

Representing Implicit Positive Meaning of Negated Statements in AMR

Katharina Stein Lucia Donatelli

Department of Language Science and Technology
Saarland Informatics Campus
Saarland University, Germany

{kstein, donatelli}@coli.uni-saarland.de

Abstract

Abstract Meaning Representation (AMR) has become popular for representing the meaning of natural language in graph structures. However, AMR does not represent scope information, posing a problem for its overall expressivity and specifically for drawing inferences from negated statements. This is the case with so-called “positive interpretations” of negated statements, in which implicit positive meaning is identified by inferring the opposite of the negation’s focus. In this work, we investigate how potential positive interpretations (PPIs) can be represented in AMR. We propose a logically motivated AMR structure for PPIs that makes the focus of negation explicit and sketch an initial proposal for a systematic methodology to generate this more expressive structure.

1 Introduction

There has been a recent and growing interest in representing the meaning of natural language in graph structures. One semantic framework that has become popular is the semantic representation language of Abstract Meaning Representation (AMR) (Banarescu et al., 2013). AMRs are attractive for their ease of readability, and they are intuitive representations due to their focus on predicate-argument structures. Additionally, annotating meaning with graph structures is easier than following many traditional formal semantic frameworks (Bos, 2016) and often allows for increased expressive capacity.

An important application for meaning representations are inference tasks that play a crucial role in many NLP tasks such as information extraction and question answering (Abzianidze and Bos, 2019). Such inference tasks are where AMR often falls short due to its lack of precise scope representation. For example, from (1-a) it can be inferred that “*The boy stole the money*” and hence the boy is a thief; from (1-b), it can be inferred that “*The boy did not steal the money*” and is likely not a thief.

- (1) a. The boy stole the money yesterday.
- b. The boy did not steal the money yesterday.

Yet, the information provided by the scope of the negation alone is not sufficient when drawing inferences. Negated statements can convey implicit positive meaning, or Potential Positive Interpretation (PPI) (Blanco and Moldovan, 2011), that is usually understood intuitively by humans. For example, an intuitive interpretation for (1-b) is that “*someone other than the boy stole the money yesterday*” (and the boy is not a thief). It could also be inferred, however, that “*the boy stole something other than money*,” (and is still a thief) depending on whether *the boy* or *the money* is the focus of the negation. As a third alternative, if the focus of the negation is *yesterday*, we can indeed infer that “*the boy stole the money*,” though we do not know when. Both scope and focus of negation need to be detected to understand these nuanced interpretations and draw correct inferences, an existing weakness in many current Natural Language Inference (NLI) tasks (Naik et al., 2018).

The goal of this paper is therefore (i) to investigate how the PPIs of negative statements can be represented in AMRs and (ii) to investigate whether these structures can be generated from the AMR of the negated sentence in a systematic way based on the work by Blanco and Sarabi (2016) and Sarabi and Blanco (2016). To this end, we propose a novel, logically-motivated AMR structure that makes both scope and focus of negation explicit. We additionally sketch an initial proposal for how to transform existing AMRs into this more expressive format. Initial assessment results support previously made observations that AMR lacks expressive capacity when it comes to scope. This lack of expressivity is especially problematic for PPIs whose concrete meaning results from an interaction of the negation and information structure. Our methodology

	Semantic Role	Dependency	Potential Positive Interpretations
1	ARG ₀	<i>nsubj</i>	[Someone] stole my money yesterday, <i>but not the young boy</i>
2	ARG ₁	<i>obj</i>	The young boy stole [something] yesterday, <i>but not my money</i>
3	<i>verb</i>	<i>root</i>	The young boy [some verb] my money yesterday, <i>but not stole</i>
4	ARGM-TMP	<i>obl:tmod</i>	The young boy stole my money [at some point in time], <i>but not yesterday</i>
5		<i>nsubj</i>	The young [some person] stole my money yesterday, <i>but not the young boy</i>
6		<i>amod</i>	The [some adjective] boy stole my money yesterday, <i>but not the young boy</i>
7		<i>nmod:poss</i>	The young boy stole [someone's] money yesterday, <i>but not my money</i>
8		<i>obj</i>	They young boy stole my [something] yesterday, <i>but not my money</i>
Positive Counterpart: The young boy stole my money yesterday.			

Table 1: The PPIs generated using semantic roles (1-4) and syntactic dependencies (1-8) for the negated statement: "The young boy did not steal my money yesterday".

enables us to express the negation's foci in AMR without modifying the original AMR. To our knowledge we are the first to propose a representation of both scope and focus in AMR in order to express the full meaning of negated statements.

This paper is structured as follows. Section 2 presents background about focus of negation, previous work on generating PPIs and about scope in AMR. Based on this, we propose an approach to represent the PPIs of a negated statement in AMRs and a methodology to systematically generate these AMRs (Section 3) which we then verify with corpus examples (Section 4). In Section 5 we conclude and discuss next steps.

2 Background and Related Work

2.1 Focus of Negation

Deriving the meaning of negated statements in NLP tasks is more challenging than for statements lacking negation. This complexity is due in part to the ambiguity of negation with respect to its scope and focus. According to Huddleston and Pullum (2002), the focus is the part of the scope that is intended to be interpreted as being false. Blanco and Moldovan (2011, 2014) state that negated statements can carry implicit positive meaning. For example, "cows do not eat *meat*" has the implicit meaning that "cows eat something other than meat". They state that implicit positive meaning arises because of a speaker's specific choice to formulate the negated statement, i.e. if the positive implication is not intended the speaker would say something different (for example, "cows do not eat").


The ability to detect the positive meaning by detecting the scope and focus of negation yields valuable information for NLP tasks, especially for

those involving inferences (Blanco and Moldovan, 2011, 2014). Selecting the smallest suitable focus, i.e. a fine-grained focus, often yields a larger amount of positive meaning (Blanco and Moldovan, 2012, 2014). Indeed, this has been found in NLI benchmarks: strong negation words are often responsible for models prediction contradiction for neutral or entailed statements (Naik et al., 2018). Representing scope and focus in a more precise manner can further assist in fact detection, building correct timelines, and identifying discourse structure (Altuna et al., 2017; Anand and Martell, 2012).

2.2 Positive Interpretations of Negated Statements

Blanco and Sarabi (2016) and Sarabi and Blanco (2016) developed methodologies that make the implicit positive meaning of a negated statement, or all positive interpretations derived from potential foci of negation, explicit by automatically generating its PPIs based on the focus.

Focusing on verbal negation the positive counterpart is first generated by removing negation; each PPI is then generated by selecting one of the potential foci of the negation and replacing it with an "existential placeholder". The potential foci get identified based on the semantic roles (Blanco and Sarabi, 2016) or the syntactic dependents (Sarabi and Blanco, 2016) of the verb. Table 1 shows both the coarse-grained (1-4) and the fine-grained PPIs (5-8). Sarabi and Blanco (2017) extend the syntactic dependency based methodology to address the negation of nouns and adjectives.



(a) Graph representation

$$\begin{aligned} & \text{instance}(s, \text{steal} - 01) \wedge \text{instance}(b, \text{boy}) \wedge \\ & \text{instance}(m, \text{money}) \wedge \text{instance}(y, \text{yesterday}) \\ & \wedge \text{ARG0}(s, b) \wedge \text{ARG1}(s, m) \wedge \\ & \text{time}(s, y) \wedge \text{polarity}(s, -) \end{aligned}$$

(b) Conjunction of Logical Triples

Figure 1: AMR representations for the sentence “*The boy did not steal the money yesterday.*” Information about negation’s scope and focus is not captured.

2.3 Abstract Meaning Representation

AMRs are single-rooted, directed, labeled acyclic graphs whose nodes represent the predicate and the arguments of the represented sentence and the edges represent the semantic relations between the concepts. Figure 1a shows the graphical AMR for the sentence “*The boy did not steal the money yesterday.*” Negation is represented with a `:polarity` relation between the negated concept and the polarity constant “-”. This representation does not provide any information about the scope of the negation, seen clearly when represented as a conjunction of logical triples helpful to draw inferences (Figure 1b).

Without negation, this logical representation of an AMR can successfully be used to draw inferences by applying the conjunction elimination rule. However, as negation is represented as a conjunct, the elimination rule can yield wrong results when drawing inferences. As already stated by Bos (2020), applying the rule allows us to erroneously drop the negation. From Figure 1b, it could be inferred that “*The boy stole the money yesterday.*”. Dropping one of the other conjuncts can lead to wrong inferences, as well, if the focus of the negation is dropped but not the negation itself.

Several approaches have been proposed to address the lack of scope information in AMR (Bos, 2016, 2020; Pustejovsky et al., 2019). Pustejovsky et al. (2019) propose an extended version of AMR that introduces a node for scope that can be added on top of the AMR structure and represents the relative scope between two other nodes, e.g. a quantifier and negation. Bos (2016) defines a recur-

sive translation function that maps an AMR into first-order logic (FOL). The resulting formulas include existential quantifiers and the negation as a logical operator that scopes over the sub-formula corresponding to the sub-graph it is attached to. However, none of these approaches addresses the representation of focus of negation.

The overall focus of a statement can be represented in AMR by making the corresponding node the root and re-arranging the rest of the graph using inverse roles. However, the expressiveness of this focus representation is limited. Specifically, only a single node can be the focus, posing a problem for PPIs where the focus of the negation consists of several nodes. Additionally, leaving the negation at the original position does not make clear that focus of the negation is expressed.

Bos (2016) shows that his translation function can derive formulas where the logical negation scopes only over a sub-formula when moving the negation in the AMR to the corresponding sub-graph. For example, the formula for “*It was not a boy who stole the money yesterday*” could be correctly derived by moving the negation to `boy`. However, the resulting AMR and the corresponding FOL formula do not explicitly represent that there exists another person who is the agent of this stealing event. Furthermore, moving only the negation does not allow to differentiate between the representation of coarse-grained PPIs and PPIs where the focus is a single, internal node.

3 Method

The goal of this exploratory work is to investigate (i) how the PPIs of a negated statement can be represented in AMRs and (ii) if it is possible to generate them systematically from the original AMR of the negated statement for increased expressivity. We employ a combination of the semantic role based (Blanco and Sarabi, 2016) and the syntactic dependency based approach (Sarabi and Blanco, 2016) as the starting point for this investigation. For simplicity, we ignore tense as well as semantic differences between definite and indefinite determiners.

3.1 Representing PPIs in AMRs

The main question regarding PPIs in AMRs concerns the representation of “*but*”. Specifically, we must decide where the subgraph for the “*but not* [original text of focus]” part of the PPI should be inserted. A manual investigation of AMRs for simi-

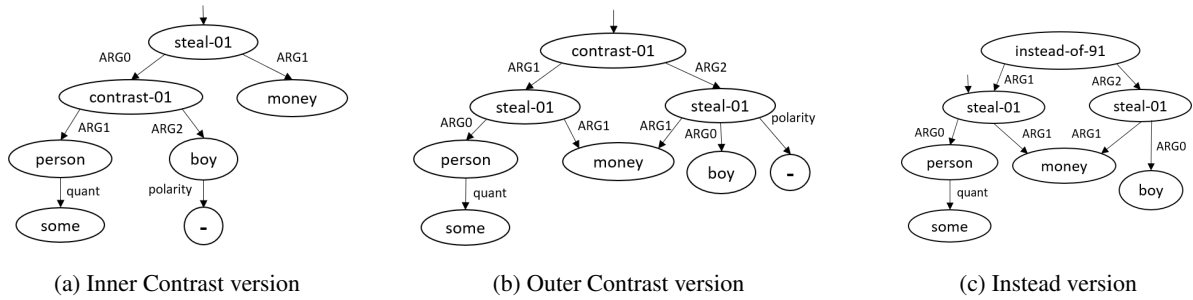


Figure 2: The three representation structures considered for representing PPIs for the example ARG0 PPI "Some person stole the money but not the boy".

	Inner Contrast	Outer Contrast	Instead
Pro	<ul style="list-style-type: none"> • "but" similar to "and" • logically motivated • makes focus explicit 	<ul style="list-style-type: none"> • most frequent way to encode "but" in similar cases 	<ul style="list-style-type: none"> • <code>instead-of-91</code> to express "A but not B"
Con	<ul style="list-style-type: none"> • only very few examples in AMR corpora with an inner <code>contrast-01</code> node 	<ul style="list-style-type: none"> • Negation attached to main predicate • focus not explicit, can only be inferred from not shared arguments 	<ul style="list-style-type: none"> • no polarity relation • only implicit negation in AMR

Table 2: The advantages and disadvantages of the three considered AMR versions with respect to the current work.

lar sentence structures in AMR corpora¹ resulted in three potential representations (see Figure 2). Two of them represent "but" using `contrast-01`: *Inner Contrast* with `contrast-01` as an inner node and *Outer Contrast* with `contrast-01` as the root node. *Instead* represents "but" with the `instead-of-91` concept.

Consider the negated statement in (2) with its potential PPIs based on different foci of the negation:

- (2) The boy did not steal the money.
 - a. ARG0: Some person stole the money but not the boy.
 - b. ARG1: The boy stole something but not the money.
 - c. ROOT: The boy did something with the money but not steal

The AMR for the negated statement corresponds to the one in Figure 1 without the `:time` relation. Figure 2 illustrates the three different AMR versions considered for the ARG0 PPI interpretation (2-a) (when ARG0 is the focus); Table 2 summarizes the main advantages and disadvantages of the versions with respect to the current work.

The `instead-of-91` relation was considered because the AMR dictionary² lists "not A but B" as an example use and the PPIs can be rephrased to e.g.

¹<https://amr.isi.edu/download/amr-bank-struct-v3.0.txt>, <https://amr.isi.edu/download/2018-01-25/amr-release-bio-v3.0.txt>

²<https://amr.isi.edu/doc/amr-dict.html>

"The boy stole not the money but something yesterday". The main difference compared to the *Contrast* versions is the absence of the `:polarity` relation, as the negation is implicitly included in the ARG2 role by definition. As one of the motivations for this work is to make the meaning of negation explicit, it would be counterintuitive to use a representation that expresses negation implicitly.

The *Inner Contrast* and the *Outer Contrast* versions have quite different overall structures. The *Outer Contrast* AMR has the `contrast-01` concept as its root and its two arguments are themselves AMRs for a complete sentence each. In particular, the ARG2 subgraph corresponds to the original negated AMR and the ARG1 subgraph to the AMR for the PPI without the "but not [original text of focus]" part, with re-entrancies for shared arguments. In the *Inner Contrast* version on the other hand, the `contrast-01` concept replaces the node that is the root of the chosen focus, e.g. the ARG0 argument for the ARG0 PPI. The AMR for the placeholder is the ARG1 argument of the `contrast-01` concept and the original ARG0 subgraph, (here *boy*), becomes the ARG2 argument with the negation directly attached to it instead of to the root predicate. The structure of this graph and the representation of "but" are similar to "and" and how "Some person and not the boy stole the money" gets represented.

Interestingly, in a manual survey of similar ex-

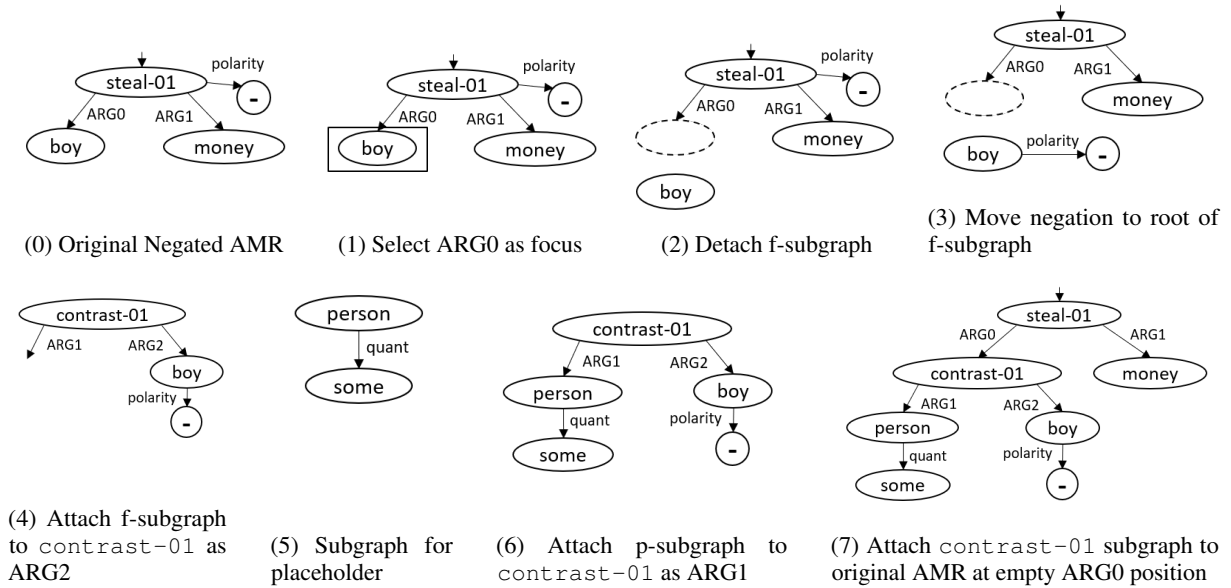


Figure 3: The steps of the initial version of the methodology to generate the AMR for the ARG0 PPI "Some person stole the money but not the boy" (7) from the AMR of the original negated statement (0).

amples encountered in AMR corpora, almost all representations correspond to the *Outer Contrast* version. Nevertheless, the *Inner Contrast* version has two advantages for the current work. First, the structure is more logically motivated than the *Outer Contrast* structure, as "but" often gets represented as "and not" in logic and hence a similar structure in AMRs seems to be preferable. Additionally, the *Inner Contrast* version makes the focus of the negation explicit as the negation is attached directly to the concept that is intended to be interpreted as false. For these reasons, we choose the *Inner Contrast* version to represent PPIs.

3.2 Generating the PPI Representations

Based on the general structure of the *Inner Contrast* AMRs and Sarabi and Blanco (2016)'s methodology, we propose to systematically generate PPI AMRs. As an initial step (0), the original AMR is traversed, starting at the node the polarity is attached to; all subgraphs rooted at the visited nodes are considered as potential foci (subgraphs), except for the polarity node. We then proceed as follows:

1. We select the focus of negation;
2. The chosen focus subgraph (f-subgraph) is detached from the original AMR;
3. The negation is moved to the root of the f-subgraph;

4. The f-subgraph is attached to a `contrast-01` node as ARG2;
5. We create a subgraph for the placeholder (p-subgraph);
6. The p-subgraph is attached to a `contrast-01` node as ARG1;
7. The `contrast-01` subgraph is then attached at the original position of the f-subgraph.

Figure 3 illustrates these steps graphically. This transformation can be executed directly on the original AMR of the negated statement and hence no positive counterpart needs to be derived as an intermediate step. As AMRs do not represent tense, inflection and auxiliaries required by the syntax, only the negatively-oriented polarity-sensitive items have to be exchanged with their positively-oriented counterparts.

4 Findings

In the following section, the validity of the initial proposal for the representation of PPIs of negated statements in AMRs and for the methodology to generate them will be verified by means of examples and extended and adapted if necessary. For reasons of space, only some of the PPIs are presented for the example statements.

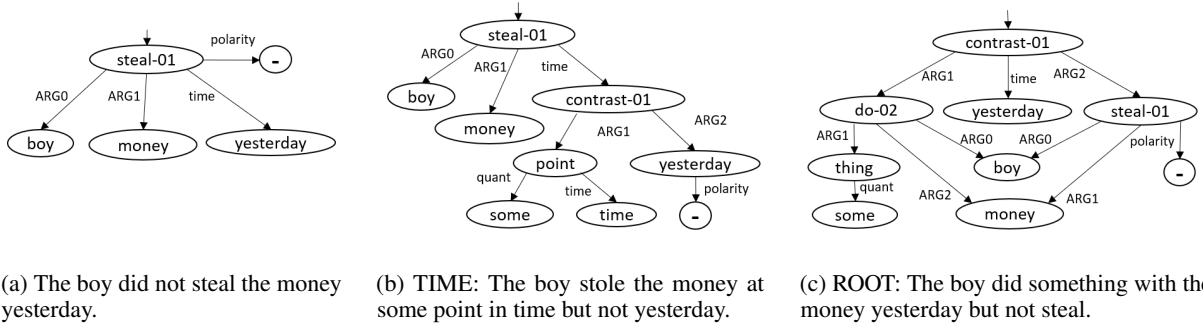


Figure 4: Verbal negation: AMRs for "The boy did not steal the money yesterday" and two of its PPIs.

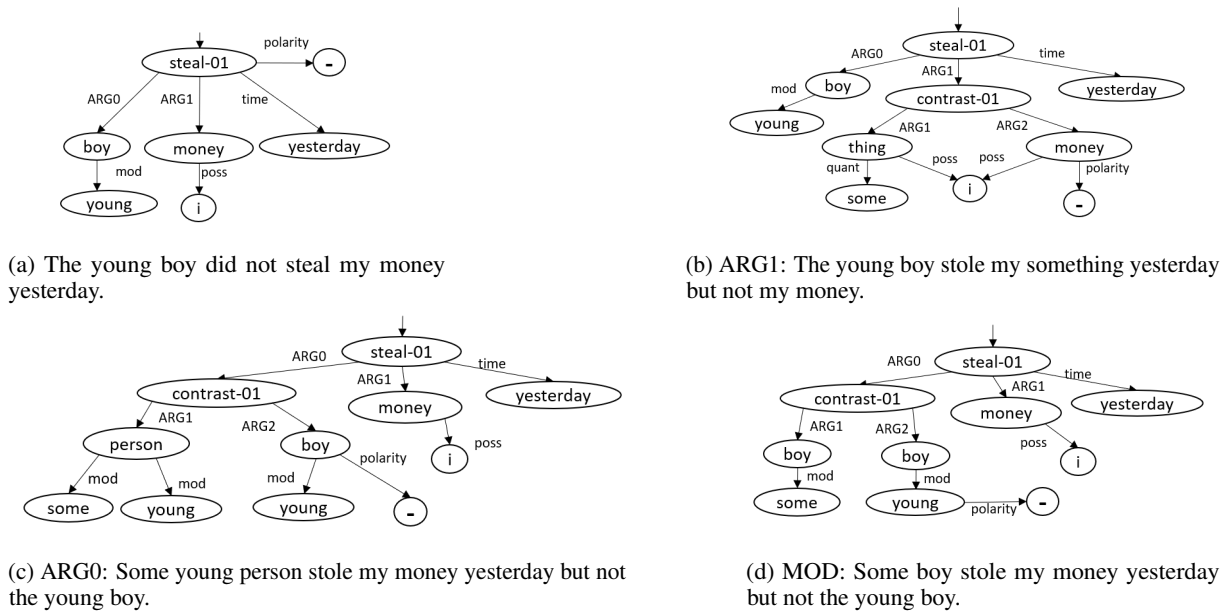


Figure 5: Fine-grained PPIs: AMRs for "The young boy did not steal my money yesterday" and three of its PPIs.

4.1 Verbal Negation

Figure 3 illustrates how to generate the AMR for an ARG0 PPI. Example (3) extends (2) with the non-core relation `:time` (Figure 4a).

- (3) The boy did not steal the money yesterday.
- (4) The young boy did not steal my money yesterday.

Figure 4b shows that the AMR for the PPI generated from `:time` has the same overall structure as when choosing a core relation as focus. When the negated verb itself is chosen as the focus, the AMR for the corresponding ROOT PPI looks different (see Figure 4c). In order to generate the AMR for ROOT PPIs, the predicate and all its arguments need to be considered as the f-subgraph and two additional steps must be added to the transformation between step (6) and (7). First, *boy* and *money* are the ARG0 and ARG1 arguments of both predicates

of the PPI respectively and hence need to be shared. Second, non-core relations modifying a complete conjunction are pulled-out in AMRs and therefore *yesterday* is moved to the root node.

Example (4) (Figure 5a) adds the modifier *young* and the possessive pronoun *my* to the ARG0 and ARG1 arguments of (3) respectively, leading to four additional, fine-grained PPIs. Figure 5b shows the AMR for the ARG1 fine-grained PPI. Similar to ROOT PPIs, fine-grained foci can be internal nodes and therefore one has to take care of the required re-entrancies, e.g. the two `:poss` relations in 5b need to link to the same entity (i.e. "I").

Whereas the previous case works similar to the coarse-grained PPIs, it gets more difficult when adjectives are involved. The two fine-grained PPIs that Sarabi and Blanco (2016)'s methodology generates for the ARG0 argument *the young boy* are "The young some person stole my money yesterday but not the young boy" and "The some_adjective

boy stole my money yesterday ...". Instead of *some_adjective*, we use *some* as a general placeholder for adjectives. Still, both PPIs sound rather unsuitable because of the co-occurrence of the determiner and *some*. Possible solutions would be to remove the determiner or to exchange it with an indefinite determiner and to remove *some* completely. To ensure consistency with negated statements not including a determiner, e.g. "*Young boys did not steal my money*", we decided to discard the determiner which also never is represented in AMR (see Figure 5c, 5d).

In case of the MOD PPI yet another issue comes up. The AMR generated by the proposed transformation steps represents only one boy that "*has some property but is not young*" because the contrast subgraph is attached at the original position of *young* beneath *boy*. However, in the MOD PPI two different boys are referenced. Therefore, the methodology in Figure 3 has to be adapted in the following way in order to generate the AMR in Figure 5d: In step (1) not only the f-subgraph but also the subgraph covering the complete semantic role to which the fine-grained focus belongs (sr-subgraph) has to be detected. Both subgraphs are equivalent in the case of coarse-grained and ROOT foci. In all steps except for step (3) "f-subgraph" needs to be replaced with "sr-subgraph". Additionally, the sr-subgraph needs to be copied before adding the negation in step (3). The p-subgraph in step (5) replaces the f-subgraph in the copied sr-subgraph and the resulting subgraph is then attached to the contrast subgraph in step (6). For the examples covered so far these modifications do not affect the resulting PPI AMRs. Hence, the adapted methodology is more suitable. (See Appendix for detailed illustrations of the adapted steps for a MOD PPI and a ROOT PPI example.)

4.2 Negated Adjectives

Sarabi and Blanco (2017) extend the methodology from Sarabi and Blanco (2016) to non-verbal negation by modifying the selection of the potential foci. When an adjective is negated, they choose all subtrees rooted in the direct dependents of the negated adjective as well as the adjective itself as potential foci. Here, we only consider the negation of predicative adjectives, i.e. where the adjective concept is the root of the AMR. Predicative adjectives are represented using a PropBank frame whenever one is available, such as for *sad* in Example (5).


(5) The child was not sad at the party.

Hence, the corresponding AMR structures are similar to those for verbal negation (see Figure 6a). However, when exchanging *sad* with *funny* the root is no longer related to *child* with a core relation but with the `:domain` relation because no PropBank frame is available (Figure 6b). Still, generating the AMRs for the PPIs works in the same way as for verbal negation except for the ROOT PPI. For predicative adjectives neither *some_adjective* nor *some* are suitable placeholders. Therefore, in order to generate the ROOT PPIs for negated adjectives first an appropriate concept for the placeholder needs to be determined (see Figure 6c).


4.3 Negated Nouns

Following Sarabi and Blanco (2017), we distinguish two cases when a noun is negated: when the negated noun is the root of the AMR and when this is not the case. If a negated sentence contains a main verb as in Examples (6) and (7) the structure of the AMR depends on the specific meaning. (6) can be rephrased as "*The boy did not steal any money yesterday*" and gets represented with the same AMR as shown in Figure 1a. Hence, it can be treated as an instance of verbal negation. (7) on the other hand cannot be rephrased this way but as "*There was no adult who spoke loudly in the church*" and therefore the noun is the root of the corresponding AMR (Figure 7e). Still, the PPIs and their AMRs are the same as for verbal negation (e.g. Figure 7f). The only adaption required in order to apply the proposed methodology to AMRs with an negated noun as root node is to make the closest predicate the root instead of the noun before selecting the foci and generating the AMRs.


Yet, there is one limitation of the PPIs generated for negated statements with the meaning "there is no X that" which concerns quantification. When generating PPIs from foci different from the negated noun itself it is important that the PPI still expresses the negated existential quantification, e.g. the MANNER PPI for (7) must not express only that "an adult spoke in some manner in church but not loudly" because this does not exclude that "there was also an adult who spoke loudly in church". We decided to replace the negation *no* with a bare plural instead of using an indefinite article as done by Sarabi and Blanco (2017) because the noun phrase needs to be generic and the use of indefinite singular generics is more re-



(a) The child was not sad at the party.




(b) The child was not funny at the party.




(c) ROOT: The child was [??] at the party but not sad.


Figure 6: AMRs for the examples of negated adjectives and the overall structure of a ROOT PPI.




(a) This man is no professor.




(b) This man is no teacher.




(c) This man is no president.



(d) ROOT: This man is a person who does something but not teach.



(e) No adult spoke loudly in the church.



(f) MANNER: Adults spoke in some manner in the church but not loudly.

Figure 7: AMRs for the examples of negated nouns and some of their PPIs.

stricted than the use of bare plural generics (Cohen, 2001). However, this adaption does not solve the issue on the AMR level because AMRs do not represent (most) determiners and do not have a deep representation for quantifiers (see Figure 7f). For the current work, the problem could be addressed by using the *Outer Contrast* version for the PPIs but nevertheless, this example supports the need to extend the expressiveness of the AMR framework.

(6) The boy stole no money yesterday .

(7) No adult spoke loudly in the church.


(8) This man is no professor.

When a copula verb is involved like in Example (8), the structure of the AMR depends on the representation of the specific noun. When no appropriate frame is available the noun itself is the root (Figure 7a). Exchanging *professor* with *teacher* leads to a different structure because *teacher* gets represented as *a person who teaches* (Figure 7b) and exchanging it with *president* results yet in another graph structure, because *president* gets represented by the *have-org-role-91* frame (Figure 7c).

The PPIs for foci that are not the negated noun

itself and their generation do not differ from the examples in the previous sections. However, when the negated noun is chosen as the focus, each of the three examples has its own particularities. For the *president* example the only difference compared to verbal negation is that it does not make sense to select the root as potential focus. Example (8) illustrates again the problem of a suitable placeholder as the AMR for the ROOT PPI "*This man is something but no professor*" is a less-than-ideal solution because of *something*. When the negated noun invokes a predicate this is easier because "*This man is no teacher*" is represented as "*This man is not a person who teaches*". However, this results in an additional potential focus, i.e. *person*. One approach to deal with this inconsistency would be to consider the predicate and *person* as one unit for nouns that invoke predicates. Additionally, the position where the contrast subgraph gets inserted into the original graph needs to be modified in order to result in the correct AMR (see Figure 7d).

The examples in this section illustrate that representing and generating PPIs for negated statements involving the copula verb *to be* that gets not represented in AMRs is complex due to differences in



(a) The girl was not sad because the boy left.

(b) CAUSE: The girl was sad because of something but not because the boy left.

(c) The girl was sad because something happened but the boy did not leave.

Figure 8: AMR for Example (9) (a) and the correct CAUSE PPI (b) and the incorrect PPI (c).

the AMR representations of specific nouns.

4.4 Limitation: Negating Causation

In the case of negated statements that include an ARGM-CAU semantic role factuality plays a role when selecting the causation as the focus. In AMRs, causation is represented with the frame `cause-01`. Figure 8a shows the AMR for Example (9).

(9) The girl was not sad because the boy left.

A first attempt to represent the CAUSE PPI "The girl was sad because of something but not because the boy left" was to embed `contrast-01` beneath the `cause-01` node in the same way as "because X and because Y" would be represented. However, the resulting graph in Figure 8c represents "The girl was sad because something happened but the boy did not leave" which implies that the boy did not leave. In contrast, the leaving of the boy is a factual event in the original sentence and all PPIs need to express that the leaving of the boy indeed happened. Figure 8b shows the correct AMR for the CAUSE PPI where the negation is attached to `cause-01`. In the case of fine-grained PPIs the proposed methodology also leads to problems with respect to the factuality of the "boy-leaving-event". Therefore, further research on investigating the interactions of negation, scope and factuality is needed for representing PPIs in AMRs.

5 Conclusion & Future Work

The way in which negation is represented in AMRs does not allow for accurate interpretation of scope or the ability to disambiguate between the potential foci of the negation. The goal of the current work was therefore, to investigate how the PPIs of negated statements can be represented in AMRs in order to explicitly represent the focus of negation

and whether these representations can be generated in a systematical way. We proposed a logically motivated representation for the PPIs where the negation is attached directly to the root node of the focus. As future work, we plan to investigate our proposed approach on sentences from existing AMR corpora.

Working through example negated statements has shown that representing positive meaning adequately in AMRs is a complex task especially when quantifier and scope interactions are involved. The results support the need for an extension of AMR that is able to handle quantifiers and scope phenomena in general, and negation in particular, in a systematic and explicit way that is suitable for inference tasks. Representing focus explicitly in AMRs does not solve the problem that modelling negation simply as a relation is inappropriate for drawing inferences. As a next step, then, we hope to combine our proposed approach with an existing approach that extends AMR's scope capacities, specifically by investigating the use of Pustejovsky et al. (2019)'s approach for determining which nodes to consider as potential foci and by translating the `contrast-01` concept and the existential placeholder subgraphs into FOL.

References


- Lasha Abzianidze and Johan Bos. 2019. [Thirty musts for meaning banking](#). In *Proceedings of the First International Workshop on Designing Meaning Representations*, pages 15–27, Florence, Italy. Association for Computational Linguistics.
- Begoña Altuna, Anne-Lyse Minard, and Manuela Speranza. 2017. [The scope and focus of negation: A complete annotation framework for Italian](#). In *Proceedings of the Workshop Computational Semantics Beyond Events and Roles*, pages 34–42, Valencia, Spain. Association for Computational Linguistics.

- Pranav Anand and Craig Martell. 2012. [Annotating the focus of negation in terms of questions under discussion](#). In *Proceedings of the Workshop on Extra-Propositional Aspects of Meaning in Computational Linguistics*, pages 65–69, Jeju, Republic of Korea. Association for Computational Linguistics.
- Laura Banarescu, Claire Bonial, Shu Cai, Madalina Georgescu, Kira Griffitt, Ulf Hermjakob, Kevin Knight, Philipp Koehn, Martha Palmer, and Nathan Schneider. 2013. [Abstract Meaning Representation for sembanking](#). In *Proceedings of the 7th Linguistic Annotation Workshop and Interoperability with Discourse*, pages 178–186, Sofia, Bulgaria. Association for Computational Linguistics.
- Eduardo Blanco and Dan Moldovan. 2011. Some issues on detecting negation from text. In *Twenty-Fourth International FLAIRS Conference*. Citeseer.
- Eduardo Blanco and Dan Moldovan. 2012. [Fine-grained focus for pinpointing positive implicit meaning from negated statements](#). In *Proceedings of the 2012 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 456–465, Montréal, Canada. Association for Computational Linguistics.
- Eduardo Blanco and Dan Moldovan. 2014. [Retrieving implicit positive meaning from negated statements](#). *Natural Language Engineering*, 20(4):501–535.
- Eduardo Blanco and Zahra Sarabi. 2016. [Automatic generation and scoring of positive interpretations from negated statements](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1431–1441, San Diego, California. Association for Computational Linguistics.
- Johan Bos. 2016. [Expressive Power of Abstract Meaning Representations](#). *Computational Linguistics*, 42(3):527–535.
- Johan Bos. 2020. [Separating argument structure from logical structure in AMR](#). In *Proceedings of the Second International Workshop on Designing Meaning Representations*, pages 13–20, Barcelona Spain (online). Association for Computational Linguistics.
- Ariel Cohen. 2001. [On the Generic Use of Indefinite Singulars](#). *Journal of Semantics*, 18(3):183–209.
- Rodney Huddleston and Geoffrey K Pullum. 2002. The cambridge grammar of english. *Language*. Cambridge: Cambridge University Press, 1.
- Aakanksha Naik, Abhilasha Ravichander, Norman Sadeh, Carolyn Rose, and Graham Neubig. 2018. [Stress test evaluation for natural language inference](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 2340–2353, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- James Pustejovsky, Ken Lai, and Nianwen Xue. 2019. [Modeling quantification and scope in Abstract Meaning Representations](#). In *Proceedings of the First International Workshop on Designing Meaning Representations*, pages 28–33, Florence, Italy. Association for Computational Linguistics.
- Zahra Sarabi and Eduardo Blanco. 2016. [Understanding negation in positive terms using syntactic dependencies](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1108–1118, Austin, Texas. Association for Computational Linguistics.
- Zahra Sarabi and Eduardo Blanco. 2017. [If no media were allowed inside the venue, was anybody allowed?](#) In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pages 860–869, Valencia, Spain. Association for Computational Linguistics.


A Appendix

0	Start with original AMR of negated statement
1	a. Select the focus and the corresponding f-subgraph b. Select the sr-subgraph to which the focus belongs
2	Detach the sr-subgraph from the original AMR and copy it
3	Move negation to the root of the f-subgraph within the detached sr-subgraph
4	Attach the resulting negated subgraph to :ARG2 of <code>contrast-01</code>
5	a. Create the appropriate subgraph for the existential placeholder b. Detach the f-subgraph from the copied sr-subgraph c. Attach the subgraph for the existential placeholder at empty position
6	Attach the resulting subgraph to :ARG1 of <code>contrast-01</code> a. Fix re-entrancies for arguments shared between :ARG1 and :ARG2 b. Pull-out shared non-core relations
7	Attach the contrast subgraph to the original AMR at empty original position of the sr-subgraph


Table 3: The steps of the proposed methodology to transform the AMR of an negated statement into the AMR of one of its PPIs.



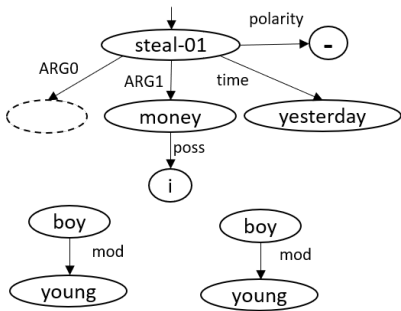
(0) Original AMR of negated statement



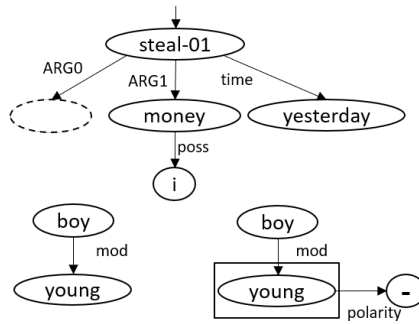
(1a) Select *young* as focus and the corresponding f-subgraph




(1b) Select the sr-subgraph to which the focus belongs




(2) Detach sr-subgraph and copy it




(3) Move negation to root of f-subgraph within the detached sr-subgraph




(4) Attach resulting negated subgraph to :ARG2 of contrast-01




(5a) Subgraph for placeholder *some*



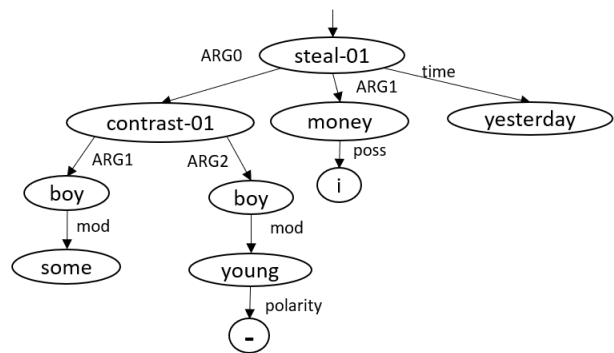
(5b) Detach f-subgraph from the copied sr-subgraph



(5c) Attach the subgraph for the placeholder at empty position



(6) Attach the resulting subgraph to :ARG1 of contrast-01



(7) Attach contrast subgraph to original AMR at empty position of the sr-subgraph

Figure 9: Illustration of generating the AMR for the MOD PPI "Some boy stole my money yesterday but not the young boy" (7) from the AMR of the negated statement "The young boy did not steal my money yesterday" (0).

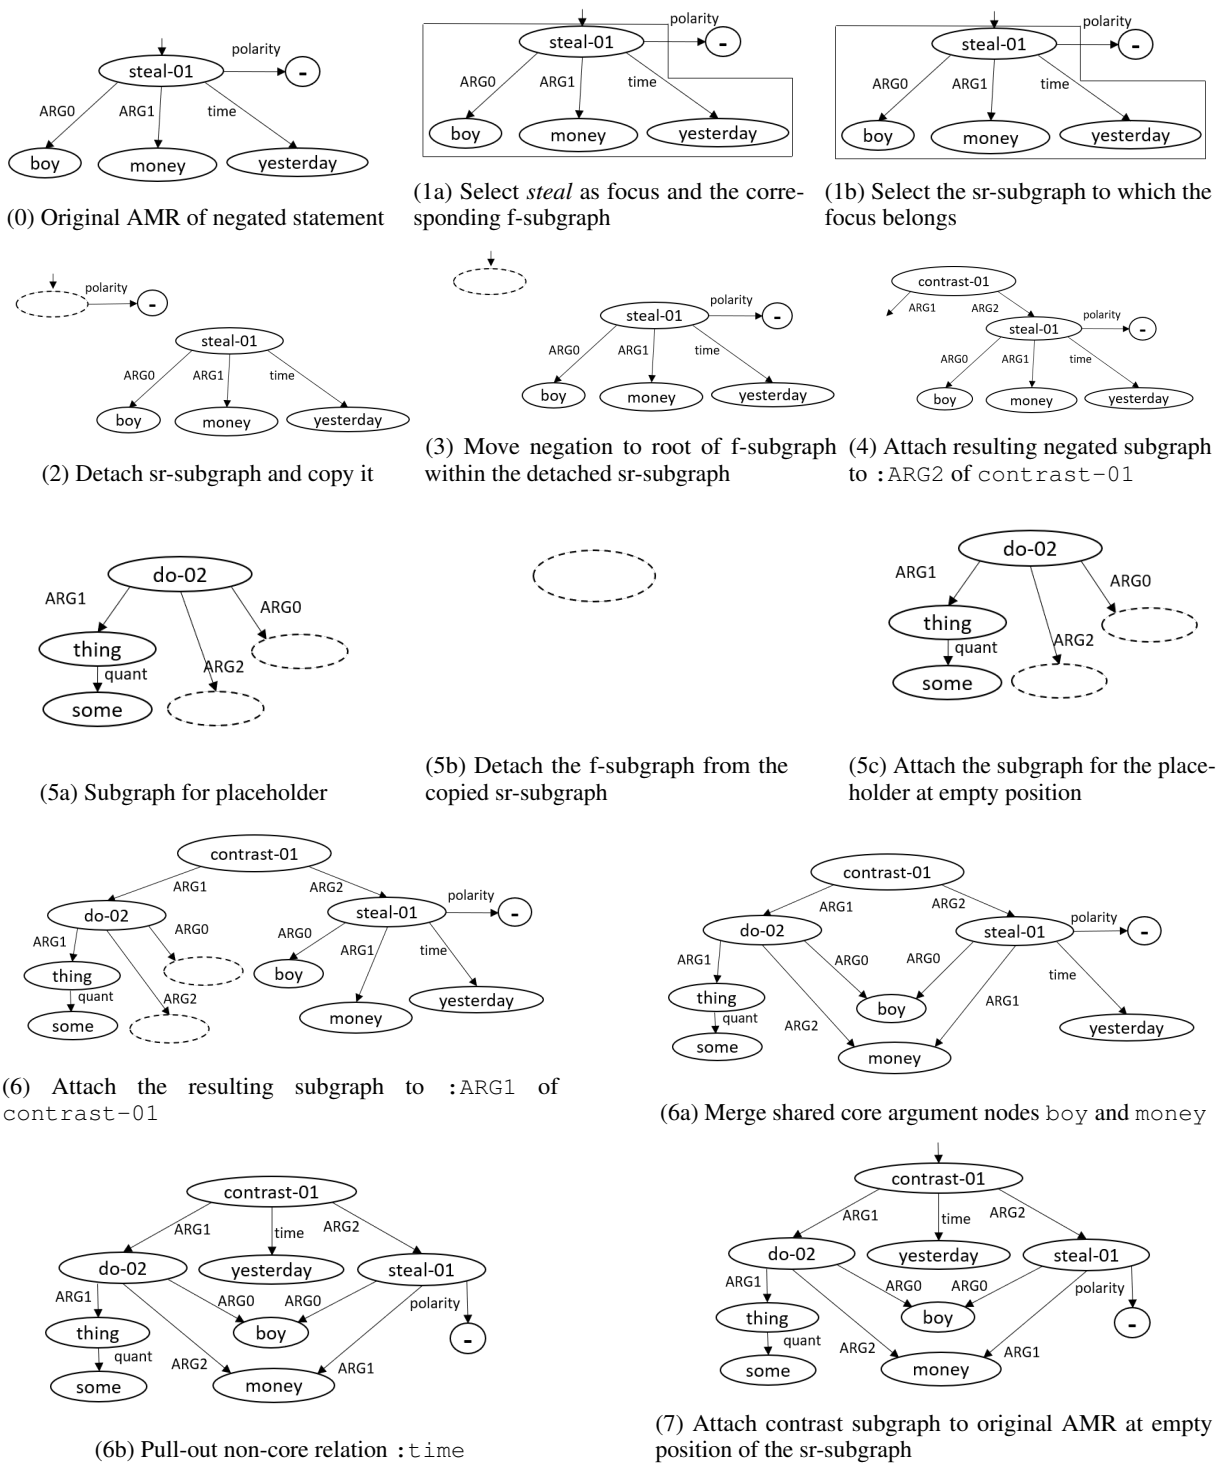


Figure 10: Illustration of generating the AMR for the ROOT PPI "The boy did something with the money yesterday but not steal" (7) from the AMR of the negated statement "The boy did not steal the money yesterday" (0). In the special case of selecting the root / the predicate frame as focus here, the f-subgraph (as well as the sr-subgraph) consists of the predicate and all its arguments except for the negation. For reasons of space the copied sr-subgraph is not shown in step (2) and (3).