

Extending Implicit Discourse Relation Recognition to the PDTB-3

Li Liang¹ Zheng Zhao² Bonnie Webber²

¹Dept of Linguistics and English Language, University of Edinburgh

²School of Informatics, University of Edinburgh

L.Liang-7@sms.ed.ac.uk

zheng.zhao@ed.ac.uk bonnie@inf.ed.ac.uk

Abstract

The PDTB-3 contains many more implicit discourse relations than the previous PDTB-2. This is in part because implicit relations have now been annotated *within* sentences as well as *between* them. In addition, some now co-occur with explicit discourse relations, instead of standing on their own. Here we show that while this can complicate the problem of identifying the *location* of implicit discourse relations, it can in turn simplify the problem of identifying their *senses*. We present data to support this claim, as well as methods that can serve as a non-trivial baseline for future state-of-the-art recognizers for implicit discourse relations.

1 Introduction

Most readers will be familiar with the PDTB-2 (Prasad et al., 2008). At the time of its creation, it was the largest public repository of annotated discourse relations (over 43K), including over 18.4K signalled by explicit discourse connectives (coordinating or subordinating conjunctions, or discourse adverbials). In the corpus, discourse relations comprise two arguments labelled *Arg1* and *Arg2*, with each relation anchored by either an explicit discourse connective or adjacency. In the latter case, annotators inserted one or more *implicit connectives* to signal the sense(s) they inferred to hold between the arguments. The size and availability of the PDTB-2 spawned work on *shallow discourse parsing*, as in the 2015 and 2016 CoNLL shared tasks (Xue et al., 2015, 2016).

With the release of the PDTB-3¹, there are now ~12.5K additional intra-sentential relations annotated (i.e., relations that lie wholly within the projection of a top-level S-node) and ~1K additional inter-sentential relations (Webber et al., 2019).

¹<https://catalog.ldc.upenn.edu/LDC2019T05>

Work on *shallow discourse parsing* (including the CoNLL shared tasks, as well as (Bai and Zhao, 2018; Dai and Huang, 2018; Rutherford et al., 2017; Shi and Demberg, 2017)) consistently shows that recognizing and sense labelling implicit discourse relations poses more of a challenge than doing so for explicit discourse relations. Hence, implicit relations are the focus of the current work.

But there is another reason as well: Work on the PDTB-2 has assumed (correctly) that non-explicit discourse relations (i.e., implicit relations, *AltLex relations* (Prasad et al., 2010) and entity relations) only hold between *adjacent sentences* as they did in the PDTB-2, so that a sentence boundary is the only position that needs to be checked for the presence of a non-explicit relation. The difficult problem lay in assigning sense-labels to implicit relations.

In Section 2, we show that, with the PDTB-3, this is no longer the case because non-explicit relations can hold *within* sentences as well as *between* them. This in turn motivates a new approach to handle implicit discourse relations in shallow discourse parsing, involving both finding them as well as identifying their senses (Section 3). After showing that the sense-distribution of implicit relations *within* sentences differs from that *between* them (cf. Section 4), we argue that one should be able to take advantage of this fact in sense-labelling these relations.² Section 5 describes two different ways of doing so, along with a way of dealing with another difference in sense distribution — that of implicit relations that co-occur with explicit relations and implicit relations that do not. While the particular methods used here for sense-labelling may not advance the state-of-the-art, it is the way we use them

²Some previous approaches to discourse parsing have also distinguished relations that occur within a sentence from those that occur across sentences (Joty et al., 2013, 2015), but it was not felt to be needed in the PDTB-2, where implicit relations only appeared across sentences.

that should deliver a new baseline for recognizing a fuller range of implicit relations and contribute to the next generation of shallow discourse parsers.³

2 Discourse Annotation in PDTB-3

Discourse annotation in the PDTB-3 differs from that in the PDTB-2 in two major ways: (1) many more discourse relations are annotated *within* sentences, and (2) there are changes in the sense hierarchy used in annotating them. While only the first requires changes to shallow discourse parsing, presenting changes to the senses used in annotating relations will allow us to show differences in the distribution of senses associated with different types of implicit discourse relations.

2.1 Additional Annotation in PDTB-3

It was a consequence of the way that the PDTB-2 was annotated, that there were over twice as many discourse relations annotated across sentences than within them. The former were either explicit relations associated with discourse adverbials or sentence-initial coordinating conjunctions⁴, or implicit relations between paragraph-internal adjacent sentences not otherwise linked by a discourse connective. Within sentences, only annotated were explicit relations associated with subordinating conjunctions, sentence-internal coordinating conjunctions, and discourse adverbials (both of whose arguments were in the same sentence). So it should not be surprising that there were many more inter-sentential relations than intra-sentential relations in the PDTB-2.

In contrast, of the over 13K additional discourse relations annotated in the PDTB-3, over 95% of them occur *within* individual sentences. Of the new relations, 5780 are implicit, some standing alone (like the implicit relations between sentences), with others co-occurring with an explicit discourse relation. Within a sentence, implicit relations occur at the boundaries of syntactic forms — for example, at the boundary between a *free adjunct* and its matrix clause (Ex. 1), or at the boundary between a *to-clause* and its matrix clause (Ex. 2), or between two punctuation-marked conjuncts (Ex. 3).

³It would not make sense to have separate processors for explicit discourse relations, as the decision process takes account of the discourse connective, thereby already learning whether the arguments are likely to occur across vs. within sentences.

⁴Despite what people may have been taught, there are over 2100 tokens of sentence-initial "But" in the Penn *WSJ* corpus and over 660 tokens of sentence-initial "And".

- (1) *Treasury bonds got off to a strong start, **advancing modestly during overnight trading on foreign markets.*** Conn=*specifically* (ARG2-AS-DETAIL) [wsj_0351]
- (2) *After a bad start, Treasury bonds were buoyed by a late burst of buying, **to end modestly higher.*** Conn=*therefore* (RESULT) [wsj_0400]
- (3) Father McKenna moves through the house *praying in Latin, **urging the demon to split.*** (CONJUNCTION) [wsj_0413]

Because implicit relations within sentences don't all occur at a single, well-defined position, this adds to the problems of shallow discourse parsing.

In addition to stand-alone implicits in the PDTB-3, annotators were allowed to indicate implicit relations that co-occur with explicit relations (Rohde et al., 2017, 2018), as a way of indicating a relation that did not derive from the explicit connective, but rather from what the annotator inferred from the arguments themselves, as in Ex. 4–6:

- (4) We've got to *get out of the Detroit mentality* **and Implicit=instead be part of the world mentality,** declares Charles M. Jordan, GM's vice president for design . . . [wsj_0956]
(EXPANSION.CONJUNCTION,
EXPANSION.SUBSTITUTION.ARG2-AS-SUBST)
- (5) . . . Exxon Corp. *built the plant* **but (Implicit=then) closed it in 1985.** [wsj_1748]
(COMPARISON.CONCESSION.ARG2-AS-DENIER,
TEMPORAL.ASYNCHRONOUS.PRECEDENCE)
- (6) . . . which [i.e., the line item veto] would enable him to *kill individual items in a big spending bill* **without (Implicit=however) having to kill the entire bill.** [wsj_1133]
(EXPANSION.MANNER.ARG2-AS-MANNER,
COMPARISON.CONCESSION.ARG2-AS-DENIER)

In Ex. 4, the annotators indicated that they inferred ARG2-AS-SUBST from the pair of arguments conjoined with *and*. The annotators took *and* itself to convey only that its arguments played the same role with respect to the prior text. It is the arguments themselves that led them to conclude that the second conjunct is meant to substitute for the first.

Similarly, in Ex. 5, the annotators indicated that they inferred the temporal relation PRECEDENCE from the pair of arguments conjoined with *but*. The annotators took *but* itself to convey CONCESSION. It is the arguments themselves that led the annotators to conclude that the second conjunct follows the first in time.

Finally, in Ex. 6, the annotators indicated that they inferred a CONCESSION relation from the pair of arguments linked by *without*. The annotators took *without* itself (like its positive version *with*) to

convey MANNER. It is only the arguments that led them to conclude that Arg2 denies an expectation raised by Arg1.

In the PDTB-3, when two relations co-occur, they are explicitly **linked** through a shared index. The consequence for shallow discourse parsing is that explicit relations now need to be checked for co-occurrence with an implicit relation.

2.2 Changes to the Sense Hierarchy

The sense hierarchy used in annotating the PDTB-3 differs from that used in annotating the PDTB-2 in three ways:

1. Rare and/or difficult to annotate senses were dropped, as with the different types of conditional senses;
 2. Sense relations at Level-3 now only encode *directionality* — for example, distinguishing ARG1-AS-SUBST (Ex. 7) from ARG2-AS-SUBST (Ex. 8)
 3. New senses were added that were found to be needed for annotating relations within sentences.
- (7) ARG1-AS-SUBST: *instead of featuring a major East Coast team against a West Coast team, it pitted the Los Angeles Dodgers against the losing Oakland A's* [wsj_0443]
- (8) He *could develop the beach through a trust, but instead is trying have his grandson become a naturalized Mexican so his family gains direct control.* [wsj_0300]

More about the senses used in annotating the PDTB-3 can be found in Webber et al. (2019). Senses are relevant to this discussion of implicit relations in shallow discourse parsing because (as set out in Section 4) implicit relations have been found to have different sense distributions depending on where they occur.

2.3 Stand-off annotation in the PDTB-3

Both the PDTB-2 and PDTB-3 use stand-off annotation. What is relevant with respect to the experiments we report here, is what information is explicit in the annotation, as opposed to having to be computed. This information includes (1) the type of the relation (Explicit, Implicit, AltLex, AltLexC, Entity, Hypophora, NoRel); (2) the byte spans of the two arguments of the relation; and (3) the explicit index (aka *link*) of relations that co-occur by virtue of sharing the same or nearly

the same arguments. The full field structure of discourse relations is set out in Section 8 of Webber et al. (2019). What has to be recovered from the argument spans and the span of the projection of the top node in each sentence-level parse tree is whether a relation occurs wholly within a single sentence or involves multiple sentences.

3 Basic Model Architecture

The sense classifiers for implicit relations used in this paper are based on a Basic Model whose properties reflect consideration of data size and the interaction between lexical information and structural information. (A full description of the Basic Model is given in Appendix A.)

The architecture of Basic Model is shown in Figure 1. It consists of two LSTMs (Hochreiter and Schmidhuber, 1997) and max-pooling layers, a hidden layer, a dense layer, and a softmax layer. Inputs to the model consist of pairs of discourse arguments, each represented as a sequence of word vectors. The output is a probability distribution of the senses between the discourse argument spans. The two sequences of word vectors are encoded by LSTMs in order to capture positional information within the sequential structure. Max-pooling on the output of the LSTMs is used to compose meaning and reduce parameters for the model, as it has been proven effective in Conneau et al. (2017). Modeling the interaction between discourse arguments follows Rutherford and Xue (2016), who argue that discourse relations can only be determined by jointly analyzing the arguments. In addition, Rutherford et al. (2017) observed the influence of different configurations on the performance of the model for the implicit sense classification task, suggesting an interaction between the lexical information in word vectors and the structural information encoded in the model itself. We follow them in adopting a 300-dimension word2vec (Mikolov et al., 2013b) word embedding and hidden size of 100 for the Basic Model.

4 Differences in the distribution of sense relations

To argue for separating the recognition of intra-sentential implicits from inter-sentential implicits, and the recognition of linked implicits from stand-alone implicits, we show how their sense distributions are different.

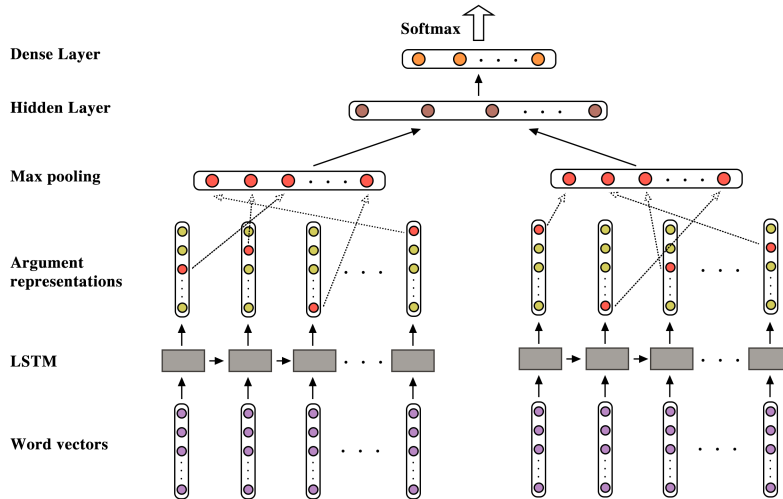


Figure 1: The overall model architecture for implicit sense classification

Table 1 compares the distribution of inter-sentential and intra-sentential implicit relations with respect to the PDTB-3’s Level-2 sense labels, along with the proportion of each label to the total inter-sentential and intra-sentential implicit relations. Besides differences in frequency — for example, relations expressing PURPOSE constitute 21.76% of intra-sentential implicit relations, while only 0.12% of inter-sentential implicits, while relations expressing INSTANTIATION constitute 8.89% of inter-sentential implicits, while only 1.4% of intra-sentential implicits — the senses of inter-sentential implicits are more unequally distributed. That is, three senses — CONTINGENCY.CAUSE, EXPANSION.CONJUNCTION and LEVEL-OF-DETAIL cover 67.08% of the inter-sentential implicits. In contrast, except for CONTINGENCY.CAUSE and PURPOSE, most of the other intra-sentential implicits are more evenly distributed. As often happens with training on an imbalanced distribution, the unequal distribution of inter-sentential relations can lead the model to predict the majority class, ignoring minority classes.

As for the 1753 implicits that co-occur with explicit relations, Table 2 shows that their sense distribution differs sharply from that of stand-alone implicit relations. For example, over 70% convey either CAUSE or ASYNCHRONOUS, while this holds of only 28.7% of stand-alone implicit relations. As such, linked implicits should be more predictable than stand-alone implicit relations.

5 Inter- and intra-sentential Implicits

Differences in the distribution of implicit relations *within* sentences and *across* sentences suggest that we exploit this difference in sense-labelling implicit relations. In this section, we first assume that we know where implicit relations are located within a sentence, so that we can simply consider their arguments. We then present work we have done towards relaxing this assumption.

Task 1: Consider the location of implicit relations in classification. There are different ways to take the location of implicit relations into consideration. Here we present two models, **Model 1** (Section 5.2) and **Model 2** (Section 5.3), both based on the basic model architecture described in Section 3. We compare them with the **Basic Model**, which uses the same classifier on all tokens. We compare their performance not just using the standard training-development-test split, where the ratio of inter- to intra-sentential implicits in the training set, WSJ section 2-21, is 12787:5014. In addition, we follow [Shi and Demberg \(2017\)](#), who argue that evaluation through cross-validation is more predictive, given the wide variation in texts that appear in different sections of the Penn *Wall Street Journal* corpus. The average ratio of inter- to intra-sentential implicits in training sets of cross-validation is 12747:4992. The scores of 3 models are weighted by the proportion of inter- and intra-sentential tokens in the test set.

Task 2: Identify the location of implicit relations. To reduce the dependency on the gold standard annotations of where implicit discourse re-

		inter-sentential	intra-sentential
Comparison	Concession	1355 (8.70%)	136 (2.19%)
	Concession+SpeechAct	7 (0.04%)	3 (0.05%)
	Contrast	700 (4.50%)	156 (2.51%)
	Similarity	14 (0.09%)	14 (0.23%)
Contingency	Cause	4153 (26.67%)	1613 (25.97%)
	Cause+SpeechAct	21 (0.13%)	1 (0.02%)
	Cause+Belief	105 (0.67%)	94 (1.51%)
	Condition	1 (0.01%)	198 (3.19%)
	Condition+SpeechAct	1 (0.01%)	1 (0.02%)
Expansion	Purpose	19 (0.12%)	1351 (21.76%)
	Conjunction	3648 (23.43%)	733 (11.80%)
	Disjunction	9 (0.06%)	21 (0.34%)
	Equivalence	286 (1.84%)	48 (0.77%)
	Exception	4 (0.03%)	1 (0.02%)
	Instantiation	1385 (8.89%)	87 (1.40%)
	Level-of-detail	2644 (16.98%)	589 (9.48%)
	Manner	4 (0.03%)	223 (3.59%)
Temporal	Substitution	221 (1.42%)	145 (2.33%)
	Asynchronous	647 (4.15%)	608 (9.79%)
	Synchronous	348 (2.23%)	188 (3.03%)
total		15572	6210

Table 1: Distribution of inter-sentential/intra-sentential implicit relations among Level 2 labels and the proportion of each label with respect to inter-sentential/intra-sentential implicit relations

lations hold within sentences, two recognizers to identify implicit relations and find argument spans are provided. The first recognizer (Section 5.4) takes syntactic features to identify sentences that contain intra-sentential relations. The second recognizer (Section 5.5) exploits the properties that some explicit relations are linked with implicit relations, checking the explicit relations for co-occurrence with implicit relations to obtain the shared arguments.

5.1 Basic Model

The Basic Model uses the same classifier on all tokens. Since we know which tokens are inter-sentential and which are intra-sentential, we can compare how well the Basic Model does on each. To compute the F_1 scores for the overall performance of the model, the scores of the model are combined, weighted by the proportion of inter- or intra-sentential tokens in the test set. This is shown on the first line of Table 3, elaborated in the confusion matrix shown in Figure 2. A Chi-squared test on the results show the performance of the Basic Model appears to depend to a statistically significant extent on whether the sense appears inter- or intra-sententially ($p=1.50e-03$).

5.2 Model 1

Model architecture: The idea behind Model 1 is to separate the classification task into intra-sentential and inter-sentential implicit sense clas-

sification, with separate classifiers for each. The model architecture and configuration of each classifier are the same as in the Basic Model (Section 3). We expect each classifier to capture different sense distributions of intra-sentential or inter-sentential implicits.

Training and evaluation: Based on their argument spans and the spans associated with each sentence in a file, tokens can be labeled as inter-sentential or intra-sentential. For the standard training-development-test framework, the tokens are allocated into separate inter-sentential/intra-sentential training, development, and test sets. The inter-sentential training set is used in training the inter-sentential implicit sense classifier, and similarly for intra-sentential classification. Test set tokens labeled as inter-sentential or intra-sentential are fed into the appropriate classifier.

Results: The second line of Table 3 presents F_1 scores for Model 1 evaluated on the main evaluation test set and by cross-validation. It shows that Model 1 improves on the Basic Model in predicting intra-sentential implicit relations. The performance of the model significantly depends on the location of relations ($p = 2.41e-09$). The confusion matrix for Model 1⁵ (cf. Figure 2) shows that labels with a relatively larger sample size in each set are predicted more often, includ-

⁵combining results of the inter-sentential and intra-sentential classifiers

		stand-alone	linked
Comparison	Concession	1401 (6.99%)	90 (5.13%)
	Concession+SpeechAct	10 (0.05%)	0 (0.00%)
	Contrast	795 (3.97%)	61 (3.48%)
	Similarity	18 (0.09%)	10 (0.57%)
Contingency	Cause	4943 (24.68%)	823 (46.95%)
	Cause+SpeechAct	22 (0.11%)	0 (0.00%)
	Cause+Belief	164 (0.82%)	35 (2.00%)
	Condition	199 (0.99%)	0 (0.00%)
	Condition+SpeechAct	2 (0.01%)	0 (0.00%)
Expansion	Purpose	1367 (6.83%)	3 (0.17%)
	Conjunction	4360 (21.77%)	21 (1.20%)
	Disjunction	30 (0.15%)	0 (0.00%)
	Equivalence	326 (16.28%)	8 (0.46%)
	Exception	4 (0.02%)	1 (0.06%)
	Instantiation	1456 (7.27%)	16 (0.91%)
	Level-of-detail	3172 (15.84%)	61 (3.48%)
	Manner	173 (0.86%)	54 (3.08%)
Temporal	Substitution	276 (1.38%)	90 (5.13%)
	Asynchronous	800 (3.99%)	455 (25.96%)
	Synchronous	511 (2.55%)	25 (1.43%)
total		20029	1753

Table 2: Distribution of linked and stand-alone implicit relations among Level 2 labels and the proportion of each label with respect to the total linked/stand-alone implicit relations

	Main evaluation metric			Cross validation
	inter-sentential	intra-sentential	overall	
Basic model	35.791	47.154	38.608	41.463
Model 1	34.973	56.666	40.222	43.418
Model 2	37.701	50.410	40.827	42.174

Table 3: F_1 scores of the different models on inter-sentential and intra-sentential implicit relation at Level 2.

ing CONTINGENCY.PURPOSE (frequent in intra-sentential implicits), EXPANSION.CONJUNCTION (frequent in inter-sentential implicits) and CONTINGENCY.CAUSE (frequent in both). The confusion matrix also shows that less frequent senses are confused with these frequent labels more often. Model 1 also reduces the ignorance problem of the Basic Model, in that it correctly classifies some samples into TEMPORAL.SYNCHRONOUS, which is a label ignored by the basic model.

5.3 Model 2

Model architecture: Model 2 treats being inter-sentential or intra-sentential as a single binary feature. Model 2 is created by modifying the Basic Model to include this feature after obtaining the combined representations of the two arguments. We concatenate the binary feature f_S with the output of the dense layer before applying the softmax function, expecting it to affect the final prediction.

Training and evaluation: The data selection follows the standard and cross-validation data split process. The evaluation assumes that each token in the test set has been given an inter-sentential or

intra-sentential feature. The scores are computed following the general process as the basic model.

Results: The third line of Table 3 shows that Model 2 improves over the Basic Model with respect to both inter- and intra-sentential implicit sense prediction, though the performance of the model still has a statistically significant dependence on the location of relations ($p = 4.53e-04$). The improvement of Model 2 on intra-sentential labels is not as dramatic as Model 1. Compared to the previous model, Model 2 doesn't sharpen its focus on those frequent labels in inter- or intra-sentential sets. Instead, the integrated feature in the representations distributes the benefits on the prediction ability of different labels more evenly. In addition, the confusion matrix in Figure 2 shows that Model 2 reduces the confusion between INSTANTIATION and LEVEL-OF-DETAIL, which Scholman and Demberg (2017) have highlighted as a common source of confusion. The confusion matrix for Model 2 also shows some attention to less frequent labels such as COMPARISON.CONTRAST, which are not predicted in either the Basic Model or Model 1.

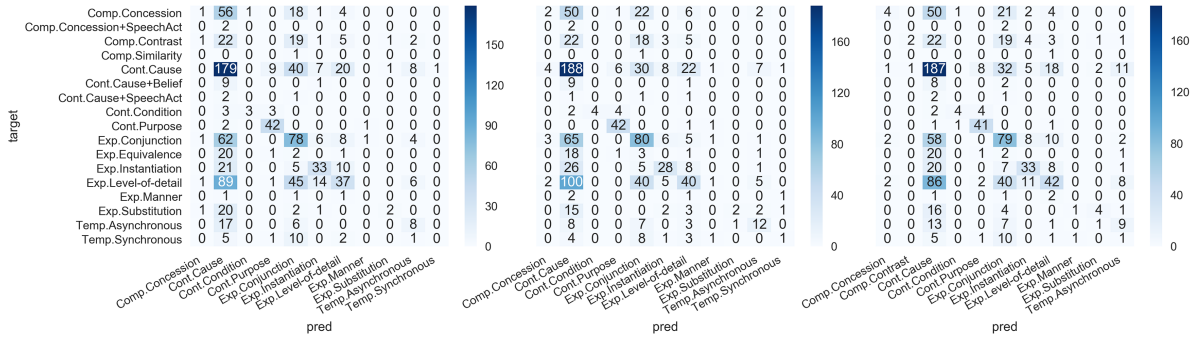


Figure 2: Confusion matrix of the Basic Model, Model 1 and Model 2

5.4 Towards finding implicits within sentences

The results presented above reflect “gold knowledge” of where implicit discourse relations hold within sentences. But in truth, their locations need to be identified before (or jointly with) labelling their senses. We have viewed this as a two-step process: Recognizing sentences that contain at least one implicit intra-sentential relation, and then recognizing the arguments to each relation. The first step has been implemented using a recognizer that takes a linearized parse tree of a sentences as the input. The second step is future work.

Model architecture: Similar to the Basic Model, inputs are represented as a sequence of word vectors, and word embeddings are initialized using pre-trained fastText (Bojanowski et al., 2017) vectors (16B tokens). These vectors are fed to a BiLSTM whose outputs are then fed to a linear layer to produce a binary label, indicating the existence of at least one implicit intra-sentential relation. Word embeddings are set to 200, hidden dimensions, to 256, and vocabulary size, to 25k.

Training and evaluation: To train our recognizer, we first created a dataset of triplets comprising a sentence from PDTB-3, its corresponding parse tree, and a binary label. We obtain the parse trees from the Penn TreeBank (PTB – Marcus et al. 1993) and set the binary label to 1 if there exist at least one implicit or AltLex relation in that sentence. For example, the sentence in Ex. 9 is labelled 1, while that in Ex. 10 is labelled 0.

- (9) MARKET MOVES, these managers don’t.
 ((S-HLN (S (NP-SBJ (NN MARKET)) (VP (VBZ MOVES))) (, ,) (S (NP-SBJ (DT these) (NNS managers)) (VP (VBP do) (RB n’t) (VP (-NONE-*?*))))))) [wsj.1825]

- (10) Oil-tool prices are even edging up.
 ((S (NP-SBJ (NN Oil-tool) (NNS prices)) (VP (VBP are) (ADVP (RB even)) (VP (VBG edging) (ADVP-DIR (RP up)))))) [wsj.0725]

Intra-sentential AltLex relations are included here because they are simply Implicit relations whose alternative lexicalization reliably signals its sense — for example, the phrases “resulting in”, “avoiding”, and “contributing to” are all taken to be alternative lexicalizations that reliably signal RESULT. This is not true of the earlier Examples 1–3, which are classed as Implicits. On the other hand, we do not label “linked” implicit relations as 1 because the visible evidence is an explicit connective signalling an explicit relation, and we don’t want that to be taken *per se* as evidence for an implicit relation. For recognizing linked implicits, we have built a separate model which will be discussed in Section 5.5.

Our training used the Adam optimizer (Kingma and Ba, 2015) with a learning rate of 1e-4. We randomly split the dataset into training (60%), development (20%) and test (20%). To understand what happens if “gold parse trees” are not used, we also created variants of the dataset using parse trees from the widely used Berkeley parser (Kitaev and Klein, 2018) and Stanford parser (Manning et al., 2014).

Results: As the dataset is heavily imbalanced, we also added a simple baseline which predicts the most frequent label. Test set results of the recognizer on the three datasets are presented in Table 4. Even though the baseline achieved an accuracy of ~0.9, it doesn’t convey any useful information, as it labels all instances as 0. We can observe that the model with gold Penn TreeBank parse trees obtain the best performance, followed by the Berkeley parser. Stanford parse trees result in worst perfor-

Parse trees	Accuracy	Precision	Recall	F_1
Baseline	0.9028	0	0	0
Gold	0.9617	0.7799	0.8968	0.8343
Berkeley	0.9473	0.7814	0.6334	0.6997
Stanford	0.9349	0.7153	0.5537	0.6242

Table 4: Results on task of identifying sentences that contain at least one intra-sentential relation, comparing gold parse trees from the PTB with the parse trees output by the Berkeley parser and by the Stanford parser. Baseline refers to the model that predicts the most frequent label.

mance. Examining these trees led us to conclude that, while the Stanford parser does well for basic syntactic structures, which are the most common, it has trouble with challenging structures such as those associated with conjunction. An example is provided in Ex. 11. Here, “steps” has been incorrectly labelled NNS, when it is actually a VBZ, heading the second conjunct. If there were only two conjuncts, explicitly conjoined with “and”, the sentence would not contain an implicit relation. With three conjuncts, however, the first two would normally be *comma-conjoined*, with the discourse relation between them taken to be implicit. But the error in PoS-tagging has eliminated evidence of a second conjunct, with an implicit discourse relation to the first conjunct. Errors in PoS-tagging and mis-parsing associated with rare constructions, means that the accuracy is lower than that of the Berkeley parser. However, as Precision, Recall, and F_1 are measured for 1 labels, these metrics are more adversely affected when compared to those of the Berkeley parser.

- (11) With three minutes left on the clock, Mr. Aikman takes the snap, steps back and fires a 21-yard pass – straight into the hands of an Atlanta defensive back.

IN CD NNS VBD IN DT NN , NNP NNP VBZ DT NN , NNS RB CC VBZ DT JJ NN : RB IN DT NNS IN DT NNP NN RB .

((S (SBAR (IN With) (S (NP (CD three) (NNS minutes)) (VP (VBD left) (PP (IN on) (NP (DT the) (NN clock)))))) (, .) (NP (NNP Mr.) (NNP Aikman)) (VP (VP (VBZ takes) (NP (NP (DT the) (NN snap)) (, .) (NP (NNS steps))) (ADVP (RB back))) (CC and) (VP (VBZ fires) (NP (DT a) (JJ 21-yard) (NN pass))) (: -) (PP (RB straight) (IN into) (NP (NP (DT the) (NNS hands)) (PP (IN of) (NP (DT an) (NNP Atlanta) (NN defensive)))))) (ADVP (RB back)) (, .))) [wsj_1411]

5.5 Recognizing “linked” implicit relations

As noted in Section 2.1, implicit relations can co-occur with explicit relations. While the location of such implicits is not identified by the recognizer

	Precision	Recall	F_1	Proportion
stand-alone	0.951	0.905	0.928	93.67%
linked	0.193	0.329	0.243	6.33%

Table 5: Precision, Recall and F_1 scores of linked/stand-alone labels predicted by the recognizer using main evaluation metrics and their proportion in test data.

described in Section 5.4, we actually know the location of their arguments, because co-occurring (aka “linked”) relations share their argument spans. Hence, recognizing explicit relations linked with implicit ones means that we also obtain argument spans of these implicits. Here we describe a first attempt to automatically discriminate explicit relations linked with implicit relations from ones that are not so linked. It comprises two steps: extracting sentences that contain explicit relations as our datasets, and then recognizing the ones linked with implicit relations.

Model architecture: To detect linked implicit relations from explicit relations, we use a naive Bayes classifier — specifically, the one provided in NLTK (Bird and Loper, 2004). Production rules are selected as input feature as it has been proven notably effective in feature-based implicit discourse relation recognition task among different features (Park and Cardie, 2012). Models trained in Task 1 will be adopted for linked sense classification.

Training and evaluation: We follow the standard split to select the training and test set. Each token in the training set consists of *Arg1*, connective and *Arg2*, and are parsed to extract syntactic productions used in parent-child nodes in the argument parse trees. The 100 most-frequent production rules are used to build a feature dictionary for input. A production rule feature is labeled as 1 in the dictionary if it appears in the parse tree of the token, otherwise it will be 0. The linked/stand-alone label is determined by whether the explicit relation shares the same index value with an implicit relation. The recognizer is evaluated by how well it distinguishes explicit relations that have a linked implicit relation from ones that don’t. Classifiers are evaluated on the recognized implicit relations.

Results: The low Recall for linked relations in Table 5 shows that the recognizer performs better on predicting stand-alone relations, which are a majority of the data. Linked implicits in the test set (WSJ Section 23) are mostly linked to conjoined clauses

or conjoined VPs, and are signaled by implicit connective like “and” (81.08%) or “but” or an adverbial. Most correctly recognized relations are VPs conjoined with “and”. All the recognized linked implicit relations are found intra-sentential. We adopt the intra-sentential classifier in Model 1 and the Basic Model to test the classifier based on the recognized results. The intra-sentential classifier achieves an F_1 score of 75, compared with 68.182 using the Basic Model. This again emphasizes that knowing the location of implicit discourse relation would benefit sense identification.

6 Conclusion and future work

We have shown that recognizing implicit discourse relations as annotated in the PDTB-3 now requires finding them, as well as figuring out what sense relation(s) holds between the arguments. However, we have also shown that the latter task is simplified by differences in the sense distribution of different implicit relations. We still have to develop a way of recognizing precisely where implicit relations hold in those sentences that can be identified as containing them, and a more accurate approach to sense labelling implicit relations that co-occur with explicit ones. We are also interested in whether these different sense distributions hold in other news corpora and other genres. While it is likely not the case that all languages show the same difference in the sense distribution of discourse relations, we would not be surprised if the discourse relations realized within sentences differed from those realized across sentences. In conclusion, we hope that the current effort will contribute to future work on shallow discourse parsing as annotated in the PDTB-3.

Acknowledgments

We would like to thank the anonymous reviewers for their valuable comments. We would also like to thank Annie Louis for her contributions to the work on recognizing the presence of sentence-internal implicit discourse relations.

References

Hongxiao Bai and Hai Zhao. 2018. Deep enhanced representation for implicit discourse relation recognition. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 571–583, Santa Fe, New Mexico, USA. Association for Computational Linguistics.

Steven Bird and Edward Loper. 2004. *NLTK: The natural language toolkit*. In *Proceedings of the ACL Interactive Poster and Demonstration Sessions*, pages 214–217, Barcelona, Spain. Association for Computational Linguistics.

Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2017. *Enriching word vectors with subword information*. *Transactions of the Association for Computational Linguistics*, 5:135–146.

Alexis Conneau, Douwe Kiela, Holger Schwenk, Loïc Barrault, and Antoine Bordes. 2017. Supervised learning of universal sentence representations from natural language inference data. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 670–680, Copenhagen, Denmark.

Zeyu Dai and Ruihong Huang. 2018. Improving implicit discourse relation classification by modeling inter-dependencies of discourse units in a paragraph. 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 141–151, New Orleans, Louisiana. Association for Computational Linguistics.

Sepp Hochreiter and Jürgen Schmidhuber. 1997. *Long short-term memory*. *Neural Comput.*, 9(8):1735–1780.

Shafiq Joty, Giuseppe Carenini, Raymond Ng, and Yashar Mehdad. 2013. *Combining intra- and multi-sentential rhetorical parsing for document-level discourse analysis*. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics*, pages 486–496, Sofia, Bulgaria.

Shafiq Joty, Giuseppe Carenini, and Raymond T. Ng. 2015. *CODRA: A novel discriminative framework for rhetorical analysis*. *Computational Linguistics*, 41(3):385–435.

Diederik P. Kingma and Jimmy Ba. 2015. *Adam: A method for stochastic optimization*. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.

Nikita Kitaev and Dan Klein. 2018. *Constituency parsing with a self-attentive encoder*. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2676–2686, Melbourne, Australia. Association for Computational Linguistics.

Christopher D Manning, Mihai Surdeanu, John Bauer, Jenny Rose 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.

- Mitchell P. Marcus, Beatrice Santorini, and Mary Ann Marcinkiewicz. 1993. *Building a large annotated corpus of English: The Penn Treebank*. *Computational Linguistics*, 19(2):313–330.
- Tomas Mikolov, Kai Chen, G.s Corrado, and Jeffrey Dean. 2013a. Efficient estimation of word representations in vector space. *Proceedings of Workshop at ICLR*, 2013.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013b. Distributed representations of words and phrases and their compositionality. In *Proceedings of the 26th International Conference on Neural Information Processing Systems - Volume 2*, NIPS’13, page 3111–3119, Red Hook, NY, USA. Curran Associates Inc.
- Joonsuk Park and Claire Cardie. 2012. Improving implicit discourse relation recognition through feature set optimization. In *Proceedings of the 13th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 108–112, Seoul, South Korea. Association for Computational Linguistics.
- Rashmi Prasad, Nikhil Dinesh, Alan Lee, Eleni Miltasakaki, Livio Robaldo, Aravind Joshi, and Bonnie Webber. 2008. The penn discourse treebank 2.0. In *Proceedings of the Sixth International Language Resources and Evaluation (LREC’08)*, pages 2961–2968. European Language Resources Association (ELRA).
- Rashmi Prasad, Aravind Joshi, and Bonnie Webber. 2010. Realization of discourse relations by other means: Alternative lexicalizations. In *Proceedings of the 23rd International Conference on Computational Linguistics (COLING)*, Beijing, China.
- Hannah Rohde, Anna Dickinson, Nathan Schneider, Christopher Clark, Annie Louis, and Bonnie Webber. 2017. Exploring substitutability through discourse adverbials and multiple judgments. In *Proceedings, 12th International Conference on Computational Semantics (IWCS 2017)*, Montpellier, France.
- Hannah Rohde, Alexander Johnson, Nathan Schneider, and Bonnie Webber. 2018. Discourse coherence: Concurrent explicit and implicit relations. In *Proceedings of the 56th Annual Meeting of the ACL*.
- Attapol Rutherford, Vera Demberg, and Nianwen Xue. 2017. A systematic study of neural discourse models for implicit discourse relation. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics (EACL 2017)*, pages 281–291.
- Attapol Rutherford and Nianwen Xue. 2016. Robust non-explicit neural discourse parser in English and Chinese. In *Proceedings of the CoNLL-16 shared task*, pages 55–59, Berlin, Germany.
- Merel Scholman and Vera Demberg. 2017. Examples and specifications that prove a point: Identifying elaborative and argumentative discourse relations. *Dialogue & Discourse*, 8:56–83.
- Wei Shi and Vera Demberg. 2017. On the need of cross validation for discourse relation classification. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 150–156, Valencia, Spain. Association for Computational Linguistics.
- Bonnie Webber, Rashmi Prasad, Alan Lee, and Aravind Joshi. 2019. The penn discourse treebank 3.0 annotation manual. <https://catalog.ldc.upenn.edu/docs/LDC2019T05/PDTB3-Annotation-Manual.pdf>.
- Nianwen Xue, Hwee Tou Ng, Sameer Pradhan, Rashmi Prasad, Christopher Bryant, and Attapol Rutherford. 2015. The CoNLL-2015 shared task on shallow discourse parsing. In *Proceedings of the Nineteenth Conference on Computational Natural Language Learning - Shared Task*, pages 1–16, Beijing, China. Association for Computational Linguistics.
- Nianwen Xue, Hwee Tou Ng, Sameer Pradhan, Attapol Rutherford, Bonnie Webber, Chuan Wang, and Hongmin Wang. 2016. CoNLL 2016 shared task on multilingual shallow discourse parsing. In *Proceedings of the CoNLL-16 shared task*, pages 1–19, Berlin, Germany. Association for Computational Linguistics.

A Specifics of the Basic Model

Here we describe the basic model architecture for implicit relation sense classification in PDTB-3. The configuration for the model is chosen based on consideration of data size and the interaction between lexical information and structural information. A further analysis on the predictive performance of the basic model on each labels is provided as well.

A.1 Model architecture

Figure 1 (repeated here as Figure 3) illustrates the overall model architecture of the neural implicit sense classifier that consists of two LSTM and max-pooling layers, a hidden layer, a dense layer, and a softmax layer. The input for the model is the discourse argument pairs with additional labels⁶, and the output is a probability distribution of the senses between the discourse argument spans.

Word vectors: In our model, arguments *Arg1* and *Arg2* are viewed as two sequences of word vectors with length of n_1 and n_2 . Word vectors for the word in arguments are taken from word embeddings.

$$Arg1 : [x_1^1, x_2^1, \dots, x_{n_1}^1] \quad (1)$$

$$Arg2 : [x_1^2, x_2^2, \dots, x_{n_2}^2] \quad (2)$$

Argument representations: The two sequences of word vectors are encoded by LSTM respectively. The hidden states H_{Arg1} and H_{Arg2} of LSTM are taken. The max-pooling function is employed to compose meaning in the hidden states and reduce parameters for the model, as it has been proven effective in (Conneau et al., 2017). As shown in eq. 6, it will select the maximum value along the sequence at each dimension of the hidden states. $a_j^1(a_j^2)$ represents a maximum value from all the values in a sequence with length of $n_1(n_2)$ at dimension j of the hidden states H_{Arg1} (H_{Arg2}). By concatenating the output of max-pooling function, we have abstract representations A_{Arg1} and A_{Arg2} of arguments *Arg1* and *Arg2* individually.

$$H_{Arg1} = [h_1^1, h_2^1, \dots, h_{n_1}^1] \quad (3)$$

$$H_{Arg2} = [h_1^2, h_2^2, \dots, h_{n_2}^2] \quad (4)$$

⁶These labels are not used in the basic model described in this work, but serve for statistical tests and further experiments.

$$a_j^1 = \max_{k \in n_1}(H_{Arg2}_{j_k}) \quad (5)$$

$$a_j^2 = \max_{k \in n_2}(H_{Arg1}_{j_k}) \quad (6)$$

$$A_{Arg1} = [a_1^1, a_2^1, \dots, a_{hidden_size}^1] \quad (7)$$

$$A_{Arg2} = [a_1^2, a_2^2, \dots, a_{hidden_size}^2] \quad (8)$$

Inter-argument interaction modeling: The modeling of the interaction between two discourse argument representations follows (Rutherford and Xue, 2016), which argues that discourse relations can only be determined by jointly analyzing the arguments. In our model, argument representations A_{Arg1} and A_{Arg2} are weighted by W_1 and W_2 separately. The combination of the weighted argument representations is then transformed non-linearly with *tanh* function in the first hidden layer H_{hid} . It is then fed into a dense layer H_{dense} ⁷. Finally, we predict the discourse relation sense using a softmax function.

$$H_{hid} = \tanh(W_1 \cdot A_{Arg1} + W_2 \cdot A_{Arg2} + b_{hid}) \quad (9)$$

$$H_{dense} = \tanh(W_{dense} \cdot H_1 + b_{dense}) \quad (10)$$

$$output = \text{softmax}(W_{output} \cdot H_{dense} + b_{output}) \quad (11)$$

A.2 Configuration

Implementation: The model is implemented with PyTorch. The cost function is the standard cross-entropy loss function and Adam optimizer with an initial learning rate of 0.001 and a batch size of 32. We determine convergence if the performance of the model on the development set does not improve after more than 3 epochs.

One problem that challenges the training of the model is the limitation on the size of the data. We introduce other resources to overcome it and adopt different techniques to avoid overfitting. Word vectors are directly taken from Word2vec embeddings (Mikolov et al., 2013a) trained with the skip-gram algorithm on Brown corpus, and are fixed during training. To avoid overfitting, we apply a 0.25 dropout ratio to the input of the LSTM layer. Batch normalization is added to normalize the activation between the hidden layer and the dense layer to accelerate the training speed and further prevent overfitting with regularization.

⁷The default size of the dense layer is $hidden_size/5$.

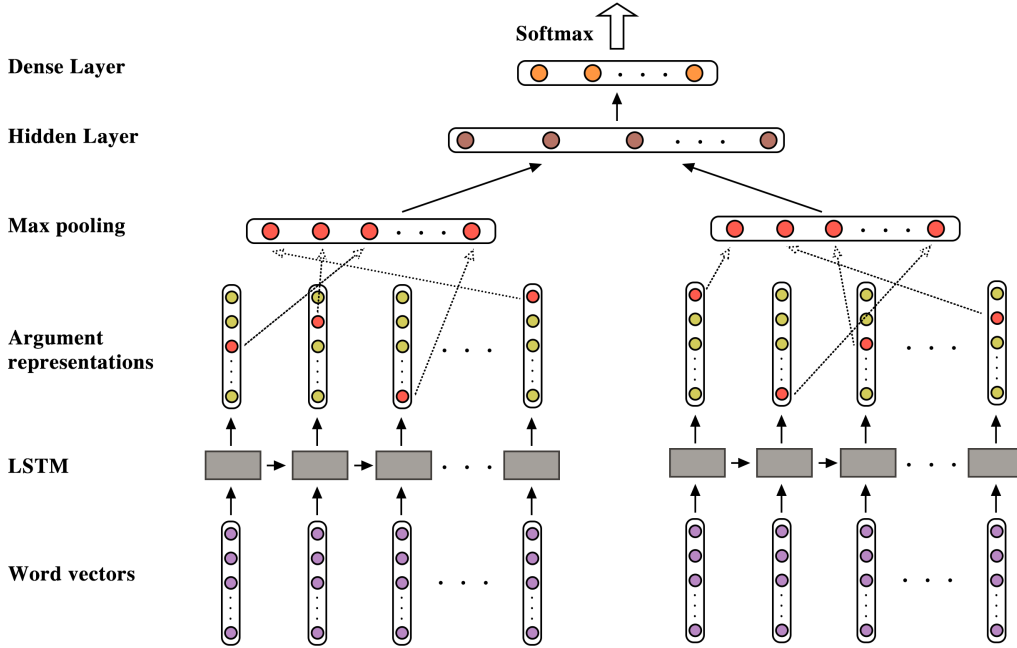


Figure 3: The overall model architecture for implicit sense classification

Hyperparameter Settings: (Rutherford et al., 2017) observed the influence of different configurations on the performance of the model for the implicit sense classification task, suggesting an interaction between the lexical information in word vectors and the structural information encoded in the model itself. To determine the configuration for our model, we trained our model with different combinations of the dimension of word embedding (50, 300) and hidden size (50, 100), and evaluate it on Level 2 labels on the WSJ section 23. Table 6 presents the performance of the model with different configurations. The baseline is Most Frequent Sense heuristic, using the most frequent sense CONTINGENCY.CAUSE in the training data for each target. Our result is in line with their finding of sequential LSTM model, showing larger hidden size 100 is effective when it is accompanied with 300-dimension word embedding. Based on the performance on Level 2 labels, we choose 300-dimension Word2vec word embedding and hidden size 100 as our configuration for the Basic Model.

Our model scores 34.778 at Level 3 (31-way classification). Using cross-validation, our model obtains 41.463 at Level 2.

A.3 Discussion

It is worth examining the performance of the model on each Level 2 label individually. Table 7 displays the precision, recall and F_1 scores of each label

along with its proportion in the test data.

The classifier obtains relatively higher scores on some types of labels. The first type is senses with larger sample size in the corpus, suggesting the imbalanced classification problem. Two senses occur frequently in the corpus (CONTINGENCY.CAUSE and EXPANSION.CONJUNCTION) are recognized with high Recall, but low Precision. This could indicate a strong signal, but one that is likely to be ambiguous. Other less frequent labels are constantly misclassified into these frequent labels. For example, the amount of EXPANSION.MANNER samples is largely reduced by our method dealing with multi-label instances, and the classifier fails to recognize the minority class. Another type of senses achieving high scores are those occurring predominantly in intra-sentential relations (CONTINGENCY.PURPOSE and CONTINGENCY.CONDITION) or in inter-sentential relations (EXPANSION.INSTANTIATION and EXPANSION.LEVEL-OF-DETAIL). The model recognize these senses with high Precision, but different levels of Recall, which could be due to a difference in the strength of evidence signalling the relation. Additionally, TEMPORAL.ASYNCHRONOUS sense that associates with much higher proportion in linked relations than stand-alone ones obtain similar Recall and Precision scores.

	embedding size	hidden size	F_1
Our model	50	50	36.492
	50	100	37.097
	300	50	37.601
	300	100	38.608
Baseline	-	-	28.024

Table 6: F_1 scores of the model with different configurations on predicting Level 2 sense labels

	Precision	Recall	F_1	Proportion
Comparison.Concession	20.000	1.235	2.326	8.17%
Comparison.Concession+SpeechAct	0.000	0.000	0.000	0.20%
Comparison.Contrast	0.000	0.000	0.000	5.14%
Comparison.Similarity	0.000	0.000	0.000	0.10%
Contingency.Cause	35.098	67.547	46.194	26.71%
Contingency.Cause+Belief	0.000	0.000	0.000	1.01%
Contingency.Cause+SpeechAct	0.000	0.000	0.000	0.30%
Contingency.Condition	75.000	33.333	46.154	0.91%
Contingency.Purpose	73.684	93.333	82.353	4.54%
Expansion.Conjunction	34.211	48.750	40.206	16.13%
Expansion.Equivalence	0.000	0.000	0.000	2.42%
Expansion.Instantiation	51.562	47.826	49.624	6.96%
Expansion.Level-of-detail	42.045	19.171	26.335	19.46%
Expansion.Manner	0.000	0.000	0.000	0.30%
Expansion.Substitution	50.000	7.692	13.333	2.62%
Temporal.Asynchronous	27.586	25.806	26.667	3.12%
Temporal.Synchronous	0.000	0.000	0.000	1.92%

Table 7: Precision, Recall and F_1 scores of different labels predicted by the basic model using main evaluation metric and their proportions in test data