

# Guidelines for Construction of Natural Language Inference Problems

## 1 Spatial reasoning

A sentence with a spatial preposition (like *across*, *in*, *towards*) can have spatial entailments in virtue of that preposition. (1) is an example of a spatial entailment (notation:  $\Rightarrow$ ) involving the preposition *from* (and *in* in the entailment).

- (1) John saw Mary *from* the garden  $\Rightarrow$  John was *in* the garden

The sentence preceding (following) the inference relation sign (in case of (1), this is  $\Rightarrow$ ) is called *premise* (*hypothesis*, respectively). (2) is an entailment based on two premises, involving the preposition *south of* in interaction with *between*

- (2) Antwerp is south of Amsterdam & Dordrecht is between Amsterdam and Antwerp  $\Rightarrow$  Amsterdam is north of Dordrecht

The opposite of an entailment is a contradiction (3), (4) (notation:  $\not\Rightarrow$ ):

- (3) John saw Mary *from* the garden  $\not\Rightarrow$  Mary was *in* the garden  
(4) Antwerp is south of Amsterdam & Dordrecht is between Amsterdam and Antwerp  $\not\Rightarrow$  Dordrecht is north of Amsterdam

The relation between sentences with prepositions can also be neutral (5), (6) (notation: #):

- (5) John saw Mary *in* the garden # John was *in* the garden  
(6) Nijmegen is west of Dordrecht & Dordrecht is between Antwerp and Amsterdam # Nijmegen is between Antwerp and Amsterdam

The inference examples like (1)-(6) are called *natural language inference* (NLI) problems.

## 2 Towards a database of spatial reasoning patterns

The goal is to create a database of NLI problems, with the labels  $\Rightarrow$  (entailment),  $\not\Rightarrow$  (contradiction), and # (neutral), involving spatial prepositions, and possibly other spatial expressions. This document provides an initial set, most of them drawn from Nam [1995] and Zwarts and Winter [2000]. The task is to create more examples with different labels and variations. Ultimately, that database should present a challenge for machine learning approaches to spatial reasoning based on textual information. This document gives an overview (not meant to be exhaustive) of inference patterns with prepositions.

But before diving into the examples of spatial reasoning, we would like to say more about the link between spatial reasoning and machine learning. Nowadays, it is often claimed that machines can reason with natural language. While this might seem true taking into account only the scores obtained by state-of-the-art NLI models on NLI benchmarks, the truth is that often NLI models get high scores by learning spurious correlations between the sentences of NLI problems. These correlations are supported by biases unintentionally introduced in the benchmarks during data collection. For example, here is a short list of reported spurious correlations learned by NLI models after trained on a large set of three-way classified NLI problems:

1. NLI problems with negation words (e.g., *no*, *not*, and *never*) are often classified as contradiction.
2. If a hypothesis sentence contains general concepts, like *animal*, *person*, and *something*, there is high chance that the problem is classified as entailment.
3. Word order is not always important for NLI models as their prediction often doesn't change when words are shuffled in the sentence of NLI problems.
4. If there is a little word overlap between a premise and a hypothesis, the problem is most likely classified as neutral.
5. When a hypothesis is much longer than its premise, for NLI models it is a good indicator to predict the neutral relation.
6. NLI models usually predict the entailment relation when a hypothesis is a proper part of the premise.

We think that it is useful for data creators to know what kind of biases can be exploited by machine learning models, so that they will not introduce any biases in new data or, put differently, introduced biases will be balanced in such a way that makes it difficult for machines to take advantage of them.

### 3 Spatial prepositions

We assume that a preposition describes a relation between two entities, the *trajector* and the *landmark*. In a basic predication like *John was in the garden* the subject gives the trajector (John) and the object gives the landmark (the garden). There is a fundamental distinction between *place* prepositions and *path* prepositions. A place preposition, like *in*, helps to describe a static, locative relation between the trajector and the landmark, while a path preposition like *into* helps to describe a more complex relation between the trajector and the landmark. The trajector can *move* along the path (7), it can *extend* along the path (8), or it can have another sort of direction along the path (9).

- (7) John went into the garden.
- (8) The path led into the garden.
- (9) John looked into the box.

## 4 Spatial entailment patterns

### 4.1 Argument orientation

An argument orientation entailment identifies the trajector of a prepositional phrase (PP). (10) shows that Mary is the trajector of the place PP *in the garden* and (11) shows that *the ball* is the trajector of the path PP *into the box*.

(10) John saw Mary in the garden  $\Rightarrow$  Mary was in the garden

(11) John threw the ball into the box  $\Rightarrow$  The ball went into the box

Small variations lead to neutral (12) and contradictory (13) versions. The following examples are obtained by replacing a single word in (10), where the new words are underlined:

(12) John saw Mary in the garden # John was in the garden

(13) John saw Mary from the garden  $\not\Rightarrow$  Mary was in the garden

### 4.2 Path/place relations

In addition to argument orientation, a path/place entailment shows the place on which a path is based. (14) shows that *across* involves being on the landmark at some point of the path.

(14) John walked across the street  $\Rightarrow$  John was on the street

Some entailments involve more clearly both argument orientation and a path/place relation:

(15) The man dragged the barrel across the street  $\Rightarrow$  The barrel was on the street

A neutral pattern is found in (16) and a contradictory pattern in (17)

(16) John walked around the field # John was on the field

(17) John saw Mary from the garden  $\not\Rightarrow$  Mary was in the garden

### 4.3 Symmetry and asymmetry

There are different ways to demonstrate the difference between ‘symmetric’ and ‘non-symmetric’ paths. Symmetric paths are expressed by prepositions like *across*, *around*, *through* and asymmetric paths by prepositions like *from*, *into*, *out of*, *towards*.

#### 4.3.1 With *back*

One entailment type uses the adverb *back*:

(18) The boy walked across the street and went back immediately  $\Rightarrow$  The boy walked across the street twice

(19) The boy walked across the street and went back immediately  $\not\Rightarrow$  The boy walked across the street once

(20) The boy walked into the room and went back immediately  $\Rightarrow$  The boy walked out of the room

- (21) The boy walked into the room and went back immediately # The boy walked into the room twice

#### 4.3.2 With *from*

The symmetry of a path can also be demonstrated when *from* introduces an additional point of view.

- (22) The house is across the street from the bus stop  $\Rightarrow$  The bus stop is across the street from the house
- (23) The house is across the street from the bus stop  $\not\Rightarrow$  The street is across the bus stop from the house
- (24) The house is across the street from the bus stop # The bus stop is behind the house

#### 4.3.3 With *opposite*

We also find the following ‘symmetry’ entailment, with the expression *opposite sides*:

- (25) John saw Mary through the window  $\Rightarrow$  John and Mary were on opposite sides of the window
- (26) John showed the picture to Mary through the window # The picture and John were on the opposite sides of the window
- (27) John hugged Mary behind the window  $\not\Rightarrow$  John and Mary were on the opposite sides of the window

#### 4.3.4 With *and then*

Asymmetric paths allow ‘chaining’, when the endpoint of one path is identified with the starting point of another path.

- (28) John flew from Los Angeles to San Diego, and then to Las Vegas  $\Rightarrow$  John flew from Los Angeles to Las Vegas
- (29) John flew from Los Angeles to San Diego, and then he drove to Las Vegas  $\not\Rightarrow$  John flew from Los Angeles to Las Vegas
- (30) John drove towards Los Angeles and then to Las Vegas # John drove from Los Angeles to Las Vegas

### 4.4 Part-whole entailments

If we know that Paris is part of (or in) France, then we can make the entailment in (31) and (32).

- (31) The house is in Paris  $\Rightarrow$  The house is in France
- (32) The house is outside France  $\Rightarrow$  The house is outside Paris

(33) and (34) are neutral and contradictory examples, respectively.

- (33) The house is outside Paris # The house is outside France
- (34) The house is outside France  $\not\Rightarrow$  The house is in Paris

Roughly speaking, the question here is whether a spatial relation persists when we expand or shrink the landmark. This is also relevant for path prepositions (35).

- (35) Mary traveled across France # Mary traveled across Europe

#### 4.5 Vector monotonicity

The preposition *between* can be used to create entailments in combination with a few other prepositions:

- (36) The tree is behind the house & Mary is between the tree and the house  $\Rightarrow$  Mary is behind the house

- (37) Mary is behind the house & Mary is between the tree and the house  $\Rightarrow$  The tree is behind the house

Intuitively, here the question is whether a spatial relation persists when we get closer to the landmark or further away, in the same direction, like the *behind* direction. Examples of non-entailment variants are (38) and (39).

- (38) Mary is in front of the house & Mary is between the tree and the house  $\not\Rightarrow$  The tree is behind the house

- (39) Mary is near the house & Mary is between the tree and the house # The tree is near the house

#### 4.6 Point of view

The *from* point of view PPs that we mentioned in relation to symmetric paths also play a role with places that require an deictic frame of reference, like *to the left of* and *to the right of*.

- (40) Bill is to the right of the tree from John  $\Rightarrow$  John is to the left of the tree from Bill

Bill and John are on opposite sides of the tree, so they have opposite points of view. Notice what happens when John is looking at a trashcan standing right next to a tree:

- (41) The trashcan is to the right of the tree from John  $\not\Rightarrow$  The tree is to the right of the trashcan from John

## 5 The task

The task is to expand the database of examples (labeled as entailments, contradictions, or neutral cases) with new examples, with different labels, so that we show how the variation of prepositions, verbs, and nouns affects spatial reasoning patterns. At the moment we have no systematic 'procedure' for that. We think that it is best to follow the following steps:

1. Take an entailment ( $\Rightarrow$ ).
2. Try to form an understanding of what it is that makes this entailment work.

3. Identify the lexical and structural elements that contribute to the entailment: not only the prepositions, but also the verb involved, the nouns, argument structure, temporal structure, ....
4. Create an example that is contradictory ( $\not\Rightarrow$ ) or neutral ( $\#$ ).

Here is an example.

1. Take the entailment: **The cat is between the house and the fence & The house is between the fence and the tree  $\Rightarrow$  The cat is between the fence and the tree**
2. We could say that this entailment is about an ordering of four things in a row: Tree - House - Cat - Fence (or Fence - Cat - House - Tree).
3. We could also identify the schema of the entailment:  **$C$  is between  $H$  and  $F$  &  $H$  is between  $F$  and  $T \Rightarrow C$  is between the  $F$  and  $T$**
4. The entailment depends on the preposition *between* and different configurations with four NPs. If we keep the preposition constant, we can vary those configurations with the same NPs.
5. We can get a contradiction when we change the conclusion to no longer match the premises:  
**The cat is between the house and the fence & The house is between the fence and the tree  $\not\Rightarrow$  The tree is between the house and the cat**
6. We get a neutral case when we set up a situation where the cat and the tree end up being 'unordered' with respect to each other:  
**The cat is between the house and the fence & The tree is between the fence and the house  $\#$  The tree is between the house and the cat**

**No redundant problems** No need to create redundant inference problems that fall in the same schema. For instance, one can create a new but redundant entailment problem by replacing *cat*, *fence*, *house*, and *tree* with *black cat*, *car*, *traffic light*, and *bike* respectively, in the running example.

**Prevent biases in new problems** [section 2](#) lists several structural and word-level biases that can be exploited by machine learning models to mimic reasoning. Try not to introduce such biases. If you think that you have introduced some biases, try to balance them by introducing a bias that neutralizes the previous ones to some extent. Most probably you will still unintentionally introduce some bias. That is okay. But it is very important that the extent of a bias is not large. In other words, if a machine learning model picks up the bias, it shouldn't help it to correctly classify a large amount of NLI problems.

**World knowledge** Some inference problems might require world knowledge. For example, the inference problems in [subsection 4.4](#) rely on the knowledge that *Paris is in France*. On the other hand, we don't consider semantic properties of the cardinal directions (*north*, *east*, *south*, and *west*) and body relative directions (*left*, *right*, *up*, *down*, etc.) as world knowledge. In case an inference problem requires world knowledge, include the world knowledge as an additional premise and flag it explicitly.

At the moment we have no sharp ideas about how to vary the examples. A lot of creativity is needed here, playing around with the examples, finding out what the effect is of different types of changes.

Be creative about types of reasoning when creating a new NLI problems. We know it is tempting to create examples that use spatial reasoning similar to the ones demonstrated by the provided examples, but try to think outside the box. Don't allow the provided examples to frame and limit your creativity, but think of them as the tip of the "spatial reasoning" iceberg. Try not to write down all your inference problems in one sitting as it might negatively affect your creativity. Diversity of the reasoning required by your inference problems is more valued than the amount of the problems.

## References

Seungho Nam. *The Semantics of Locative Prepositional Phrases in English*. PhD thesis, University of California, 1995.

Joost Zwarts and Yoad Winter. Vector space semantics: A model-theoretic analysis of locative prepositions. *Journal of Logic, Language and Information*, 9:169–211, 2000.