

# Monotonicity Marking from Universal Dependency Trees

Zeming Chen      Qiyue Gao

Department of Computer Science and Software Engineering,  
Rose-Hulman Institute of Technology

{chenz16, gaoq}@rose-hulman.edu

## Abstract

Dependency parsing is a tool widely used in the field of Natural Language Processing and computational linguistics. However, there is hardly any work that connects dependency parsing to monotonicity, which is an essential part of logic and linguistic semantics. In this paper, we present a system that automatically annotates monotonicity information based on Universal Dependency parse trees. Our system utilizes surface-level monotonicity facts about quantifiers, lexical items, and token-level polarity information. We compared our system’s performance with existing systems in the literature, including NatLog and ccg2mono, on a small evaluation dataset. Results show that our system outperforms NatLog and ccg2mono.

## 1 Introduction

The number of computational approaches for Natural Language Inference (NLI) has rapidly grown in recent years. Most of the approaches can be categorized as (1) Systems that translate sentences into first-order logic expressions and then apply theorem proving (Blackburn and Bos, 2005). (2) Systems that use blackbox neural network approaches to learn the inference (Devlin et al., 2019; Liu et al., 2019). (3) Systems that apply natural logic as a tool to make inferences (MacCartney and Manning, 2009; Hu et al., 2020; Angeli et al., 2016; Abzianidze, 2017). Compared to neural network approaches, systems that apply natural logic are more robust, formally more precise, and more explainable. Several systems contributed to the third category (MacCartney and Manning, 2009; Hu et al., 2020; Angeli et al., 2016) to solve the NLI task using monotonicity reasoning, a type of logical inference that is based on word replacement. Below is an example of monotonicity reasoning:

1. (a) All students<sub>↓</sub> carry a MacBook<sub>↑</sub>.

(b) All students carry a laptop.

(c) All new students carry a MacBook.

2. (a) **Not all** new students<sub>↑</sub> carry a laptop.

(b) Not all students carry a laptop.

As the example shows, the word replacement is based on the polarity mark (arrow) on each word. A monotone polarity (↑) allows an inference from (1a) to (1b), where a more general concept *laptop* replaces the more specific concept *MacBook*. An antitone polarity (↓) allows an inference from (1a) to (1c), where a more specific concept *new students* replaces the more general concept *students*. The direction of the polarity marks can be reversed by adding a downward entailment operator like *Not* which allows an inference from (2a) to (2b). Thus, successful word placement relies on accurate polarity marks. To obtain the polarity mark for each word, an automatic polarity marking system is required to annotate a sentence by placing polarity mark on each word. This is formally called the polarization process. Polarity markings support monotonicity reasoning, and thus are used by systems for Natural Language Inference and data augmentations for language models. (MacCartney and Manning, 2009; Hu et al., 2020; Angeli et al., 2016).

In this paper, we introduce a novel automatic polarity marking system that annotates monotonicity information by applying a polarity algorithm on a universal dependency parse tree. Our system is inspired by ccg2mono, an automatic polarity marking system (Hu and Moss, 2018) used by Hu et al. (2020). In contrast to ccg2mono, which derives monotonicity information from CCG (Lewis and Steedman, 2014) parse trees, our system’s polarization algorithm derives monotonicity information using Universal Dependency (Nivre et al., 2016) parse trees. There are several advantages of using UD parsing for polarity marking rather than

CCG parsing. First, UD parsing is more accurate since the amount of training data for UD parsing is larger than those of CCG parsing. The high accuracy of UD parsing should lead to more accurate polarity annotation. Second, UD parsing works for more types of text. Overall, our system opens up a new framework for performing inference, semantics, and automated reasoning over UD representations. We will introduce the polarization algorithm’s general steps, a set of rules we used to mark polarity on dependency parse trees, and comparisons between our system and some existing polarity marking tools, including NatLog (MacCartney and Manning, 2009; Angeli et al., 2016) and ccg2mono. Our evaluation focuses on a small dataset used to evaluate ccg2mono (Hu and Moss, 2020). Our system outperforms NatLog and ccg2mono. In particular, our system achieves the highest annotation accuracy on both the token level and the sentence level.

## 2 Related Work

Universal Dependencies (UD) (Nivre et al., 2016) was first designed to handle language tasks for many different languages. The syntactic annotation in UD mostly relies on dependency relations. Words enter into dependency relations, and that is what UD tries to capture. There are 40 grammatical dependency relations between words, such as nominal subject (**nsubj**), relative clause modifier (**acl:recl**), and determiner (**det**). A dependency relation connects a headword to a modifier. For example, in the dependency parse tree for *All dogs eat food* (figure 1), the dependency relation **nsubj** connects the modifier *dogs* and the headword *eat*. The system presented in this paper utilizes Universal Dependencies to obtain a dependency parse tree from a sentence. We will explain the details of the parsing process in the implementation section.

There are two relevant systems of prior work: (1) The NatLog (MacCartney and Manning, 2009; Angeli et al., 2016) system included in the Stanford CoreNLP library (Manning et al., 2014); (2) The ccg2mono system (Hu and Moss, 2018). The NatLog system is a natural language inference system, a part of the Stanford CoreNLP Library. NatLog marks polarity to each sentence by applying a pattern-based polarization algorithm to the dependency parse tree generated by the Stanford dependency parser. A list of downward-monotone and non-monotone expressions are defined along

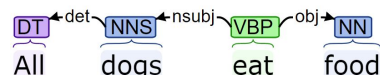


Figure 1: A dependency parse tree for "All dogs eat food."

with an arity and a Tregex pattern for the system to identify if an expression occurred.

The ccg2mono system is a polarity marking tool that annotates a sentence by polarizing a CCG parse tree. The polarization algorithm of ccg2mono is based on van Benthem (1986)’s work and Moss (2012)’s continuation on the soundness of internalized polarity marking. The system uses a marked/order-enriched lexicon and can handle application rules, type-raising, and composition in CCG. The main polarization contains two steps: mark and polarize. For the mark step, the system puts markings on each node in the parse tree from leaf to root. For the polarize step, the system generates polarities to each node from root to leaf. Compared to NatLog, an advantage of ccg2mono is that it polarizes on both the word-level and the constituent level.

## 3 Universal Dependency to Polarity

### 3.1 Overview

Our system’s polarization algorithm contains three steps: (1) Universal Dependency Parsing, which transforms a sentence to a UD parse tree, (2) Binarization, which converts a UD parse tree to a binary UD parse tree, and (3) Polarization, which places polarity marks on each node in a binary UD parse tree.

### 3.2 Binarization

To preprocess the dependency parse graph, we designed a binarization algorithm that can map each dependency tree to an s-expression (Reddy et al., 2016). Formally, an s-expression has the form (exp1 exp2 exp3), where exp1 is a dependency label, and both exp2 and exp3 are either (1) a word such as *eat*; or (2) an s-expression such as (**det** *all dogs*). The process of mapping a dependency tree to an s-expression is called binarization. Our system represents an s-expression as a binary tree. A binary tree has a root node, a left child node, and a right child node. In representing an s-expression, the root node can either be a single word or a dependency label. Both the left and the right child nodes can either be a sub-binary-tree, or null. The

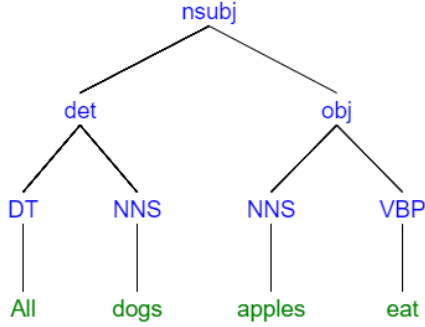


Figure 2: A binarized dependency parse tree for "All dogs eat apples."

system always puts the modifiers on the left and the headwords on the right. For example, the sentence *All dogs eat apples* has an s-expression

**(nsubj (det All dogs) (obj eat apples))**

and can be shown as a binary tree in figure 2. In the left sub-tree (*All dogs*), the dependency label **det** will be the root node, the modifier *all* will be the left child, and the headword *dogs* will be the right child.

Our binarization algorithm employs a dependency relation hierarchy to impose a strict traversal order from the root relation to each leaf word. The hierarchy allows for an ordering on the different modifier words. For example, in the binary dependency parse tree **(nsubj (det All dogs) (obj eat**

relation	level-id	relation	level-id
conj-sent	0	obl:tmod	50
advcl-sent	1	obl:npm	50
advmod-sent	2	cop	50
case	10	det	55
mark	10	det:predet	55
expl	10	acl	60
discourse	10	acl:relel	60
nsubj	20	appos	60
csbj	20	conj	60
nsubj:pass	20	conj-np	60
conj-vp	25	conj-adj	60
ccomp	30	obj	60
advcl	30	iobj	60
advmod	30	cc	70
nmod	30	amod	75
nmod:tmod	30	nummod	75
nmod:npm	30	compound	80
nmod:poss	30	compound:prt	80
xcomp	40	fixed	80
aux	40	conj-n	90
aux:pass	40	conj-vb	90
obl	50	flat	100

Table 1: Universal Dependency relation hierarchy. The smaller a relation’s level-id is, the higher that relation is in the hierarchy.

apples)), the nominal subject (**nsubj**) goes above the determiner (**det**) in the tree because **det** is lower than **nsubj** in the hierarchy. We originally used the binarization hierarchy from Reddy et al. (2016)’s work, and later extended it with additional dependency relations such as oblique nominal (**obl**) and expletive (**expl**). Table 1 shows the complete hierarchy where the level-id indicates a relation’s level in the hierarchy. The smaller a relation’s level-id is, the higher that relation is in the hierarchy.

### Algorithm 1 Binarization

---

```

1: root ← GET_ROOT_NODE(G)
2: T ← COMPOSE(root)
3: return T
4:
5: function COMPOSE(node):
6:   C ← GET_CHILDREN(node)
7:   Cs ← SORT_BY_PRIORITY(C)
8:   if |Cs| == 0 then
9:     B ← BINARYDEPENDENCYTREE()
10:    B.val = node
11:    return B
12:  else
13:    top ← C.pop()
14:    B ← BINARYDEPENDENCYTREE()
15:    B.val = RELATE(top, node)
16:    B.left = COMPOSE(top)
17:    B.right = COMPOSE(node)
18:    return B
19:  end if
20: end function

```

---

### 3.3 Polarization

The polarization algorithm places polarities on each node of a UD parse tree based on a lexicon of polarization rules for each dependency relation and some special words. Our polarization algorithm is similar to the algorithms surveyed by Lavalle-Martínez et al. (2018). Like the algorithm of Sanchez (1991), our algorithm computes polarity from leaves to root. One difference our algorithm has is that often, the algorithm computes polarity following a left-to-right inorder traversal (left→root→right) or a right-to-left inorder traversal (right→root→left) in addition to the top-down traversal. In our algorithm, each node’s polarity depends both on its parent node and its sibling node (left side or right side), which is different from algorithms in Lavalle-Martínez et al. (2018)’s paper. Our algorithm is deterministic, and thus never fails.

The polarization algorithm takes in a binarized UD parse tree  $\mathcal{T}$  and a set of polarization rules, both dependency-relation-level ( $\mathcal{L}$ ) and word-level ( $\mathcal{W}$ ). The algorithm outputs a polarized UD parse tree  $\mathcal{T}^*$  such that (1) each node is marked with

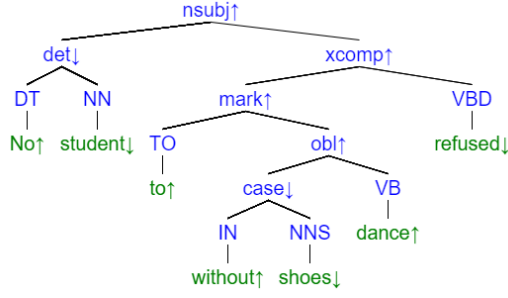


Figure 3: Visualization of a polarized binary dependency parse tree for a triple negation sentence *No student refused to dance without shoes.*

a polarity of either monotone ( $\uparrow$ ), antitone ( $\downarrow$ ), or no monotonicity information ( $=$ ), (2) both  $\mathcal{T}$  and  $\mathcal{T}^*$  have the same universal dependency structure except the polarity marks. Figure 3 shows a visualization of the binary dependency parse tree after polarization completes. The general steps of the polarization start from the root node of the binary parse tree. The system will get the corresponding polarization rule from the lexicon according to the root node’s dependency relation. In each polarization rule, the system applies the polarization rule and then continues the above steps recursively down the left sub-tree and the right sub-tree. Each polarization rule is composed from a set of basic building blocks include rules for negation, equalization, and monotonicity generation. When the recursion reaches a leaf node, which is an individual word in a sentence, a set of word-based polarization rules will be retrieved from the lexicon, and the system polarizes the nodes according to the rule corresponding to a particular word. More details about word-based polarization rules will be covered in section 3.4.2, Polarity Generation. An overview of the polarization algorithm and a general scheme of the implementation for dependency-level polarization rules are shown in Algorithm 2.

### 3.4 Polarization Rules

Our polarization algorithm contains a lexicon of polarization rules corresponding to each dependency relation. Each polarization rule is composed from a set of building blocks divided into three categories: negation rules, equalization rules, and monotonicity generation rules. The generation rules will generate three types of monotonicity: monotone ( $\uparrow$ ), antitone ( $\downarrow$ ), and no monotonicity information ( $=$ ) either by initialization or based on the words.

---

### Algorithm 2 Polarization

---

**Input:**  $\mathcal{T}$ : binary dependency tree  
 $\mathcal{L}$ : dependency-level polarization rules  
 $\mathcal{W}$ : word-level polarization rules  
**Output:**  $\mathcal{T}^*$ : polarized binary dependency tree

```

1: if  $\mathcal{T}$ .is_tree then
2:   relation  $\leftarrow \mathcal{T}$ .val
3:   POLARIZATION_RULE(.)  $\leftarrow \mathcal{L}$ [relation]
4:   POLARIZATION_RULE( $\mathcal{T}$ )
5: end if
6:
7:  $\triangleright$  General scheme of a polarization rule’s implementation for a dependency relation
8: function POLARIZATION_RULE( $\mathcal{T}$ )
9:    $\triangleright$  Initialize or inherit polarities
10:  if  $\mathcal{T}$ .mark  $\neq$  NULL then
11:     $\mathcal{T}$ .right.mark =  $\mathcal{T}$ .mark
12:     $\mathcal{T}$ .left.mark =  $\mathcal{T}$ .mark
13:  else
14:     $\mathcal{T}$ .right.mark =  $\uparrow$ 
15:     $\mathcal{T}$ .left.mark =  $\uparrow$ 
16:  end if
17:
18:   $\triangleright$  Polarize sub-trees
19:  POLARIZATION( $\mathcal{T}$ .left)
20:  POLARIZATION( $\mathcal{T}$ .right)
21:   $\triangleright$  Or, for relations like nsubj:
22:   $\triangleright$  POLARIZATION( $\mathcal{T}$ .right)
23:   $\triangleright$  POLARIZATION( $\mathcal{T}$ .left)
24:
25:   $\triangleright$  Apply negation and equalization rules
26:  if NEGATE is applicable then
27:    NEGATE( $\mathcal{T}$ )
28:  end if
29:  if EQUALIZE is applicable then
30:    EQUALIZE( $\mathcal{T}$ )
31:  end if
32:
33:   $\triangleright$  Apply word-level rules
34:  if not  $\mathcal{T}$ .is_tree and  $\mathcal{T}$ .val  $\in \mathcal{W}$ .keys then
35:    WORD_RULE(.)  $\leftarrow \mathcal{W}$ [ $\mathcal{T}$ .val]
36:    WORD_RULE( $\mathcal{T}$ )
37:  end if
38: end function

```

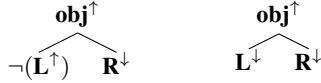
---

#### 3.4.1 Building Blocks

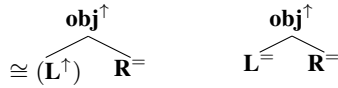
**Negation and Equalization** The negation rule and the equalization rule are used by several core dependency relations such as **nmod**, **obj**, and **acl:recl**. Both negation and equalization have two ways of application: backward or top-down. A backward negation rule is triggered by a downward polarity ( $\downarrow$ ) on the right node of the tree (marked below as R), flipping every node’s polarity under the left node (marked below as L). Similarly, a backward equalization rule is triggered by a no monotonicity information polarity ( $=$ ) on the tree’s right node, and it marks every node under the left node as  $=$ . Examples for trees before and after applying a backward and forward negation and equalization are shown as follows:



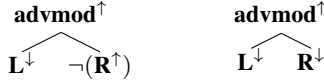
- Backward Negation:



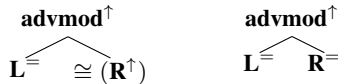
- Backward Equalization:



- Forward Negation:

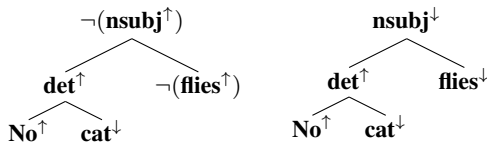


- Forward Equalization:



where  $\neg$  means negation and  $\cong$  means equalization.

A top-down negation is used by the polarization rule like determiner (**det**) and adverbial modifier (**advmod**). It starts at the parent node of the current tree, and flips the arrow on each node under that parent node excluding the current tree. This top-down negation is used by **det**, **case**, and **advmod** when a negation operators like *no*, *not*, or *at-most* appears. Below is an example of a tree before and after applying the top-down negation:



**Polarity Generation** The polarity is generated by words. During the polarization, the polarity can change based on a particular word that can promote the polarity governing the part of the sentence to which it belongs. These words include quantifiers and verbs. For the monotonicity from quantifiers, we follow the monotonicity profiles listed in the work done by [Icard III and Moss \(2014\)](#) on monotonicity, which built on [van Benthem \(1986\)](#). Additionally, to extend to more quantifiers, we observed polarization results generated by *ccg2mono*. Overall, we categorized the quantifiers as follows:

- Universal Type

Every  $\downarrow \uparrow$  Each  $\downarrow \uparrow$  All  $\downarrow \uparrow$

- Negation Type

No  $\downarrow \downarrow$  Less than  $\downarrow \downarrow$  At most  $\downarrow \downarrow$

- Exact Type

Exactly  $n = =$  The  $= \uparrow$  This  $= \uparrow$

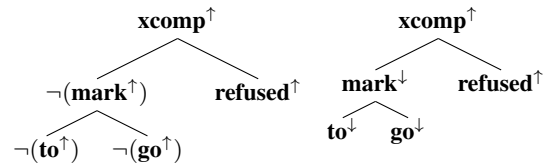
- Existential Type

Some  $\uparrow \uparrow$  Several  $\uparrow \uparrow$  A, An  $\uparrow \uparrow$

- Other Type

Most  $= \uparrow$  Few  $= \downarrow$

Where the first mark is the monotonicity for the first argument after the quantifier and the second mark is the monotonicity for the second argument after the quantifier. For verbs, there are upward entailment operators and downward entailment operators. Verbs that are downward entailment operators, such as *refuse*, promote an antitone polarity, which will negate its dependents. For example, for the phrase *refused to go*, *refused* will promote an antitone polarity, which negates *to dance*:



In addition to quantifiers and verbs, some other words also change the monotonicity of a sentence. For example, words like *not*, *none*, and *nobody* promote an antitone polarity. Our system also handles material implications with the form *if x then y*. Based on [Moss \(2012\)](#), the word *if* promotes an antitone polarity in the antecedent and positive polarity in the consequent. For background on monotonicity and semantics, see [van Benthem \(1986\)](#), [Keenan and Faltz \(1984\)](#), and also [Karttunen \(2012\)](#).

### 3.4.2 Dependency Relation Rules

Each dependency relation has a corresponding polarization rule. All the rules start with initializing the starting node as upward monotone polarity ( $\uparrow$ ). Alternatively, if the starting node has a polarity marked, each child node will inherit the root node's polarity. Each rule's core part is a combination of the default rules and monotonicity generation rules. In this section, we will briefly show three major types of dependency relation rules in the polarization algorithm. The relative clause modifier relation will represent rules for modifier relations. The determiner relation rule will represent rules containing monotonicity generation rules. The Object and open clausal complement rule will represent rules containing word-level polarization rules.

---

**Algorithm 3** Polarize\_acl:relcl

---

**Input:**  $\mathcal{T}$ : binary dependency sub-tree**Output:**  $\mathcal{T}^*$ : polarized binary dependency sub-tree

```
1: if  $\mathcal{T}.\text{mark} \neq \text{NULL}$  then
2:    $\mathcal{T}.\text{right.mark} = \mathcal{T}.\text{mark}$ 
3: else
4:    $\mathcal{T}.\text{right.mark} = \uparrow$ 
5: end if
6:  $\mathcal{T}.\text{left.mark} = \uparrow$ 
7:
8: POLARIZE( $\mathcal{T}.\text{right}$ )
9: POLARIZE( $\mathcal{T}.\text{left}$ )
10:
11: if  $\mathcal{T}.\text{right.mark} == \downarrow$  then
12:   NEGATE( $\mathcal{T}.\text{left}$ )
13: else if  $\mathcal{T}.\text{right.mark} == =$  then
14:   EQUALIZE( $\mathcal{T}.\text{left}$ )
15: end if
```

---

**Relative Clause Modifier** For the relative clause modifier relation (**acl:relcl**), the relative clause depends on the noun it modifies. First, the polarization will first be performed on both the left and right nodes, and then, depending on the polarity of the right node, a negation or an equalization rule will be applied. The algorithm first applies a top-down inheritance if the root already has its polarity marked; otherwise, it initializes the left and right nodes as monotone. The algorithm polarizes both the left and right nodes. Next, the algorithm checks the right node’s polarity. If the right node is marked as antitone, a backward negation is applied. Alternatively, if the right node is marked as no monotonicity information, a backward equalization is applied. During the experiments, we noticed that if the root node is marked antitone, and the left node inherits that, a negation later will cause a double negation, producing incorrect polarity marks. To avoid this double negation, we exclude the left node from the top-down inheritance rule by initializing the left node directly with a monotone mark. The rule for **acl:relcl** also applies to the adverbial clause modifier (**advcl**) and the clausal modifier of noun (**acl**). An overview of the algorithm is shown in Algorithm 3.

**Determiner** For the determiner relation (**det**), each different determiner can assign a new monotonicity to the noun it modifies. First, the algorithm performs a top-down inheritance on the left node if the root already has polarity marked. Next, the algorithm assigns the polarity for the noun depending on the determiner’s type. For example, if the determiner is a universal quantifier, an antitone polarity is assigned to the right node. For negation quanti-

---

**Algorithm 4** Polarize\_det

---

**Input:**  $\mathcal{T}$ : binary dependency sub-tree $\mathcal{D}$ : determiner mark dictionary**Output:**  $\mathcal{T}^*$ : polarized binary dependency sub-tree

```
1:  $\text{det\_type} \leftarrow \text{GET\_DET\_TYPE}(\mathcal{T}.\text{left})$ 
2: if  $\mathcal{T}.\text{mark} \neq \text{NULL}$  then
3:    $\mathcal{T}.\text{left.mark} = \mathcal{T}.\text{mark}$ 
4: else
5:    $\mathcal{T}.\text{left.mark} = \uparrow$ 
6: end if
7:
8:  $\mathcal{T}.\text{right.mark} = \mathcal{D}[\text{det\_type}]$ 
9: POLARIZE( $\mathcal{T}.\text{right}$ )
10:
11: if  $\text{det\_type} == \text{negation}$  then
12:   NEGATE( $\mathcal{T}.\text{parent}$ )
13: end if
```

---

fiers like *no*, its right node also receives an antitone polarity. Thus, a top-down negation is applied at the determiner relation tree’s parent. Algorithm 4 shows an overview of the algorithm.

**Object and Open Clausal Complement** For the object relation (**obj**) and the open clausal complement relation **xcomp**, both the verb and the noun would inherit the monotonicity from the parent in the majority of cases. The inheritance procedure is the same as the one used in **acl:relcl**’s rule. Similarly, after the inheritance, the rule will polarize both the right sub-tree and the left sub-tree. Differently, since **obj** and **xcomp** both have a verb under the relation, they require a word-level polarization rule that will check the verb determine if the verb is a downward entailment operator, which prompts an antitone monotonicity. The algorithm takes in a dictionary that contains a list of verbs and their

---

**Algorithm 5** Polarize\_obj

---

**Input:**  $\mathcal{T}$ : binary dependency sub-tree**Output:**  $\mathcal{T}^*$ : polarized binary dependency sub-tree

```
1: if  $\mathcal{T}.\text{mark} \neq \text{NULL}$  then
2:    $\mathcal{T}.\text{right.mark} = \mathcal{T}.\text{mark}$ 
3: else
4:    $\mathcal{T}.\text{right.mark} = \uparrow$ 
5: end if
6:  $\mathcal{T}.\text{left.mark} = \uparrow$ 
7:
8: POLARIZE( $\mathcal{T}.\text{right}$ )
9: POLARIZE( $\mathcal{T}.\text{left}$ )
10:
11:  $\triangleright$  Word-level polarization rule for downward entailment operators
12: if IS_DOWNWARD_OPERATOR( $\mathcal{T}.\text{right.mark}$ ) then
13:   NEGATE( $\mathcal{T}.\text{left}$ )
14: end if
15:
```

---

implicatives. The dictionary is generated from the implicative verb dataset made by [Ross and Pavlick \(2019\)](#). If a verb is a downward entailment operator, which has a negative implicative, the rule will apply a negation rule on the left sub-tree to flip each node’s arrow in the left sub-tree. An overview of the algorithm is shown in Algorithm 5.

#### 4 Comparison to Existing Systems

We conducted several preliminary comparisons to two existing systems. First, we compared to NatLog’s monotonicity annotator. Natlog’s annotator also uses dependency parsing. The polarization algorithm does pattern-based matching for finding occurrences of downward monotonicity information, and the algorithm only polarizes on word-level. In contrast, our system uses a tree-based polarization algorithm that polarizes both on word-level polarities and constituent level polarities. Our intuition is that the Tregex patterns used in NatLog is not as common or as easily understandable as the binary tree structure, which is a classic data structure wildly used in the field of computer science.

According to the comparison on a list of sentences, NatLog’s annotator does not perform as well as our system. For example, for a phrase *the rabbit, rabbit* should have a polarity with no monotonicity information (=). However, NatLog marks *rabbit* as a monotone polarity (↑). NatLog also incorrectly polarizes sentences containing multiple negations. For example, for a triple negation sentence, *No newspapers did not report no bad news*, NatLog gives: *No<sup>↑</sup> newspapers<sup>↓</sup> did<sup>↓</sup> not<sup>↓</sup> report<sup>↑</sup> no<sup>↑</sup> bad<sup>↑</sup> news<sup>↑</sup>*. This result has incorrect polarity marks on multiple words, where *report*, *bad*, *news* should be ↓, and *no* should be ↑. Both of the scenarios above can be handled correctly by our system.

Comparing to *ccg2mono*, our algorithm shares some similarities to its polarization algorithm. Both of the systems polarize on a tree structure and rely on a lexicon of rules, and they both polarize on the word-level and the constituent level. One difference is that *ccg2mono*’s algorithm contains two steps, the first step puts markings on each node, and the second step puts polarities on each node. Our system does not require the step of adding markings and only contains the step of adding polarities on each node.

Our system has multiple advantages over *ccg2mono*. For parsing, our system uses UD pars-

ing, which is more accurate than CCG parsing used by *ccg2mono* due to a large amount of training data. Also, our system covers more types of text than *ccg2mono* because UD parsing works for a variety of text genres such as web texts, emails, reviews, and even informal texts like Twitter tweets. ([Silveira et al., 2014](#); [Zeldes, 2017](#); [Liu et al., 2018](#)). Our system can also work for more languages than *ccg2mono* since UD parsing supports more languages than CCG parsing.

Overall, our system delivers more accurate polarization than *ccg2mono*. Many times the CCG parser makes mistakes and leads to polarization mistakes later on. For example, in the annotation *The<sup>↓</sup> market<sup>↓</sup> is<sup>↓</sup> not<sup>↓</sup> impossible<sup>↓</sup> to<sup>↓</sup> navigate<sup>↓</sup>*, *ccg2mono* incorrectly marks every word as ↓. Our system, on the other hand, uses UD parsing which has higher parsing accuracy than CCG parsing, and thus leads to fewer polarization mistakes compared to *ccg2mono*. For the expression above, our system correctly polarizes it as *The<sup>↑</sup> market<sup>=</sup> is<sup>↑</sup> not<sup>↑</sup> impossible<sup>↓</sup> to<sup>↑</sup> navigate<sup>↑</sup>*.

Our system also handles multi-word quantifiers better than *ccg2mono*. For example, for a multi-word quantifier expression like *all of the dogs*, *ccg2mono* mistakenly marks *dogs* as =. Our system, however, can correctly mark the expression: *all<sup>↑</sup> of<sup>↑</sup> the<sup>↑</sup> dogs<sup>↓</sup>*.

Moreover, the core of *ccg2mono* does not include aspects of verbal semantics of downward-entailing operators like *forgot* and *regret* ([Moss and Hu, 2020](#)). For example *ccg2mono*’s polarization for *Every<sup>↑</sup> member<sup>↓</sup> forgot<sup>↑</sup> to<sup>↑</sup> attend<sup>↑</sup> the<sup>↑</sup> meeting<sup>=</sup>* is not correct because it fails to flip the polarity of *to attend the*. In contrast, our system produces a correct result: *Every<sup>↑</sup> member<sup>↓</sup> forgot<sup>↑</sup> to<sup>↓</sup> attend<sup>↓</sup> the<sup>↓</sup> meeting<sup>=</sup>*.

All three systems have difficulty polarizing sentences containing numbers. A scalar number *n*’s monotonicity information is hard to determine because it can present different contexts: a single number *n*, without additional quantifiers or adjectives, can either mean *at least n*, *at most n*, *exactly n*, and *around n*. These contexts are syntactically hard to identify for a dependency parser or a CCG parser because it would require pragmatics and some background knowledge which the parsers do not have. For example, in the sentence *A dog ate 2 rotten biscuits*, the gold label for 2 is = which indicates that the context is "exactly 2". However, our system marks this as "↓" since it considers the

sentence	type
More <sup>↑</sup> dogs <sup>↑</sup> than <sup>↑</sup> cats <sup>↓</sup> sit <sup>=</sup>	comparative
Less <sup>↑</sup> than <sup>↑</sup> 5 <sup>↑</sup> people <sup>↓</sup> ran <sup>↓</sup>	less-than
A <sup>↑</sup> dog <sup>↑</sup> who <sup>↑</sup> ate <sup>↑</sup> two <sup>=</sup> rotten <sup>↑</sup> biscuits <sup>↑</sup> was <sup>↑</sup> sick <sup>↑</sup> for <sup>↑</sup> three <sup>↓</sup> days <sup>↓</sup>	number
Every <sup>↑</sup> dog <sup>↓</sup> who <sup>↓</sup> likes <sup>↓</sup> most <sup>↓</sup> cats <sup>=</sup> was <sup>↑</sup> chased <sup>↑</sup> by <sup>↑</sup> at <sup>↑</sup> least <sup>↑</sup> two <sup>↓</sup> of <sup>↑</sup> them <sup>↑</sup>	every:most:at-least
Even <sup>↑</sup> if <sup>↑</sup> you <sup>↓</sup> are <sup>↓</sup> addicted <sup>↓</sup> to <sup>↓</sup> cigarettes <sup>↓</sup> you <sup>↑</sup> can <sup>↑</sup> smoke <sup>↑</sup> two <sup>↓</sup> a <sup>↑</sup> day <sup>↑</sup>	conditional:number

Table 2: Example sentences in Hu and Moss (2020)’s evaluation dataset

context as "at least 2", which is different from the gold label.

## 5 Experiment

**Dataset** We obtained the small evaluation dataset used in the evaluation of ccg2mono (Hu and Moss, 2020) from its authors. The dataset contains 56 hand-crafted English sentences, each with manually annotated monotonicity information. The sentences cover a wide range of linguistic phenomena such as quantifiers, conditionals, conjunctions, and disjunctions. The dataset also contains hard sentences involving scalar numbers. Some example sentences from the dataset are shown in Table 2.

**Dependency Parser** In order to obtain a universal dependency parse tree from a sentence, we utilize a parser from Stanza (Qi et al., 2020), a Python natural language analysis package made by Stanford. The neural pipeline in Stanza allow us to use pretrained neural parsing models to generate universal dependency parse trees. To achieve optimal performance, we trained two neural parsing models: one parsing model trained on Universal Dependency English GUM corpus (Zeldes, 2017). The pretrained parsing model achieved 90.0 LAS (Zeman et al., 2018) evaluation score on the testing data.

**Experiment Setup** We evaluated the polarization accuracy on both the token level and the sentence level, in a similar fashion to the evaluation for part-of-speech tagging (Manning, 2011). For both levels of accuracy, we conducted one evaluation on all tokens (*acc(all-tokens)* in Table 3) and another one on key tokens including content words (nouns, verbs, adjectives, adverbs), determiners, and numbers (*acc(key-tokens)* in Table 3). The key tokens contain most of the useful monotonicity information for inference. In token-level evaluation, we counted the number of correctly annotated tokens for *acc(all-tokens)* or the number of correctly annotated key tokens for *acc(key-tokens)*. In sentence-level evaluation, we counted the number of cor-

system	Token-level		
	NatLog	ccg2mono	ours
acc(all-tokens)	69.9	76.0	<b>96.5</b>
acc(key-tokens)	68.1	78.0	<b>96.5</b>
system	Sentence-level		
	NatLog	ccg2mono	ours
acc(all-tokens)	28.0	44.6	<b>87.5</b>
acc(key-tokens)	28.6	50.0	<b>89.2</b>

Table 3: This table shows the polarity annotation accuracy on the token level and the sentence level for three systems: NatLog, ccg2mono, and our system. The token level accuracy counts the number of correctly annotated tokens, and the sentence level accuracy counts the number of correctly annotated sentences. Two types of accuracy are used. For *acc(all-tokens)*, all tokens are evaluated. For *acc(key-tokens)*, only key tokens (content words + determiners + numbers) are evaluated.

rect sentences. A correct sentence has all tokens correctly annotated for *acc(all-tokens)* or all key tokens correctly annotated for *acc(key-tokens)*. We also evaluated our system’s robustness on the token level. We followed the robustness metric for evaluating multi-class classification tasks, which uses precision, recall, and F1 score to measure a system’s robustness. We calculated these three metrics for each polarity label: monotone(↑), antitone(↓), and None or no monotonicity information(=). The robustness evaluation is also done both on all tokens and on key tokens.

## 6 Evaluation

Table 3 shows the performance of our system, compared with NatLog and ccg2mono. Our evaluation process is the same as Hu and Moss (2020). From Table 3, we first observe that our system consistently outperforms ccg2mono and NatLog on both the token level and the sentence level. For accuracy on the token level, our system has the highest accuracy for the evaluation on all tokens (96.5) and the highest accuracy for the evaluation on key tokens (96.5). Our system’s accuracy on key tokens is higher than the accuracy on all tokens, which demonstrates our system’s good performance on polarity annotation for tokens that are more signif-



		All Tokens								
system	NatLog			cgg2mono			ours			
Polarity	Monotone	Antitone	None	Monotone	Antitone	None	Monotone	Antitone	None	
precision	71.4	43.5	70.7	86.0	75.6	58.0	<b>97.6</b>	<b>96.5</b>	<b>91.7</b>	
recall	87.3	15.9	63.9	77.8	78.3	74.6	<b>97.2</b>	<b>89.4</b>	<b>87.3</b>	
F1-score	78.6	23.3	67.1	81.7	76.9	65.3	<b>97.4</b>	<b>97.6</b>	<b>89.4</b>	
		Key Tokens								
system	NatLog			cgg2mono			ours			
Polarity	Monotone	Antitone	None	Monotone	Antitone	None	Monotone	Antitone	None	
precision	68.7	70.9	42.1	85.2	78.7	62.7	<b>96.9</b>	<b>96.4</b>	<b>94.2</b>	
recall	88.6	61.5	14.0	80.3	79.3	73.7	<b>97.9</b>	<b>98.5</b>	<b>86.0</b>	
F1-score	77.4	65.9	21.1	82.7	79.0	67.7	<b>97.4</b>	<b>97.4</b>	<b>89.9</b>	

Table 4: Token level robustness comparison between NatLog, cgg2mono, and our system. The robustness score is evaluated both on all tokens and on key tokens (content words + determiners + numbers). For each of the three polarities: monotone( $\uparrow$ ), antitone( $\downarrow$ ), and None or no monotonicity information( $=$ ), the relative precision, recall and F1 score are calculated.

icant to monotonicity inference. For accuracy on the sentence level, our system again has the highest accuracy for the evaluation on all tokens (87.5) and the highest accuracy for the evaluation on key tokens (89.2). Such results suggest that our system can achieve good performance on determining the monotonicity of the sentence constituents. Overall, the evaluation validates that our system has higher polarity annotation accuracy than existing systems. We compared our annotations to cgg2mono’s annotation and observed that of all the tokens in the 56 sentences, if cgg2mono annotates it correctly, then our system also does so. This means, our system’s polarization covers more linguistic phenomena than cgg2mono. Table 4 shows the robustness score of our system and the two existing systems. Our systems has much higher precision and recall on all three polarity labels than the other two systems. For the F1 score, our system again has the highest points over the other two systems. The consistent and high robustness scores show that our system’s performance is much more robust on the given dataset than existing systems.

## 7 Conclusion and Future Work

In this paper, we have demonstrated our system’s ability to automatically annotate monotonicity information (polarity) for a sentence by conducting polarization on a universal dependency parse tree. The system operates by first converting the parse tree to a binary parse tree and then marking polarity on each node according to a lexicon of polarization rules. The system produces accurate annotations on sentences involving many different linguistic phenomena such as quantifiers, double negation, relative clauses, and conditionals. Our

system had better performance on polarity marking than existing systems including cgg2mono (Hu and Moss, 2018) and NatLog (MacCartney and Manning, 2009; Angeli et al., 2016). Additionally, by using UD parsing, our system offers many advantages. Our system supports a variety of text genres and can be applied to many languages. In general, this paper opens up a new framework for performing inference, semantics, and automated reasoning over UD representations.

For future work, an inference system can be made that utilizes the monotonicity information annotated by our system, which is similar to the Mon-aLog system (Hu et al., 2020). Several improvements can be made to the system to obtain more accurate annotations. One improvement would be to incorporate pragmatics to help determine the monotonicity of a scalar number.

## Acknowledgements

This research is advised by Dr. Lawrence Moss from Indiana University and Dr. Michael Wollowski from Rose-hulman Institute of Technology. We thank their helpful advises and feedback on this research. We also thank the anonymous reviewers for their insightful comments.

## References

- Lasha Abzianidze. 2017. [LangPro: Natural language theorem prover](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 115–120, Copenhagen, Denmark. Association for Computational Linguistics.
- Gabor Angeli, Neha Nayak, and Christopher D. Manning. 2016. [Combining natural logic and shallow](#)

- reasoning for question answering. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 442–452, Berlin, Germany. Association for Computational Linguistics.
- Johan van Benthem. 1986. *Essays in Logical Semantics*, volume 29 of *Studies in Linguistics and Philosophy*. D. Reidel Publishing Co., Dordrecht.
- P. Blackburn and Johan Bos. 2005. Representation and inference for natural language - a first course in computational semantics. In *CSLI Studies in Computational Linguistics*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. **BERT: Pre-training of deep bidirectional transformers for language understanding**. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Hai Hu, Qi Chen, Kyle Richardson, Atreyee Mukherjee, Lawrence S Moss, and Sandra Kübler. 2020. MonaLog: a lightweight system for natural language inference based on monotonicity. In *Proceedings of the Society for Computation in Linguistics (SCiL) 2020*, pages 319–329.
- Hai Hu and Larry Moss. 2018. **Polarity computations in flexible categorial grammar**. In *Proceedings of the Seventh Joint Conference on Lexical and Computational Semantics*, pages 124–129, New Orleans, Louisiana. Association for Computational Linguistics.
- Hai Hu and Lawrence S. Moss. 2020. **An automatic monotonicity annotation tool based on ccg trees**. In *Second Tsinghua Interdisciplinary Workshop on Logic, Language, and Meaning: Monotonicity in Logic and Language*.
- Thomas F. Icard III and Lawrence S. Moss. 2014. **Recent progress on monotonicity**. In *Linguistic Issues in Language Technology, Volume 9, 2014 - Perspectives on Semantic Representations for Textual Inference*. CSLI Publications.
- Lauri Karttunen. 2012. Simple and phrasal implicatives. In *Proceedings of the First Joint Conference on Lexical and Computational Semantics - Volume 1: Proceedings of the Main Conference and the Shared Task, and Volume 2: Proceedings of the Sixth International Workshop on Semantic Evaluation, SemEval '12*, page 124–131, USA. Association for Computational Linguistics.
- Edward L. Keenan and Leonard M. Faltz. 1984. *Boolean Semantics for Natural Language*. Springer.
- J. Lavallo-Martínez, M. Montes y Gómez, L. Pineda, Héctor Jiménez-Salazar, and Ismael Everardo Bárcenas Patiño. 2018. Equivalences among polarity algorithms. *Studia Logica*, 106:371–395.
- Mike Lewis and Mark Steedman. 2014. **A\* CCG parsing with a supertag-factored model**. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 990–1000, Doha, Qatar. Association for Computational Linguistics.
- Wei Liu, Lei Li, Zuying Huang, and Yinan Liu. 2019. **Multi-lingual Wikipedia summarization and title generation on low resource corpus**. In *Proceedings of the Workshop MultiLing 2019: Summarization Across Languages, Genres and Sources*, pages 17–25, Varna, Bulgaria. INCOMA Ltd.
- Yijia Liu, Yi Zhu, Wanxiang Che, Bing Qin, Nathan Schneider, and Noah A. Smith. 2018. **Parsing tweets into Universal Dependencies**. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 965–975, New Orleans, Louisiana. Association for Computational Linguistics.
- Bill MacCartney and Christopher D. Manning. 2009. **An extended model of natural logic**. In *Proceedings of the Eight International Conference on Computational Semantics*, pages 140–156, Tilburg, The Netherlands. Association for Computational Linguistics.
- Christopher Manning, Mihai Surdeanu, John Bauer, Jenny Finkel, Steven Bethard, and David McClosky. 2014. **The Stanford CoreNLP natural language processing toolkit**. In *Proceedings of 52nd Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 55–60, Baltimore, Maryland. Association for Computational Linguistics.
- Christopher D. Manning. 2011. Part-of-speech tagging from 97linguistics? In *CICLing*.
- L. Moss. 2012. The soundness of internalized polarity marking. *Studia Logica*, 100:683–704.
- Lawrence S. Moss and Hai Hu. 2020. Syllogistic logics with comparative adjectives. Unpublished ms., Indiana University.
- Joakim Nivre, Marie-Catherine de Marneffe, Filip Ginter, Yoav Goldberg, Jan Hajič, Christopher D. Manning, Ryan McDonald, Slav Petrov, Sampo Pyysalo, Natalia Silveira, Reut Tsarfaty, and Daniel Zeman. 2016. **Universal Dependencies v1: A multilingual treebank collection**. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, pages 1659–1666, Portorož, Slovenia. European Language Resources Association (ELRA).
- Peng Qi, Yuhao Zhang, Yuhui Zhang, Jason Bolton, and Christopher D. Manning. 2020. **Stanza: A python natural language processing toolkit for many human languages**. In *Proceedings of the 58th Annual Meeting of the Association for Computational*

*Linguistics: System Demonstrations*, pages 101–108, Online. Association for Computational Linguistics.

Siva Reddy, Oscar Täckström, Michael Collins, Tom Kwiatkowski, Dipanjan Das, Mark Steedman, and Mirella Lapata. 2016. [Transforming dependency structures to logical forms for semantic parsing](#). *Transactions of the Association for Computational Linguistics*, 4:127–140.

Alexis Ross and Ellie Pavlick. 2019. [How well do NLI models capture verb veridicality?](#) In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2230–2240, Hong Kong, China. Association for Computational Linguistics.

V. Sanchez. 1991. Studies on natural logic and categorical grammar.

Natalia Silveira, Timothy Dozat, Marie-Catherine de Marneffe, Samuel Bowman, Miriam Connor, John Bauer, and Christopher D. Manning. 2014. A gold standard dependency corpus for English. In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC-2014)*.

Amir Zeldes. 2017. [The GUM corpus: Creating multilayer resources in the classroom](#). *Language Resources and Evaluation*, 51(3):581–612.

Daniel Zeman, Jan Hajič, Martin Popel, Martin Potthast, Milan Straka, Filip Ginter, Joakim Nivre, and Slav Petrov. 2018. [CoNLL 2018 shared task: Multilingual parsing from raw text to Universal Dependencies](#). In *Proceedings of the CoNLL 2018 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies*, pages 1–21, Brussels, Belgium. Association for Computational Linguistics.