department of Homeland Security Grant N0014-07-1-0152, the DARPA Machine Reading program under contract FA8750-09-C-0172, and the NSF grant IIS-0712764.

References

Yejin Choi, Eric Breck, and Claire Cardie. 2006. Joint extraction of entities and relations for opinion recognition. In *EMNLP '06: Proceedings of the 2006 Conference on Empirical Methods in Natural Language Processing*, pages 431–439, Morristown, NJ, USA. Association for Computational Linguistics.

S. Kim and E. Hovy. 2006. Extracting Opinions, Opinion Holders, and Topics Expressed in Online News Media Text. In *Proceedings of ACL/COLING Workshop on Sentiment and Subjectivity in Text*.

W. Lehnert, J. Black, and B. Reiser. 1981. Summarizing Narratives. In *Proceedings of the Seventh International Joint Conference on Artificial Intelligence*.

W. G. Lehnert. 1981. Plot Units and Narrative Summarization. *Cognitive Science*, 5(4):293–331.

G. Miller. 1990. Wordnet: An On-line Lexical Database. *International Journal of Lexicography*, 3(4).

E. Riloff and W. Phillips. 2004. An Introduction to the Sundance and AutoSlog Systems. Technical Report UUCS-04-015, School of Computing, University of Utah.

V. Stoyanov and C. Cardie. 2008. Topic Identification for Fine-Grained Opinion Analysis. In *Conference on Computational Linguistics (COLING 2008)*.

Hiroya Takamura, Takashi Inui, and Manabu Okumura. 2005. Extracting semantic orientations of words using spin model. In *ACL '05: Proceedings of the 43rd Annual Meeting on Association for Computational Linguistics*.

A. Wierzbicka. 1987. *English speech act verbs: a semantic dictionary*. Academic Press, Sydney, Orlando.

T. Wilson, P. Hoffmann, S. Somasundaran, J. Kessler, J. Wiebe, Y. Choi, C. Cardie, E. Riloff, and S. Patwardhan. 2005. OpinionFinder: A system for subjectivity analysis. In *Proceedings of HLT/EMNLP 2005 Interactive Demonstrations*.

Theresa Wilson. 2005. Recognizing contextual polarity in phrase-level sentiment analysis. In *In Proceedings of HLT-EMNLP*.