

NegPar: a parallel corpus annotated for negation

Qianchu Liu¹ Federico Fancellu² Bonnie Webber²

¹ LTL, University of Cambridge

² ILCC, School of Informatics, University of Edinburgh

q1261@cam.ac.uk f.fancellu@sms.ed.ac.uk bonnie@inf.ed.ac.uk

Abstract

Although the existence of English corpora annotated for negation has allowed for extensive work on monolingual negation detection, little is understood on how negation-related phenomena translate across languages. The current study fills this gap by presenting *NegPar*, the first English-Chinese parallel corpus annotated for negation in the narrative domain (a collection of stories from Conan Doyle’s Sherlock Holmes). While we followed the annotation guidelines in the CONANDOYLE-NEG corpus (Morante and Daelemans, 2012), we reannotated certain scope-related phenomena to ensure more consistent and interpretable semantic representation.

To both ease the annotation process and analyze how similar negation is signaled in the two languages, we experimented with first projecting the annotations from English and then manually correcting the projection output in Chinese. Results show that projecting negation via word-alignment offers limited help to the annotation process, as negation can be rendered in different ways across languages.

Keywords: negation, annotation projection, English, Chinese, parallel corpora, cross-lingual semantics

1. Introduction

Considering negation when extracting information from text is crucial for a range of natural language processing (NLP) tasks, especially when one is interested in whether an action or a state described in a statement is factual or counter-factual.

To date, most work has focused on automatically identifying which tokens in a sentence constitute the three key elements of negation: *cue*, *event* and *scope*. As shown in (1), ‘not’ is marked as the presence of negation, the cue; ‘eat’ is a negated event; and [‘I’, ‘do’, ‘eat’, ‘pizza’] are part of the scope.¹ If not otherwise stated, here and in the remainder of the paper, the cue is marked in **bold** with the scope underlined and the event marked in a box.

(1) I **do not** eat pizza

Corpora annotated for negation have been developed for both English and Chinese and for a variety of domains (§ 2.). However, these are monolingual corpora, and there is little empirical work on how negation is represented across languages.

The current study aims at addressing this shortcoming by presenting the first *parallel* corpus annotated for negation. We created it by leveraging the CONANDOYLE-NEG corpus (Morante and Daelemans, 2012), a collection of four Sherlock Holmes stories. Before extending the English annotations onto the sentence-aligned Chinese translation, we edited the English annotations to better capture phenomena related to negation scope. (§ 3.2.).

To ease the annotation task and to investigate differences in the way negation is represented in the two languages, we experimented with first automatically projecting the annotations and then manually correcting them (§ 3.3. and § 4.).

¹The negated event is usually considered as part of the scope of negation

We experimented with projecting onto both Chinese characters and words via character or word alignment extracted from a large parallel corpus.

In general, the performance of automatic projection yields relatively low results with a large number of false negatives, hence lower recall. In order to ease the annotation task, more effort is required to find those instances where projection failed to predict mostly because negation is present in the Chinese translation but not in English, than to filter over-predicted instances. .

We hope that the resource and insights provided in this study will foster work in a variety of NLP tasks such as machine translation, information extraction, etc.

2. Related Work

2.1. Corpora annotated for negation

The importance of detecting negation is testified by the different domains corpora were annotated for. BIOSCOPE (Vincze et al., 2008), a collection of medical papers, abstract and clinical reports, stresses the importance of recognizing negation for information extraction from medical records; the SFU PRODUCT REVIEW CORPUS (Konstantinova et al., 2012) annotates negation on top of product reviews, acknowledging its importance for sentiment analysis tasks. Finally, CONANDOYLE-NEG (Morante and Daelemans, 2012) annotates negation in narrative texts – a collection of four stories from Conan Doyle’s Sherlock Holmes. There have also been some attempts in developing corpora annotated for negation in other languages as demonstrated by CNE SP (Chinese Negation and Speculation corpus) (Zou et al., 2015), which closely follows the annotation style of the BIOSCOPE corpus.

However, tailoring the annotation style to a specific domain leads these corpora to differ in what was annotated and how. For instance, the SFU and BIOSCOPE corpora consider negation scope in purely syntactic terms (that we infer

is ‘the maximum constituent c-commanded by the negation cue’), including only the tokens to the right of the cue and excluding the subject of the clause except in passive constructions. (2) exemplifies this.

- (2) a. BIOSCOPE: It helps activation, **not inhibition of the cell**.
 b. SFU: This book **wasn’t** published before the year 2000

The CNESP corpus follows a similar notion of scope but also includes the subject, unless this creates a discontinuity in the scope; the subject ‘it’ in (2.a) would therefore *not* be annotated because it would break the scope into two discontinuous spans.

CONANDOLYE-NEG considers a more semantic notion of scope which can be a discontinuous span of text and is therefore able to capture long range dependencies that other corpora do not take into account. Unlike other corpora, CONANDOLYE-NEG also annotates affixal negation (e.g. ‘**impatient**’) as well as the negated event, which allows for better coverage of the negated elements in a sentence.

2.2. Annotation projection

To our knowledge, there has not been any previous work on projecting negation across languages. However previous studies have experimented with projecting semantic annotations via word-alignment information extracted from large parallel corpora. Hwa et al. (2002) used word alignment to project parses from English to Chinese and later improved the performance by implementing a set of linguistically-informed post-processing rules. Padó and Lapata (2009) used word alignment information in their constituent-based projection algorithm to transfer semantic role labels from English to German.

3. The Annotation Task

3.1. Creating the corpus

We built our annotated parallel corpus by aligning the four stories in CONANDOLYE-NEG (‘The Hound of the Baskervilles’, ‘The Adventure of Wisteria Lodge’, ‘The Adventure of the Cardboard Box’ and ‘The Adventure of the Red Circle’) to their Chinese translation by Mengyuan Lin.²

We annotate negation as a tuple of three elements: the **cue**, which is the word (e.g. ‘not’), multiword unit (e.g. ‘neither...nor’) or morpheme (e.g. ‘im’-patient) that are inherently expressing negation; the **scope**, which is the sentence span affected by negation; and the negated **event**, which is the part of the scope that is directly or most prominently negated. The current annotations closely follow the same CONLL format of CONANDOLYE-NEG, marking each negation instance in a sentence as a set of three columns, for cue, event and scope respectively.

The annotation and the guidelines for Chinese were created by one of the current authors, who is a native Mandarin speaker with linguistics background. Issues arising during

	English	Chinese
# sents	5520	5005
# neg. sents	1227(22.22%)	1442(28.81%)
# cues	1421	1782
# events	911	1168
# scopes	1304	1762

Table 1: Comparison between the English and the Chinese text of NEGPAR in terms of total number of sentences (# sents.), negated sentences (# neg.sents.), number of cues (#cues), events (#events) and negated scope spans (#scopes).

the annotation process were discussed with the other two authors. All authors collaborated equally in reannotating the English side.

The annotations were carried out on the basis of the projection from CONANDOLYE-NEG. When a negation instance in English is translated as such into Chinese, we correct the projection if the projection is wrong. If a negation instance in English is not translated as negation in Chinese, the extra negation projection is deleted from the Chinese corpus. Finally, we add those annotations for instances of negation in Chinese that are not present in the corresponding English text. The annotation process took approximately one month to complete.

The corpus with the full annotation guidelines can be found at <https://github.com/qianchu/NegPar>.

A quantitative comparison between the English and the Chinese sides of our corpus is presented in Table 1. Overall, there are more negation instances in the Chinese translation than in English as the percentage of negated sentences, the number of cues, scopes and events in Chinese are consistently higher than in English.

3.2. Reannotating the English side

Before projecting the CONANDOLYE-NEG onto Chinese, we reread the original guidelines of Morante and Daelemans (2012) to verify how phenomena related to negation scope are handled.

We divide this discussion in two parts, the first addressing the phenomena we think need improvement and the second addressing those that were not considered by the original guidelines.

3.2.1. Alternative annotation choices

Morphological negation. Some cases of morphological negation (as in ‘**unhappy**’, ‘**impatient**’, etc.) do not have scope in the traditional semantic sense of the word. The statements ‘I am happy’ and ‘I am unhappy’ can in fact be false at the same time (‘I am neither happy nor unhappy’ is a valid statement), which is why they are contraries rather than real contradictions. Although keeping them annotated, we do not consider them any further.

However, as a matter of consistency, we reconsider the status of negation affixes in adverbs. Whereas the original English guidelines state that ‘If the negated affix is attached to an adverb that modifies a verb, the negation scopes over the entire clause’ (Morante et al., 2011), we found that in many

²The Chinese translation can be found at <http://www.kanunu8.com/book3/8105/>.

such cases, only the adverb is in the scope. Such inconsistencies are shown in (3).

- (3) a. [...] tossing restlessly from side to side
b. [...] glaring helplessly at the frightful thing which was hunting him down

In a case like (3.b) we felt that the wide scope reading does not lead to a correct interpretation because the event actually takes place and it is just the manner it takes place that is negated. For this reason, in such cases, we just **annotate the adverb as being in the scope of negation**; therefore, (3.b) will be reannotated in the same way as (3.a).

Except/save/no...but. By stating that ‘often exception items function as neutralizers of the polarity of the statement expressed in the main sentence where they occur’, Morante et al. (2011) annotated these as part of the scope, as well as the material they introduce. This is exemplified in (4):

- (4) [...] Marx knew nothing of his customer save that he was a good payer

This annotation style does not reflect the fact that ‘save’ excludes from the set of things negated in the matrix clause, which should be interpreted as positive. (4) can be in fact paraphrased as ‘*It is not the case that (Marx knew nothing about his customer). He knew that he was a good payer*’. This use of ‘save’ contrasts from its use as a negation cue as shown in (5)

- (5) Mr. Sherlock Holmes, who was usually very late in the mornings, save upon those not infrequent occasions when he was up all night, was seated at the breakfast table.

Here, ‘save’ is used to neutralize positive polarity to highlight the set of instances where an event did not take place; (5) implies in fact that ‘he was usually very late but he wasn’t late on those not infrequent occasions when he was up all night’.

For this reason, we distinguish two types of exceptions: **The ‘exception to nothing’, which is positive, as shown in (4), and is excluded from the scope of negation; and the ‘exception to usually’, with a negative meaning as shown in (5), where we include only the ‘except’ phrase in the scope of negation, as the original guidelines already do.**

3.2.2. Additional phenomena

Neg raising. Neg raising, i.e. the phenomenon that a negation in the matrix clause of a sentence is interpreted as negating the complement clause, is not covered by the annotation guidelines. Neg raising is encountered with verbs expressing the speaker’s opinion, such as ‘think’, ‘believe’, ‘want’, ‘seem’, etc.. In cases like (6), the original annotations consider the entire sentence under the scope of negation; however, it is not the thinking that should be negated but the object of the thought.

- (6) I do not think it is likely = I think that it is not likely.

In cases of neg raising we annotate as part of the scope the subordinate only. (6) is therefore reannotated as(7)

- (7) I do **not** think it is likely

Quantifiers. The interaction between quantification and negation scope at a string level is not considered at all in the original CONANDOYLE-NEG guidelines.

Cases where ‘not’ directly precedes lexical items like ‘all’ and ‘every’ are correctly annotated, as in the following.

- (8) Money is not everything.
(= *It is not the case that money is everything*).

However, let us consider the following example annotated according to the original guidelines.

- (9) The fellow might have had other reasons for thinking that all was not well

The original guidelines paraphrased the construction under the scope as ‘It was not the case that the fellow was thinking that all was well’, which is why ‘thinking’ as well is in the scope. Besides the problem with neg raising, from a logical perspective, the universal quantifier should scope over negation and not vice versa, given that $\forall thing(x) \rightarrow \neg \exists e.well(e) \wedge Topic(e, x)$

In cases like (10), we therefore **exclude the lexical item representing universal quantification** from the scope to yield the following annotation.

- (10) The fellow might have had other reasons for thinking all was not well

Modals. The interaction between the scope of negation and modality is another phenomenon the guidelines do not mention. Some cases, as the one shown in (11), are correctly handled: negation correctly scopes over the modal.

- (11) You need not to fear to speak the truth. = *It is not the case that you need to fear to speak the truth.*

We however found two cases of deontic modality where the annotations fail to capture this interaction as shown below.

- (12) You certainly must not go alone \neq *It is not the case that you certainly must go alone.*

Having negation take scope over ‘must’ leads to the incorrect interpretation that the person could go alone.

In cases like (12), we adopt a strategy similar to the one used for quantifiers and **exclude the lexical item representing modality from the scope.**

3.3. Annotating the Chinese side

We include here a brief summary of the annotation guidelines in Chinese, where we report annotation examples subdivided into the three components we considered: cue, scope and event. The full guidelines can be found at the same link presented in §3.1..

3.3.1. Cue

We annotated a total of 45 negation cue *types* in Chinese including adverbs, auxiliary verbs and prefixes. Amongst these, we found 10 *core* negation cues³; these negation cues can stand independently to function as adverbs (similarly to the English ‘not’) but they can also be compounded with either adjectives or adverbs to form multiword cues eg.并不, meaning ‘not’ with a contrast emotion.

The most common cues in Chinese are 不, 没有, 没 and 无 as shown in Figure 1. 不 is an adverb cue equivalent to ‘not’, and it is the most widely used cue in Chinese. 没(有) is a negated auxiliary verb that indicates non-completion, roughly equivalent to English ‘did not/have not’. 无 is the classical form of 没有. In modern Chinese, 无 is also used as a negation prefix that translates the English suffix ‘-less’. e.g. 无线(‘no wire=wireless’).

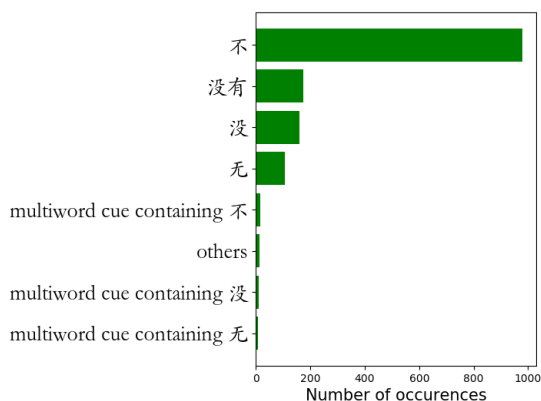


Figure 1: The distribution of Chinese cues in NegPar

Implicit negated verbs. English verbs such as ‘refuse’ and ‘fail’ implicitly express negation and are treated as negation cues in the English guidelines. However, since they are not annotated consistently in the English side, we decide *not* to mark these verbs as cues in Chinese.

Infix cue in verb-complement constructions. In Chinese, the negation cue 不 can appear as an infix in verb-complement construction. The complement usually indicates the result or the direction of an action expressed by the verb, as well as expresses a potential form. According to Li and Thompson (1989), one can interpret infixal negation in Chinese as introducing a result that is ‘unachievable’, roughly equivalent to English ‘cannot’. (13) exemplifies this construction alongside its annotation, where only the infix 不 is considered the cue.

- (13) 他说 得 清楚
He speaks can clear
‘He can speak clearly.’

- 他说 不 清楚
He speaks cannot clear
‘He cannot speak clearly.’

Non-functional negation cue. We do *not* annotate as cues any instances of non-functional negation, i.e. expressions that include markers of negation but have positive meaning.

Certain fixed expressions belong to this category; as shown in (14), the expression ‘can’t help’ and the Chinese counterpart ‘不得不/不能不’, despite including a negation marker, have a positive meaning (i.e. the action they introduce has taken place or will take place).

- (14) 我 不得不 讨 厌 他
I not-should-not hate he
‘I couldn’t help hating him’

Chinese displays a wide array of double negative constructions that have positive meaning. Some of these are *chengyu*, idioms of four characters, where the first and the third characters are sometimes negation markers; For example, the idiom 无往不利, literally meaning ‘there are no places where victory is not achieved’, has a positive meaning that is translated into ‘always successful’.

Similar to English, another issue arising from identifying non-functional negation in Chinese is that negation affixes sometimes do not introduce negation. For instance, in the word ‘disgrace’, the affix ‘dis-’ is *not* considered a cue because the meaning of the whole word is not ‘no grace’, whereas ‘impatient’ is opposed to ‘patient’.

This problem is mostly related and can be solved through *semantic transparency*; that is, if the meaning of the whole cannot be analyzed from the meaning of the parts, we do not annotate part of the word as a negation cue. In the case of 无聊 ‘boring’, literally meaning ‘no chatting’, we do not annotate the negation marker 无 as a cue.

Another criterion that we used is *obsolescence*: if part of a word that is modified by a negation affix is now obsolete, we do not annotate the affix as a cue. This is the case of 不然 ‘otherwise’, where the second character possessed the meaning of ‘like this’ in classical Chinese, but not in modern Chinese.

Finally, we exclude negation cues used in rhetorical and yes-no questions which often take the form of ‘modal + cue + modal’. Some of these, as in (15), are roughly equivalent to the English ‘shall we...?’

- (15) 咱们 要 不 要 向 后 退 ？
we want not want towards back retreat ?
‘Shall we move further back?’

Discontinuous cues. The Chinese equivalent of ‘neither...nor’, ‘既不...也不’, is also a type of discontinuous cues as shown in (16). It is worth mentioning however that omission is a feature of this construction which can be reduced further to ‘不...不’, therefore preserving only the core cues.

- (16) 对 他 (既)不 应该 可 怜 , (也)不 应该 原 谅
towards him not should pity , not should excuse
‘[...] for whom there was neither pity nor excuse’

³To see a list of the core negation cues with examples, please go to the link provided in §3.1.

3.3.2. Scope

Sentential negation. If negation is sentential, i.e. the predicate of a simple clause is negated, we annotate the entire clause under the scope of negation. In the case of two or more coordinated clauses where only one is negated, we annotate as inside the scope only the negated clause. If there is any material missing from the negated clause but retrievable from other parts of the sentence, this retrievable part is annotated as well. (17) exemplifies the annotation of coordination.

- (17) 我把他 弃 而 不 顾 了
I BA him abandon and not care ASP
'I abandoned and did not care about him'

Subordination. If the negated event is only in the subordinate clause, we only consider the subordinate clause for scope annotation but not the matrix clause. If the negated event is only in the matrix clause, subordinates are usually excluded from the scope of negation.

However, Chinese allows for it-cleft constructions like the one in (18), where only the subordinate clause, which appears before the event of the main clause, is in the scope of negation.

- (18) 您 不 是 因 为 知 道 了 这 一 点
才 感 到 高 兴
you not BE because know ASP this point
then feel happy
It is not because you know this that you feel happy.

Relative clauses. If a negation appears in a relative clause, we annotate only the relative clause in the scope of negation but not the head noun it modifies. Unlike English, where the clause follows the head, Chinese displays the opposite order with the particle ‘的’ in between the head and the relative clause. This is exemplified in (19), where the adjective marker ‘的’ separates the relative clause ‘不爱出风头’ and the head ‘人’.

- (19) 他 是 个 不 爱 出 风 头 的 人
he is CL not like show-off DE person
'He is a person who does not like to show off.'

Nominal and adjectival predicates. When negation directly denies a state which is also the main predicate of a clause, the entire clause is under the scope. Whereas in English, these constructions are formed by copula followed by an adjective (“He is impatient”), Chinese do not require a copula. This is shown in (20).

- (20) 这 样 不 公 正
This not fair
'This is unfair.'

With respect to negated adjectives, one important difference between Chinese and English is the status of affixal negation. While affixal negation in English often creates contraries rather than contradictions, hence not forming a scope, in Chinese, an adjective and its negated counterpart usually cannot be false at the same time, therefore abiding by the ‘Law of the Excluded Middle’⁴.

⁴Notice however the definition of affixal negation in Chinese is less clear.

- (21) I am neither patient nor impatient
*我既不耐心也不不耐心
I neither patient nor impatient

As to nominal predicates, where a noun phrase follows a copula (similar to the English “He is not a patient man”), we also annotate the entire clause in the scope, as shown in (22)

- (22) 他 不 是 一 个 耐 心 的 人
He not BE a CL patient DE person
'He is not a patient man.'

Sentence final particles. Chinese is characterized by sentence-final mood particles that serve to express the attitude or mood of the speaker towards the whole sentence. Given that these particles are not affected by the presence of a negation cue, they are *not* included in the scope of negation; this decision is also supported by theories that define these particles as complementisers out of the IP (Paul, 2014). For example in (23), the emphatic mood introduced by the final ‘呀’ is not affected by the negation in the sentence.

- (23) 不 要 等 他 过 了 山 呀!
not need wait he past ASP mountain MOOD
'Do not wait until he has past the mountain!'

Comparative constructions. In Chinese, comparison is expressed in most cases through the co-verb ‘比’, which takes as subject and object the two things compared, followed by the dimension they are compared along. This is the case in (24), where the subject and the object are compared for their age; In such cases where the negation scopes over the co-verb ‘比’, we annotate as scope the entire clause.

- (24) 约 翰 森 先 生 年 纪 不 比 你 大
Johnson Mr. age not compare you old
'Mr. Johnson is not older than you.'

However, when negation scopes over only the dimension being compared (25), we distinguish this from the previous case by excluding from the scope the object of the comparison. This is in line with the English annotation for ‘Compared to usual, my sleep hasn’t been deep’ where the phrase ‘compared to usual’ is excluded from the scope.

- (25) 我 的 觉 睡 的 比 平 常 还 要 不 踏 实
I of sleep sleep compare normal even not sound
'Compared to usual, my sleep has not been deep'

3.3.3. Event

We annotate an event as being negated if it is factual and if it did not happen; ‘eventuality’ includes here both events and states. What the annotation considers as an event is a minimal unit in a negated phrase, usually corresponding to its head. An example of a negated event in verbal predicate is shown inside a box in (26)(we omit the scope just for presentational purposes). Although one could consider 吃羊肉, ‘eat mutton’, as the entire event, the event actually annotated is just its minimal unit, that is, the head verb 吃 ‘to eat’

(26) 我 不 吃 羊肉
 I not eat mutton
 ‘I do not eat mutton.’

Existential and copulative constructions. In existential constructions, we do *not* mark the verb 有, ‘there is/are’, as an event; instead, we mark as the event the head of the nominal phrase following the existential verb as shown in (27). This treatment is consistent with the English guidelines in Morante et al. (2011).

(27) 这里 没有 人
 here not there-is people
 ‘There is nobody here’

As shown in (27), the existential construction in Chinese also encodes universal quantification (i.e. ‘nobody’). When universal quantification applies to the subject of the clause, we annotate as event only the head of the verbal predicate. In (28), we therefore mark as event ‘动’, ‘to move’ but not the aspect marker ‘在’ (continuous action).

(28) 没有 人 在 动
 not there-is person ASP-cont. move
 ‘Nobody is moving’

Modality. Given that we annotate factuality, we do *not* annotate as events those verbs in the scope of certain modals, in particular where the speaker is uncertain about the happening of an event. In English, this excludes most of epistemic and deontic modality (i.e. verbs introduced by auxiliaries such as ‘should’, ‘would’, etc.). In both English and Chinese, we do *not* annotate negated events in the scope of modals except for modality expressing the subject internal ability. This is the case of the modal ‘能’ which is annotated only when expressing participant internal abilities (29.1) but not when expressing conjecture about a non-factual event (29.2).

(29) 1. 我 不 能 打 篮球
 I not can play volleyball
 ‘I cannot play volleyball’
 2. 我 不 能 再 这 样 了
 I not can again like-this ASP
 ‘I couldn’t be like this anymore.’

Supposition or presumption. We also examine the semantics of the verb that is directly negated by the cue, when annotating events. If the verb suggests that the speaker is certain about the content of the predicate, we treat the head of the predicate as factual and annotate the negated event in the clause. If the verb suggests that the predicate is only supposed or presumed by the speaker, we do *not* annotate events in the predicate. This contrast is exemplified in (30.1) and (30.2), through the verbs ‘知道’, ‘to know’, where we annotate the event, and ‘相信’, ‘to believe’, where we do not.

(30) 1. 我 知 道 您 决 不 愿 意 [...]
 I know you not want [...]
 ‘I know that you do not wish [...]’

2. 我 相 信 您 决 不 愿 意 [...]
 I believe you not want [...]
 ‘I believe that you do not wish [...]’

4. Annotation Projection

4.1. Methodology

The goal of the annotation projection is to investigate whether we can ease the burden of annotating from scratch in the presence of parallel text.

Annotations are projected using word alignment information computed by the IBM model 2, as implemented in the fast_align toolkit (Dyer et al., 2013).

The training data for the alignment model consists of the aligned sentence pairs in NegPar and the English-Chinese UN parallel corpus (Ziemski et al., 2016); the Chinese side of the corpus was also tokenized using the Stanford Word Segmenter (Tseng et al., 2005). We used the symmetrical two-way alignment results as the basis for projection. In our work, we experimented with two types of alignment models: (1).English word to Chinese word (word-level projection); (2).English word to Chinese character (character-level projection). An example of the former is shown in Fig. 2 where all elements are projected correctly, except for the scope projection of the subject ‘我’(‘I’).

For both levels of projection, we report precision, recall, F1 measure and number of gold (both English and Chinese) and projected spans for cue, event and scope independently, as they were projected as such.

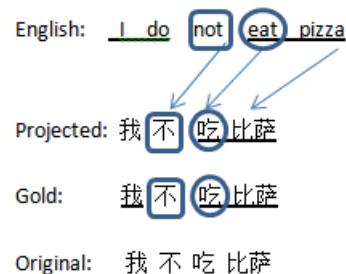


Figure 2: An example of projection from an English sentence to its Chinese translation: cue is indicated with a box around, event is circled and the scope is underlined. Word alignment (English word to Chinese word) is indicated by arrows.

4.2. Results in the projection task

The results for cue, scope and event projection on the development set are summarized in Table 2.

Considering F1 alone, we found word-level projection to yield better results for event and scope but not for cue projection. This can be explained by the fact that cues often span subword units (as in the case of morphological negation) and word-level projection might end up over-predicting cues in Chinese (hence the relatively lower precision). However, in terms of easing the annotation process, a relatively lower recall means that more work is required to find elements that the projection has missed, rather than

	Word-level			Character-level			gold-en	gold-zh	projection (% gold-zh)
	precision	recall	F1	precision	recall	F1			
cue	0.39	0.47	0.43	0.49	0.42	0.45	175	230	169(73%)
event	0.37	0.32	0.34	0.40	0.27	0.32	123	153	103(67%)
scope	0.64	0.48	0.55	0.64	0.44	0.50	170	226	168(73%)

Table 2: Performance of the annotation projection on the *dev* set both at word and character level for cue, event and scope respectively.

filtering out over-predicted elements (reflected in lower precision) which might be done with some post-processing heuristics.

In general, even in the presence of parallel data, detecting negation using annotation projection does not lead to good results. The number of projected spans vs. gold spans in Chinese and English suggests that this is in part due to differences in how negation is translated.

4.3. Error Analysis

To delve deeper into the errors in projecting the annotations from English, we carry out an error analysis. In doing so, we consider character-based projection for the cue and word-based projection for event and scope spans.

Cue. we first break down the performance of cue projection according to different target and source cues. First of all we found performance to vary across different cues as shown in Fig. 3. Compared to precision, recall is lower for two-character cues where projection often seems to miss either of the characters; this is the case of both 没有 and 并不. The low performance of 无 and 未 might be caused by 无 and 未 being common components in Chinese chengyu (idioms) that are frequently used to translate positive phrases in English as in (32) (8 cases).

Furthermore, we analyzed projection errors based on source cues as shown in Fig 3.(b). The relatively higher performance compared with Fig 3.(a) is not surprising as we exclude here a large number of errors from cases where negation expressed in Chinese but not in English. In general, when the English cues are correctly projected onto the Chinese cues, the projection might also include additional surrounding characters in Chinese as cues; therefore the recall is higher than precision. It is especially the case for the negation pronoun 'nothing' which maps to the negative polarity item (NPI) 什么都/一点也/一点/一, 'any/anything'. This accounts for 5 errors, one of which we report below (31).

- (31) projected.: 从他那里什么都得不到
gold-zh: 从他那里什么都得不到
from him there anything DOU get not POT
'gold-en: can get **nothing** from him'

Finally, we wanted to analyze those cases where projection fails due to the fact that negation is present in Chinese but not in English. we found this to happen for 2 different reasons:

1. A negation instance in Chinese is paraphrased in positive terms in English; this often concerns just a pair of contrary adjectives or adverbs as shown in (32), but also extends to entire clauses as shown in (33)




Figure 3: Performance of annotation projection plotted against Chinese (a) and English (b) cues

- (32) gold-zh: 他安然无恙
he safe not sick
gold-en: 'He is safe and sound.'
- (33) gold-zh: 惊慌的脸上没有一点血色
panicked DE face not there-is one-bit blood-color
gold-en: 'Every tinge of colour struck from his astonished face'

2. Some lexical items in English can be interpreted as inherently expressing negation and thus can be translated as cues in Chinese, but they are usually not annotated as cues on the English side. This is the case in (34); along with 'hardly', we found other expressions such as 'rather than', 'absence', 'out of question' and 'refused' that can be translated into negation in Chinese.

- (34) gold-zh: 这件事的前前后后不可能是为了 [...]
This CL thing of everything not can be for [...]
gold-en: 'The whole proceeding could hardly be for [...]

Event. Out of the 153 gold events we found that only 13.3% were correctly projected from English, with 38.8%

of the cases where the projection does not detect an event at all. These false negatives are caused by the fact that Chinese translates positive terms in English as negation (same as the cases in (32)~(34)) but in some cases are just due to English words aligning to a null token.

On the other hand, in 16% of the cases we observed that the event from English is projected onto a completely different span of the sentence. Some of these cases are however not due to alignment errors but because the Chinese side uses different constructions with emphasis on different events from English. For instance, in (35), the English guidelines annotate as event the nominal predicate ‘colour’ where this is translated in Chinese as a verb 说上 ‘say’.

(35) gold-zh: 那张脸既不黑 [...] 说不上
 上 是什么颜色
 That CL face either not black [...] say not
 up be what colour
 gold-en: It was n’t black [...] nor any colour

Finally, in 13.3% of the cases, the projection only partially matches the gold annotation for an event. 13 of these 25 cases we found that the projection includes the negation cue inside the event. These are often cases such as (36), where a word containing morphological negation in English is projected onto both the cue and the event in Chinese.

(36) projected: ‘[...] 我们还弄不清楚的罪行’
 gold-zh: [...] 我们还弄不清楚的罪行
 [...] we still manage not clear DE crime
 gold-en: ‘They were all confederates in the same
 un-known crime.’

Scope. Out of the 226 instances of negation scope we found that only 3 (0.01%) were fully and correctly projected, with 39% of the cases where the projection returns nothing. As discussed above, this is probably due to positive phrases being translated into negation in Chinese.

We found only 5 cases (0.02%) where the scope in English is projected to a completely different span from the gold span in the Chinese sentence. The majority of the errors (145/226 – 77%) concerns partial overlap, where the projection covers the gold scope only in part. A closer analysis shows that the projection tends to often miss the NPI 什么 (12 cases). 什么 corresponds to the English ‘any’ when in the scope of negation. In all the cases where projection fails to include this element in the scope, English uses the determiner ‘no’ or pronoun ‘nothing’ instead of an overt NPI eg. ‘any’; therefore 什么 is usually mapped to the negation cue rather than being marked as a scope element; this is exemplified in (37).

(37) projected: 这里面没有什么
 gold-zh: 这里面没有 什么
 Here not there-is anything
 gold-en: **Nothing in all this**

Finally, to allow comparison with future work, Table 3 and 4 report the performance of annotation projection on the test set as well.

	precision	recall	F1
cue	0.372	0.428	0.398
scope	0.574	0.381	0.458
event	0.299	0.209	0.246

Table 3: Results from word-level projection of negation on the test data

	precision	recall	F1
cue	0.478	0.382	0.425
scope	0.583	0.312	0.406
event	0.338	0.180	0.235

Table 4: Results from character-level projection of negation on the test data

5. Conclusion

We have introduced *NegPar*, the first English-Chinese parallel corpus annotated for negation. The corpus is based on a pre-existing negation corpus for English, CONANDOYLE-NEG, whose annotation we have amended and extended onto its Chinese translation. To ease the burden of annotating from scratch, we have experimented with projecting the English annotations onto Chinese via word-alignments. Results have shown that the projection offers limited help as it would not deal with the different ways negation can be translated across languages.

Future work could investigate whether differences in translating negation vary according to languages and domains. The annotation projection algorithm could also be improved for example by considering syntactic mapping.

6. Appendix A: Abbreviations

ASP: aspect marker
 BA: object marker
 CL: classifier
 DE: adjective marker
 MOOD: mood particle
 ASP-cont: Present tense marker
 DOU: particle (all)
 POT: verb complement expressing the potential form

7. Acknowledgements

We would like to thank the three anonymous reviewers for their useful comments. This project was also funded by the European Union’s Horizon 2020 research and innovation programme under grant agreement No 644402 (HimL). We would also like to thank Peterhouse College Cambridge for funding for the conference trip.

8. Bibliographical References

Dyer, C., Chahuneau, V., and Smith, N. A. (2013). A simple, fast, and effective reparameterization of ibm model 2. Association for Computational Linguistics.
 Hwa, R., Resnik, P., Weinberg, A., and Kolak, O. (2002). Evaluating translational correspondence using annotation projection. In *Proceedings of the 40th Annual*

- Meeting on Association for Computational Linguistics, pages 392–399. Association for Computational Linguistics.
- Konstantinova, N., De Sousa, S. C., Díaz, N. P. C., López, M. J. M., Taboada, M., and Mitkov, R. (2012). A review corpus annotated for negation, speculation and their scope. In LREC, pages 3190–3195.
- Li, C. N. and Thompson, S. A. (1989). Mandarin Chinese: A functional reference grammar. Univ of California Press.
- Morante, Roser and Daelemans, Walter. (2012). ConanDoyle-neg: Annotation of negation in Conan Doyle stories.
- Morante, R., Schrauwen, S., and Daelemans, W. (2011). Annotation of negation cues and their scope: Guidelines v1. Computational linguistics and psycholinguistics technical report series, CTRS-003.
- Padó, S. and Lapata, M. (2009). Cross-lingual annotation projection for semantic roles. Journal of Artificial Intelligence Research, 36(1):307–340.
- Paul, W. (2014). Why particles are not particular: Sentence-final particles in chinese as heads of a split cp. Studia Linguistica, 68(1):77–115.
- Tseng, H., Chang, P., Andrew, G., Jurafsky, D., and Manning, C. (2005). A conditional random field word segmenter for sighthan bakeoff 2005. In Proceedings of the fourth SIGHAN workshop on Chinese language Processing.
- Vincze, V., Szarvas, G., Farkas, R., Móra, G., and Csirik, J. (2008). The bioscope corpus: biomedical texts annotated for uncertainty, negation and their scopes. BMC bioinformatics, 9(11):S9.
- Ziemski, M., Junczys-Dowmunt, M., and Poulliquen, B. (2016). The united nations parallel corpus v1. 0. In LREC.
- Zou, B., Zhu, Q., and Zhou, G. (2015). Negation and speculation identification in chinese language. In ACL (1), pages 656–665.

9. Language Resource References

- Morante, Roser and Daelemans, Walter. (2012). ConanDoyle-neg: Annotation of negation in Conan Doyle stories.