

Composing Conversational Negation

Razin A. Shaikh*

Mathematical Institute,
University of Oxford

razin.shaikh@maths.ox.ac.uk

Lia Yeh*

Quantum Group,
University of Oxford

lia.yeh@cs.ox.ac.uk

Benjamin Rodatz*

Computer Science,
University of Oxford

benjamin.rodatz@cs.ox.ac.uk

Bob Coecke

Cambridge Quantum

bob.coecke@cambridgequantum.com

Negation in natural language does not follow Boolean logic and is therefore inherently difficult to model. In particular, a broader understanding of what is being negated must be taken into account. In previous work, we proposed a framework for the negation of words that accounts for ‘worldly context’. This paper extends that proposal to now account for the compositional structure inherent in language within the DisCoCirc framework. We compose the negations of single words to capture the negation of sentences. We also describe how to model the negation of words whose meanings evolve in the text.

1 Introduction

Negation in language is a complicated operation. Differing views of negation in language are a recurring subject of debate amongst linguists, epistemologists, and psychologists. One view maintains that negation in language conveys denial of a proposition [11]. Another view on negation in language asserts that it is the collective notion of plausible alternatives [29]; this view traces its origins to as far back as Plato’s view of not-being-X as *otherness*-than-X [24]. An explanation compatible with both views is that there can be different stages at which the negation is interpreted, for instance, initially denying information, and later searching for alternatives [32]. This search for alternatives differentiates negation in conversation from simple logical negation.

Consider the sentences:

- a) This is not a hamster; this is a guinea pig.
- b) This is not a hamster; this is a planet.

Both sentences are grammatically and logically correct. However, most users of the English language will agree that there is something wrong with sentence b). Unlike sentence a), which seems reasonable without context, sentence b) must undergo a highly unusual ‘contextual pressure’ [22] to be believable—imagine a sci-fi flick about hamster-sized planets. The plausibility of different alternatives to a negated word naturally has a grading [29, 22]. Our previous work modelled operational conversational negation of a word, and experimentally validated that it positively correlates with human judgment of this grading [33].

This paper will show how to perform the conversational negation of sentences using the conversational negation of words. This is analogous to how Coecke et al. [8] developed the compositional categorical framework DisCoCat to obtain the meaning of a sentence from the meaning of its words. To extend from negation of words to negation of a sentence, a new

*Equal contribution

challenge appears: ambiguity arises not only from the meaning of a negated word but also from which word(s) in the sentence the negation is principally applied to. As an example, take the sentence “Johnny did not travel to Manchester by train” [30]. Envision that Johnny is given emphasis—the natural conclusion is that someone else, instead of Johnny, went to Manchester by train. Correspondingly, if the emphasized word were **Manchester** or **train**, then the respective conclusions would be that Johnny went elsewhere or Johnny took another mode of transportation. Therefore, we note that the conversational negation of this sentence is arrived at from the conversational negation of its constituent words, i.e. constituent negation [14]. We also see that the grammatical structure is unaltered between the non-negated and negated forms of the sentence. This is in line with other attempts in deciphering the meaning of negation in natural language [16, pg. 104].

Another challenge we tackle in this paper is performing conversational negation of words or entities whose meanings are not fixed; instead, the meanings evolve as we obtain more information in the text. For instance, in the text “Waffle is a dog. Waffle is fluffy. Waffle likes to play fetch.”, the meaning of the entity “Waffle” evolves with each sentence; i.e. we first learn that “Waffle” refers to a dog, then that it is fluffy, and finally, that it likes to play fetch. We want to perform the conversational negation, i.e. model alternatives, of the entity “Waffle” whose meaning updates with time. The negation of an evolving entity has ambiguity similar to the negation of sentences: we do not know to which information/property of the evolving entity the negation is applied. Coecke [5] extended the DisCoCat framework to DisCoCirc, which models evolving meanings by allowing us to compose sentences in the text. We will show that in the DisCoCirc formalism, the negation of evolving meanings is, in fact, the same thing as the negation of sentences. Hence, we will use the same framework of conversational negation that we derived for sentences.


Structure In Section 2, we present the compositional DisCoCirc framework and the meaning category of positive operators and completely positive maps. In Section 3, we summarize the framework for conversational negation of words that we originally proposed in [33]. In Section 4, we go from conversational negation of words to conversational negation of sentences. In Section 5, we model the conversational negation of entities whose meanings evolve in the text. Finally, in Section 6, we discuss the implications of the ideas presented in this paper and outline directions for future work.

2 Compositional language meaning


2.1 DisCoCirc

The DisCoCat framework [8] combines grammar (cf. categorial grammar [23]) and meanings (cf. vector embeddings in machine learning) within one compositional framework that enables one to compute the meaning of a sentence from the meanings of its words. To achieve this, it exploits the common compact closed categorical structure, be it of vectors and linear maps, or of positive operators and completely positive maps [2, 31]. The DisCoCirc framework [5] improved on DisCoCat, (1) by enabling one to compose sentences into larger text, just as gates are composed in circuits; (2) by allowing meanings to evolve as that text evolves; (3) by specifying the sentence type as the tensored spaces of those entities that evolve. For our purposes, a

DisCoCirc diagram has two key ingredients: (1) meaning states; (2) updates [7]:



For example, here we have the noun meanings *Alice* and *Bob*, which initially are separate, being updated with the verb meaning *love*. Alternatively, we can have noun-wires with open input, which we then update to being *Alice* and *Bob* respectively, and then update by *love*:



These open input wires allow for pre-composition with other sentences.

2.2 The meaning category $\mathbf{CPM}(\mathbf{FHilb})$

While the DisCoCirc framework allows for various encodings of meaning, for this paper, we work with the category $\mathbf{CPM}(\mathbf{FHilb})$ of positive operators and completely positive maps, as done in [26, 7, 31]. Positive operators are complex matrices, which are equal to their own conjugate transpose (Hermitian) and have non-negative eigenvalues (positive semidefinite). Completely positive maps are linear maps from positive operators to positive operators. The compact closed category $\mathbf{CPM}(\mathbf{FHilb})$ can be seen as an extension of the category of finite-dimensional Hilbert spaces \mathbf{FHilb} . It can be obtained from \mathbf{FHilb} via the CPM construction originally introduced by Selinger [35]. In this construction, for any given unit vector $|v\rangle$ of a finite-dimensional Hilbert space, we can obtain a pure state positive operator by taking the outer product $|v\rangle\langle v|$. All other positive operators can be obtained as a linear combination of pure states.

In contrast to vectors, which have no inherent ordering structure [1], positive operators can be viewed as an extension of vector spaces that allows for encoding lexical entailment structure such as proposed in [2, 25]. We use these entailment measures to capture hyponymy: a word w_1 is a hyponym of w_2 if w_1 is a type of w_2 ; then, w_2 is a hypernym of w_1 . For example, *dog* is a hyponym of *animal*, and *animal* is a hypernym of *dog*. These entailment measures are often graded and take values between 0 (no entailment) and 1 (full entailment).

Additionally, positive operators can be used to encode ambiguity—words having multiple meanings—via mixing [31, 27], i.e. taking weighted sums over the different meanings. This ambiguity can later be disambiguated through additional meaning updates [7, 27]. We will use this property of positive operators to encode the ambiguity of negation in Section 4. For analysis and model performance for encoding and resolving ambiguity in compositional distributional semantics, we refer the reader to the extensive literature [3, 20, 17, 31, 7, 27].

3 Conversational negation of words

This section summarizes the operation for conversational negation of words as we proposed in [33]. As pointed out in Section 1 with the example of hamster-sized planets, negation in conversation not only denies information—a planet is indeed not a hamster—but additionally utilizes the listener’s understanding of the world to weigh possible alternatives. This builds on the assumption that plausible alternatives to a word should appear in a similar context [30]. To present the framework for conversational negation, we will first present logical negation operations and an encoding of worldly knowledge.

3.1 Logical negation

Logical negation (denoted by \neg) should fulfill certain properties such as the double negative ($\neg(\neg p) = p$) and the contrapositive ($p \sqsubseteq q \Leftrightarrow \neg q \sqsubseteq \neg p$). Lewis [26] proposes the operation $\neg X := \mathbb{I} - X$, mapping in the case of projectors to the orthogonal subspace, as Widdows and Peters did for vectors [36]. In [33], we proposed another logical negation based on generalizing the matrix inverse, that satisfies the contrapositive condition for the (graded) Löwner order [2].

In some sense, the result of the logical negation of a word is akin to a mixture of everything that is not that word. Despite this aligning with the set-theoretic notion of complement sets, this is unlike how humans perceive negation. Indeed, in our prior experiments, we found that alternatives elicited by both of the proposed logical negations have a negative correlation with human intuition [33]. We remedied this by amending logical negation with the worldly context of the negated word to achieve a positive correlation.

3.2 Worldly context

Worldly context is another primary ingredient of conversational negation. It encodes the intuitive understanding of the world most readers possess, by capturing the context in which a word tends to appear. Thus, as it encodes the space of possible alternatives to a word, worldly context can be utilized to weigh the results of the logical negation.

To build worldly context for a given word, we proposed utilizing entailment hierarchies such as the human-curated WordNet [12] or the unsupervised Hearst patterns [13]. These hierarchies give us entailment relations such as displayed in Figure 1, where each directed edge represents




Figure 1: Example of hyponymy structure as can be found in entailment hierarchies

an entailment. We thus get relations such as every hamster being a rodent, every rodent being an animal, and all animals being entities. When negating hamster, we are most likely to talk about other rodents such as guinea pigs. We are less likely to talk about other animals such as dogs and yet less likely to talk about other entities such as planets.

Building on this idea, we proposed to construct the worldly context of a word by considering its hypernym paths and taking a weighted sum over all hypernyms. Hence, for a word w with


hypernyms h_1, \dots, h_n ordered from closest to furthest, we define the worldly context \mathbf{wc}_w as:

$$\llbracket \mathbf{wc}_w \rrbracket := \sum_i p_i \llbracket h_i \rrbracket \quad (3)$$

where $p_i \geq p_{i+1}$ for all i . We denote the positive operator encoding the meaning of a word with double brackets.

3.3 Framework for conversational negation of words

In [33], conversational negation of words, written as the operation CN_{word} , is defined as:



This framework can be interpreted as the following three steps.

1. Calculate the logical negation $\neg(\llbracket w \rrbracket)$.
2. Compute the worldly context $\llbracket \mathbf{wc}_w \rrbracket$.
3. Update the meaning of $\neg(\llbracket w \rrbracket)$ by composing with $\llbracket \mathbf{wc}_w \rrbracket$ to obtain $\neg(\llbracket w \rrbracket) \curlywedge \llbracket \mathbf{wc}_w \rrbracket$.

This framework is flexible to the choice of logical negation, worldly context generation and composition operation \curlywedge . In [33], we studied and compared the performance of various choices of operations. While we only experimentally validated our negation operation on nouns, the same operation is applicable to adjectives and verbs.

4 Conversational negation of strings of words

In the distributional approach to natural language processing, a commonly used model is the bag-of-words that disregards any grammar and treats words as a structureless bag. Coecke et al. [8] proposed the DisCoCat model, which combines grammar and distributional word meanings in a categorical framework, allowing us to compose the meaning of words to get the meaning of a sentence. In this section, we do the same for conversational negation by introducing the conversational negation of sentences that is calculated using the conversational negation of words.

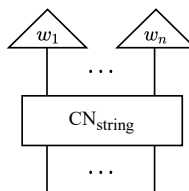
4.1 Modelling conversational negation

As pointed out by Oaksford and Stenning [30], the negation of more complex structures consisting of multiple words may be interpreted as the negation of a subset of the constituents. For example, a sentence such as “Bob did not drive to Oxford by car” could be interpreted as:

- a) Bob did not drive to Oxford by car - Alice did
- b) Bob did not drive to Oxford by car - He carpooled
- c) Bob did not drive to Oxford by car - He drove to London
- d) Bob did not drive to Oxford by car - He drove a van
- e) Bob did not drive to Oxford by car - Alice carpooled to Oxford

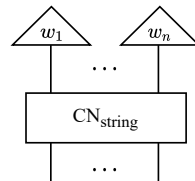
where the underlines indicate which words are being negated. The last example is one of many possible cases that negate multiple constituents. While some of these alternatives might immediately seem more plausible to the reader, the correct choice is inherently dependent on the context.

Based on this interpretation, that a negation of multiple words is negating a subset of the constituents, we extend our conversational negation framework to a string of words by utilizing conversational negation of individual words (see Section 3). As the correct interpretation of which words to negate may not usually be obvious, we create a mixture of all possible interpretations. Therefore, the negation of a string of n words $w_1 \otimes w_2 \otimes \dots \otimes w_n$ is a weighted mixture of all the interpretations where only one word is negated, all the interpretations where two words are negated, and so on. We have:



$$\begin{aligned}
 & \sum_{i=1}^n p_{\{w_i\}} (w_1 \otimes \dots \otimes \text{CN}_{\text{word}}(w_i) \otimes \dots \otimes w_n) + \\
 & := \sum_{i=1}^n \sum_{j=i+1}^n p_{\{w_i, w_j\}} (w_1 \otimes \dots \otimes \text{CN}_{\text{word}}(w_i) \otimes \dots \otimes \text{CN}_{\text{word}}(w_j) \otimes \dots \otimes w_n) \quad (5) \\
 & + \dots
 \end{aligned}$$

where in the overall mixture representing the negation, each interpretation has some weight. Formally, for $S = \{w_1, \dots, w_n\}$ and non-empty $S' \subseteq S$ which we call the **negation set**, we get:

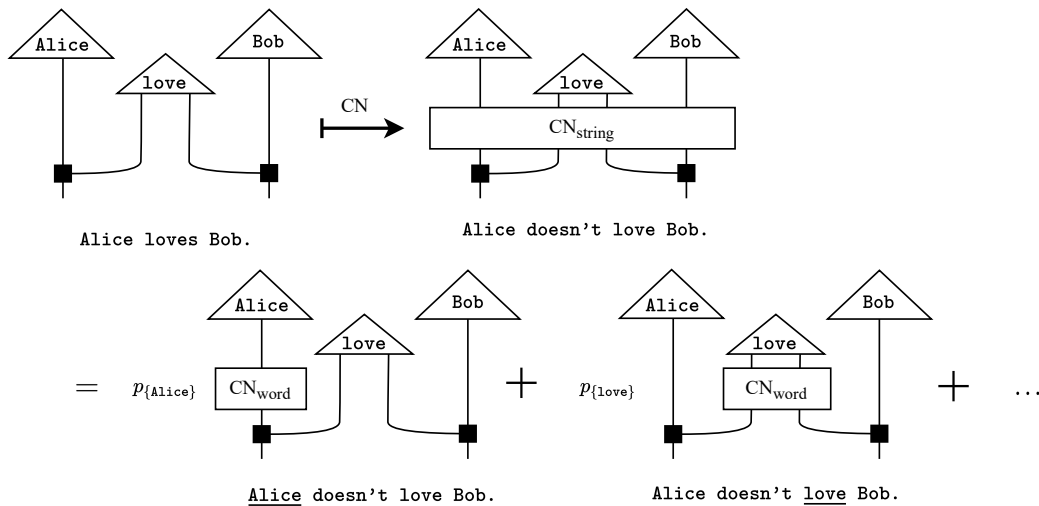


$$\sum_{S' \in \mathcal{P}(S) \setminus \{\emptyset\}} p_{S'} \bigotimes_{i=1}^n \begin{cases} w_i & \text{if } w_i \notin S' \\ \text{CN}_{\text{word}}(w_i) & \text{if } w_i \in S' \end{cases} \quad (6)$$

To apply this negation to sentences, we adopt the view that sentences are processes updating wires as presented by Coecke in [5]. These processes are built from a combination of meaning states, interacting via updates. We propose the negation of a sentence to be viewed as the same set of meaning states, first updated by the conversational negation of the words before updating the wires as if the negation was absent. This is in line with Kamp and Reyle [16, pg. 104], who also treat the grammatical structure of negated sentences as if they were not negated. Hence, we represent the conversational negation as a function called CN that maps the circuit of a sentence to a circuit where the sentence is pre-composed with $\text{CN}_{\text{string}}$.

For example, applying CN to the sentence “Alice loves Bob”, we get the circuit for the

sentence “Alice doesn’t love Bob”:



Here, we sum over all possible non-empty subsets of {Alice, love, Bob}, each of which is weighted by the appropriate scalar.

4.2 Deriving the weights

The main challenge for the conversational negation of a string is deriving the weights for the different interpretations of the negation. The choice of correct interpretation, and therefore the weights, is dependent on context. Context can be derived from many sources, such as the person who is speaking and their intentions. In spoken language, intonation could clarify the speaker’s intent by emphasizing the words meant to be negated.

Another source of context is the grammatical structure of the negated sentence itself. Given the earlier example “Bob did not drive to Oxford by car,” which mentions the mode of transport explicitly, intuitively the focus of negation is on this detail. If the speaker solely wanted to negate Bob’s destination, the sentence “Bob did not drive to Oxford” would be sufficient, not requiring any additional detail. The other example, “Alice does not love Bob,” is more ambiguous as the grammatical structure does not indicate the target of the negation.

Context from surrounding sentences can further determine the sensibility of interpretations to a negation. Consecutive sentences should have meanings consistent with each other and the reader’s general understanding of the world. The statement Bob did not drive to Oxford by car – He drove to London is reasonable with the knowledge that Oxford and London are cities an hour’s drive apart. Likewise, Bob did not drive to Oxford by car – He drove a van makes sense because car and van are similar vehicles.

Overall, no single source of context is sufficient. A combination of all contextual information—worldly, textual, physical, grammatical, intonation, etc.—is required to determine the correct interpretation.

While the different weights for the negation sets are context-dependent, some general observations can be made. Larger negation sets should tend to have smaller weights. The psychological reasoning is that humans are less able to focus on a larger number of details due to “limited information processing capacity” [10, 30]. Considering the previous example, one would require

a lot of context for the interpretation of “Alice doesn’t love Bob” to sensibly imply “Claire likes Dave”. Secondly, one can observe that the weight of a negation set should depend on the likelihood of its individual elements to be the target of the negation. If Alice being the target of the negation is unlikely, for instance if the entire text is about Alice, then the negation set of both Alice and love is also unlikely.

4.2.1 Determining weights using entailment

As mentioned earlier, one possible source of context can be the surrounding text. In a text which solely talks about Alice and Bob, the sentence “Alice doesn’t love Bob” probably intends to negate the word love, therefore asserting that Alice feels emotions other than love for Bob. Building on this intuition, we propose to use entailment measures to derive the weights for the different interpretations of a negation. If the given interpretation of the negation entails the surrounding sentences to a high degree, then the interpretation is consistent with the surrounding text, and hence it is more likely to be the intended meaning of the sentence.

We compare each possible interpretation of the negation with the surrounding sentences, where sentences closer in the text have more influence towards the final weighting than sentences that are further away. Let us consider the following, simplified scenario of a negation, followed immediately by the clarification with both sentences of the same grammatical structure:

This is not red wine
This is white wine

Here, we colour code the sentences for visual clarity, where the negated sentence is red. We want to determine the respective weights $p_{\{\text{red}\}}$, $p_{\{\text{wine}\}}$ and $p_{\{\text{red}, \text{wine}\}}$ of the negation sets $\{\text{red}\}$, $\{\text{wine}\}$ and $\{\text{red}, \text{wine}\}$.¹ As a human reader, the sentence following the negation clarifies that wine is not negated. Thus, the intended negation set is $\{\text{red}\}$.

To mathematically come to the same conclusion, we calculate the entailment between each interpretation of the negation and the follow-up sentence, i.e. how much each negation set of **not red wine** entails **white wine**. We compare the two sentences word by word (or more precisely, treating adjectives and nouns individually) and then take the product of the results. We consider:

- **not red wine** \sqsubseteq **white wine** - We first calculate the entailment of $\text{CN}_{\text{word}}(\text{red})$ with **white** which is medium as something that is not red could have many other colors, including white. The entailment of **wine** with **wine** is maximal as a wine is indeed a wine. Therefore the overall score of this interpretation is high.
Overall entailment: high
- **not red wine** \sqsubseteq **white wine** - This interpretation has medium entailment between $\text{CN}_{\text{word}}(\text{wine})$ and **wine** due to the fact that in distributional semantics, a word and its conversational negation appear in similar contexts [28, 30]. Yet the entailment between **red** and **white** is low since something being red does not entail that it is white.
Overall entailment: low
- **not red wine** \sqsubseteq **white wine** - This interpretation has a medium entailment between $\text{CN}_{\text{word}}(\text{red})$ and **white** and a medium entailment between $\text{CN}_{\text{word}}(\text{wine})$ and **wine**.
Overall entailment: medium

¹For the sake of simplicity, we ignore the fact that “This” could also be the target of the negation, as in “This is not red wine. That is!”.

Comparing the three interpretations, the first option has the highest score, matching our intuition of being the correct choice.

While this entailment method relies on the sentences having an identical grammatical structure to compare the sentences word by word, we can also directly compare entailment between two sentences. This requires composition operations and entailment measures which interact well to preserve word level entailment on the sentence level. This is still a field of active research, with some promising results presented in [19, 18, 34, 9].

5 Conversational negation of evolving meanings

Coecke [5] enhanced the DisCoCat framework to create DisCoCirc, which allows the meaning of sentences to be composed to obtain the meaning of text. In this section, we show that our conversational negation framework can be easily extended from sentences to conversational negation of text and evolving meanings.

5.1 Negating evolving meanings

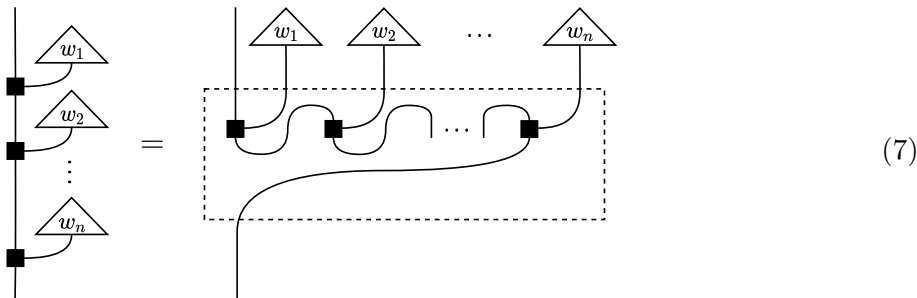
One of the key features of DisCoCirc is that it allows the meanings of entities to evolve as text evolves. The meanings are updated when the wires of the entities are composed with some meaning states. In DisCoCirc, texts that have the same meaning result in the same updates on the wires, even if they contain different sentences. For example, a text containing the following two sentences:

Bob is a scientist.
Bob is an alcoholic.

results in the same circuit as the text containing the single sentence:

Bob who is a scientist is an alcoholic.


This motivates us to think of all the meaning updates to a wire as a single large sentence. If we have a sequence of updates to a wire, we can morph the wire using the snake equation. We thereby derive a single meaning update process that updates the wire with a sequence of word meanings. Hence, for a wire whose meaning evolves through updates by words w_1, \dots, w_n , we have:



$$(7)$$

Now, to perform conversational negation, we can simply apply the CN function (see Section 4) that maps the circuit of a sentence to a circuit where the sentence is pre-composed with $\text{CN}_{\text{string}}$.

Thus, the conversational negation of a dynamic evolving entity becomes:




This idea can also be applied to text with multiple, interacting actors. Consider the text:

Alice is evil.
 Bob is old.
 Alice loves Bob.

Alternatively, we can write this as a single sentence, omitting commas for simplicity:

Alice who is evil loves Bob who is old.

The circuit for the text looks like the following:

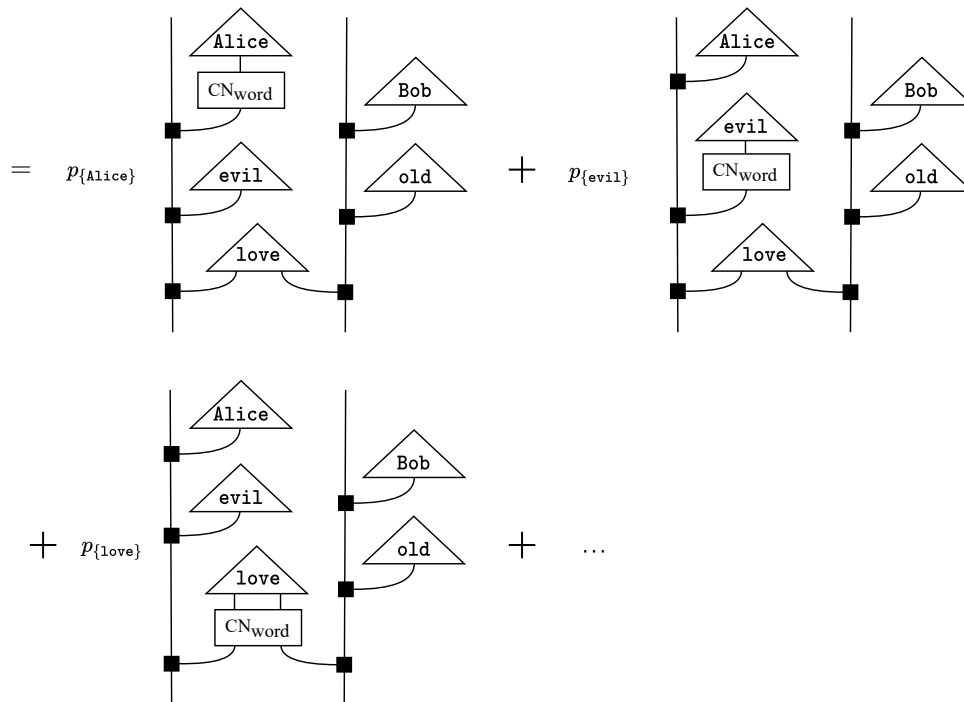
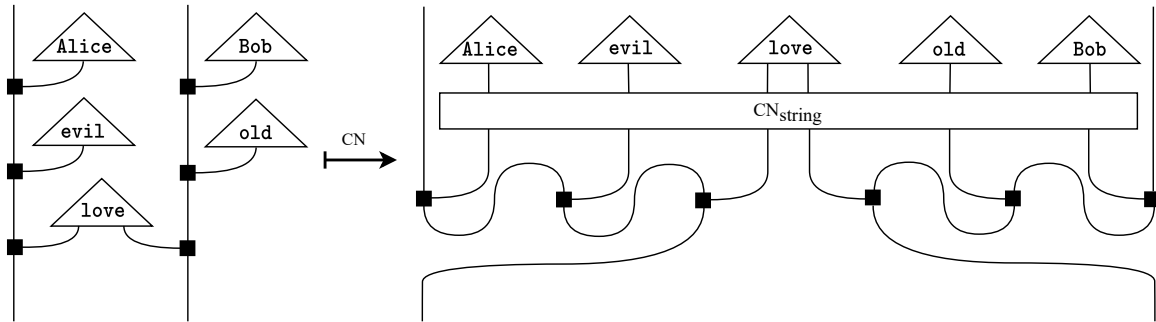


Suppose we want to find the conversational negation of the actor named Alice in the above text, i.e. we want to know that if someone is not Alice, who else could they be. Similar to the case of conversational negation of sentences, we have various possibilities based on the meaning updates that have been performed on the wire corresponding to Alice. Any subset of the words that have contributed to the meaning update of Alice — either directly or via meaning updates to entangled actors — could be negated. For example, possibilities of who “not Alice” could be include:

- a) Claire who is evil and loves old Bob
- b) Alice (different person but same name) who is virtuous and hates old Bob
- c) Dave who is evil and loves young Bob
- d) Claire who is virtuous and hates old Daisy

If we again think of the text as a single large sentence, we can simply apply the CN function

from Section 4. We get:



If we want to perform the conversational negation of the actor **Bob** instead of **Alice**, we will again get the same resulting circuit after applying the CN function. The difference to the conversational negation of **Alice** is that we will have different weights for the negation sets. For example, the weight of the negation set $\{\text{Bob, old}\}$ is likely low, for the negation of **Alice** — someone who is not **Alice** — is intuitively unlikely to be someone who is called **Alice** and is evil but loves a person with different attributes than **Bob**. In contrast, the weight of the same negation set could be reasonably high, for the negation of **Bob** — someone who is not **Bob** — could have a different name and age but still be the lover of **Alice**. Therefore, while having the same diagrammatic representation, the negation of the actor **Alice** can give a fundamentally different result than the negation of **Bob** through the weights which were chosen for the negation sets.

5.2 Example: finding alternatives

Consider the following text where the meanings of the words evolves as the text evolves:

Alice is a human.	Alice is an archaeologist.
Bob is a human.	Bob is a biologist.
Claire is a human.	Claire is a pianist.
Daisy is a dog.	Daisy is a pet.

Suppose we want to perform the conversational negation of the actor **Alice** and evaluate how much it entails the other actors: **Bob**, **Claire** and **Daisy**. Based on the given text, it is reasonable to expect that someone who is not **Alice** (a human archaeologist) is more likely to be **Bob** (a human biologist) than **Claire** (a human pianist). In fact, someone who is not **Alice** is still more likely to be **Claire** (a human pianist) than **Daisy** (a pet dog). Now we will analyze if the conversational negation presented in Section 5.1 reflects this intuition.

When we apply the conversational negation to the actor **Alice**, we get a mixture containing all possible negation sets along with their weights. These negation sets are nonempty subsets of $\{\mathbf{Alice}, \mathbf{human}, \mathbf{archaeologist}\}$. The weights of the negation sets must be determined using all contexts as discussed in Section 4.2. However, to explore the maximum entailment that can be achieved from the negation of the actor **Alice** to each of the remaining actors, we only consider the most appropriate negation sets of “not **Alice**” for each actor.

- **Bob** - Since Bob is a human biologist, the best negation set of “not **Alice**” for Bob is $\{\mathbf{Alice}, \mathbf{archaeologist}\}$. The table below shows the entailment between this negation set of “not **Alice**” and the actor **Bob**. From the table, it is clear that “not **Alice**” **highly** entails **Bob**.

“not Alice ”	Bob	Entailment
$\text{CN}_{\text{word}}(\mathbf{Alice})$	Bob	medium
human	human	1 (max)
$\text{CN}_{\text{word}}(\mathbf{archaeologist})$	biologist	high

- **Claire** - Similar to **Bob**, the best negation set for **Claire** is $\{\mathbf{Alice}, \mathbf{archaeologist}\}$. As shown in table below, “not **Alice**” **moderately** entails **Claire**.

“not Alice ”	Claire	Entailment
$\text{CN}_{\text{word}}(\mathbf{Alice})$	Claire	medium
human	human	1 (max)
$\text{CN}_{\text{word}}(\mathbf{archaeologist})$	pianist	medium

- **Daisy** - For **Daisy** the pet dog, the best negation set of “not **Alice**” is $\{\mathbf{Alice}, \mathbf{human}, \mathbf{archaeologist}\}$. From the table below, “not **Alice**” only **slightly** entails **Daisy**.

“not Alice”	Daisy	Entailment
CN _{word} (Alice)	Daisy	medium
CN _{word} (human)	dog	medium
CN _{word} (archaeologist)	pet	low

Therefore, in our proposed framework, someone who is not **Alice** has the highest chance to be (from most to least likely): **Bob**, **Claire** and **Daisy**, which indeed lines up with human intuition. However, the final result of the negation depends on the weights of the negation sets, determined by the context. Hence, if for some reason the negation set {**Alice**, **human**, **archaeologist**} has been determined to be the correct interpretation, then “not Alice” might be more closely related to **Daisy** than **Claire** after all.

6 Discussion and future work

The framework to model conversational negation, proposed in this paper, utilises a new mechanism not currently present in DisCoCirc: the external derivation of the weights required for the mixture of different interpretations of a negation. This is motivated by the observation that disambiguation in language relies on an understanding of the world not necessarily present in the text. In the case of the conversational negation of words, this understanding is captured in a single meaning state derived from existing lexical entailment hierarchies, which we called the worldly context (Section 3.2). In the case of conversational negation of sentences and evolving meanings, the weights allow for this understanding to be integrated into the circuit. While Section 4.2.1 provides an intuition of how the surrounding text can partially inform the correct choice of weights, they should also take into consideration other contexts such as worldly knowledge, intonation, or the environment of the speaker. One goal for future work is to explore these sources of context and find methods to incorporate that information into the weights. Apart from using the weights to inject context, we would like to explore embedding the sources of context directly within the representation of word meanings; for instance, by building upon the work on conceptual space models of cognition in compositional semantics [4].

The ideas presented in this paper are built on the intuitions gathered from psychology papers and the experiments performed for the conversational negation of words in [33]. To empirically validate the framework for conversational negation of sentences and evolving meaning, experiments should be devised. With the basic intuition for deriving the weights from surrounding sentences being solely presented for grammatically identical sentences, further work needs to be done to generalise this process.

In this paper, we modelled the conversational negation of a sentence as the conversational negation of its constituent words. We focused on single-word constituents, 1) for the purpose of clarity, and 2) because, in our prior work, we proposed and experimentally validated a framework for conversationally negating a single word [33]. However, our model does not forbid constituents made of multiple words. For instance, in the following example the negated constituent is “the key to the garage”:

This is not the key to the garage. - This is a toy.

If we were to broaden the scope of our conversational negation to generate plausible alternatives

to a sentence with a different grammatical structure, we would require a mechanism for parsing and negating constituents made of multiple words. In this view, it could then be possible to model the negation of non-conjunctive composition of concepts such as the compositional explanation of the “pet fish” phenomenon by Coecke and Lewis [6].

Although we have devised a treatment for modelling the conversational negation of a sentence as the conversational negations of its constituent words then composed together grammatically, there are other forms of negation in natural language. Since Klima [21] and Jackendoff [15], most linguists have treated the form of negation we model — constituent negation — as distinct from sentential negation [14]. Consider the sentence **I did not walk my dog**. Constituent negation invokes the collective mixture of possible alternative interpretations: that someone else walked the dog, I cuddled the dog, it was my friend’s dog, etc. Instead, it may be that we simply want to negate the sentence as a whole — it is untrue that I walked my dog. Without experiment, it is unclear the extent to which logical negation of the sentence meaning alone suffices to model sentential negation.

Another related challenge is to formalise a mathematical model of the logic underlying conversational NOT, AND, and OR. This requires investigating the extent to which boolean logic holds in a setting known not to follow boolean logic. A long-term goal would be to extend the conversational negation process to a *conversational logic* process, compatible with compositional distributional semantics, particularly its properties with regards to entailment.

Acknowledgements

We would like to thank Vincent Wang for the insightful discussions. We thank the anonymous reviewers for their helpful feedback. Lia Yeh gratefully acknowledges funding from the Oxford-Basil Reeve Graduate Scholarship at Oriel College in partnership with the Clarendon Fund.

References

- [1] Esma Balkir, Mehrnoosh Sadrzadeh & Bob Coecke (2016): *Distributional Sentence Entailment Using Density Matrices*. In Mohammad Taghi Hajiaghayi & Mohammad Reza Mousavi, editors: *Topics in Theoretical Computer Science*, Springer International Publishing, Cham, pp. 1–22, doi:10.1007/978-3-319-28678-5_1.
- [2] Dea Bankova, Bob Coecke, Martha Lewis & Dan Marsden (2019): *Graded hyponymy for compositional distributional semantics*. *Journal of Language Modelling* 6(2), p. 225–260, doi:10.15398/jlm.v6i2.230.
- [3] William Blacoe, Elham Kashefi & Mirella Lapata (2013): *A Quantum-Theoretic Approach to Distributional Semantics*. In: *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Association for Computational Linguistics, Atlanta, Georgia, pp. 847–857. Available at <https://www.aclweb.org/anthology/N13-1105>.
- [4] Joe Bolt, Bob Coecke, Fabrizio Genovese, Martha Lewis, Dan Marsden & Robin Piedeleu (2019): *Interacting Conceptual Spaces I: Grammatical Composition of Concepts*, pp. 151–181. Springer International Publishing, Cham, doi:10.1007/978-3-030-12800-5_9.
- [5] Bob Coecke (2020): *The Mathematics of Text Structure*. arXiv:1904.03478.

- [6] Bob Coecke & Martha Lewis (2016): *A Compositional Explanation of the ‘Pet Fish’ Phenomenon*. In Harald Atmanspacher, Thomas Filk & Emmanuel Pothos, editors: *Quantum Interaction*, Springer International Publishing, Cham, pp. 179–192, doi:10.1007/978-3-319-28675-4_14.
- [7] Bob Coecke & Konstantinos Meichanetzidis (2020): *Meaning updating of density matrices*. *FLAP* 7, pp. 745–770. Available at <https://arxiv.org/abs/2001.00862>.
- [8] Bob Coecke, Mehrnoosh Sadrzadeh & Stephen Clark (2010): *Mathematical foundations for a compositional distributional model of meaning*. *Linguistic Analysis* 36, pp. 345–384. Available at <https://arxiv.org/abs/1003.4394>.
- [9] Gemma De las Cuevas, Andreas Klingler, Martha Lewis & Tim Netzer (2020): *Cats climb entails mammals move: preserving hyponymy in compositional distributional semantics*. *CoRR* abs/2005.14134. arXiv:2005.14134.
- [10] Jonathan St B. T. Evans (1989): *Bias in human reasoning: Causes and consequences*. Lawrence Erlbaum Associates, Inc.
- [11] Jonathan St B. T. Evans, John Clibbens & Benjamin Rood (1996): *The Role of Implicit and Explicit Negation in Conditional Reasoning Bias*. *Journal of Memory and Language* 35, pp. 392–409, doi:10.1006/jmla.1996.0022.
- [12] Christiane Fellbaum (1998): *WordNet: An Electronic Lexical Database*, doi:10.7551/mitpress/7287.001.0001. Available at <https://wordnet.princeton.edu/>.
- [13] Marti A. Hearst (1992): *Automatic Acquisition of Hyponyms from Large Text Corpora*. In: *Proceedings of the 14th Conference on Computational Linguistics - Volume 2, COLING ’92*, Association for Computational Linguistics, USA, p. 539–545, doi:10.3115/992133.992154.
- [14] Laurence R. Horn & Heinrich Wansing (2020): *Negation*. In Edward N. Zalta, editor: *The Stanford Encyclopedia of Philosophy*, Spring 2020 edition, Metaphysics Research Lab, Stanford University.
- [15] Ray S. Jackendoff (1969): *An Interpretive Theory of Negation*. *Foundations of Language* 5(2), pp. 218–241, doi:10.2307/25000377.
- [16] Hans Kamp & Uwe Reyle (2013): *From discourse to logic: Introduction to modeltheoretic semantics of natural language, formal logic and discourse representation theory*. 42, Springer Science & Business Media.
- [17] Dimitri Kartsaklis & Mehrnoosh Sadrzadeh (2013): *Prior Disambiguation of Word Tensors for Constructing Sentence Vectors*. In: *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, Association for Computational Linguistics, Seattle, Washington, USA, pp. 1590–1601. Available at <https://www.aclweb.org/anthology/D13-1166>.
- [18] Dimitri Kartsaklis & Mehrnoosh Sadrzadeh (2016): *A Compositional Distributional Inclusion Hypothesis*. In Maxime Amblard, Philippe de Groote, Sylvain Pogodalla & Christian Retoré, editors: *Logical Aspects of Computational Linguistics. Celebrating 20 Years of LACL (1996–2016)*, Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 116–133, doi:10.1007/978-3-662-53826-5_8.
- [19] Dimitri Kartsaklis & Mehrnoosh Sadrzadeh (2016): *Distributional Inclusion Hypothesis for Tensor-based Composition*. In: *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, The COLING 2016 Organizing Committee, Osaka, Japan, pp. 2849–2860. Available at <https://aclanthology.org/C16-1268>.
- [20] Dimitri Kartsaklis, Mehrnoosh Sadrzadeh & Stephen Pulman (2013): *Separating Disambiguation from Composition in Distributional Semantics*. In: *Proceedings of the Seventeenth Conference on Computational Natural Language Learning*, Association for Computational Linguistics, Sofia, Bulgaria, pp. 114–123. Available at <https://www.aclweb.org/anthology/W13-3513>.
- [21] Edward S Klima (1964): *Negation in English*. In J. A. Fodor & J. J. Katz, editors: *The Structure of Language*, Englewood Cliffs: Prentice-Hall, pp. 246–323.

- [22] Germán Kruszewski, Denis Paperno, Raffaella Bernardi & Marco Baroni (2016): *There is no logical negation here, but there are alternatives: Modeling conversational negation with distributional semantics*. *Computational Linguistics* 42(4), pp. 637–660, doi:10.1162/COLI_a.00262.
- [23] Joachim Lambek (2007): *From Word to Sentence: A Pregroup Analysis of the Object Pronoun Who(m)*. *Journal of Logic, Language, and Information* 16(3), pp. 303–323, doi:10.1007/s10849-006-9035-9.
- [24] Edward N. Lee (1972): *Plato on Negation and Not-Being in the Sophist*. *The Philosophical Review* 81(3), pp. 267–304, doi:10.2307/2184327.
- [25] Martha Lewis (2019): *Compositional Hyponymy with Positive Operators*. In: *Proceedings of the International Conference on Recent Advances in Natural Language Processing (RANLP 2019)*, IN-COMA Ltd., Varna, Bulgaria, pp. 638–647, doi:10.26615/978-954-452-056-4.075.
- [26] Martha Lewis (2020): *Towards logical negation for compositional distributional semantics*. *IfCoLoG Journal of Logics and their Applications* 7(3). Available at <https://arxiv.org/abs/2005.04929>.
- [27] Francois Meyer & Martha Lewis (2020): *Modelling Lexical Ambiguity with Density Matrices*. In: *Proceedings of the 24th Conference on Computational Natural Language Learning*, Association for Computational Linguistics, Online, pp. 276–290, doi:10.18653/v1/2020.conll-1.21. Available at <https://www.aclweb.org/anthology/2020.conll-1.21>.
- [28] Saif M. Mohammad, Bonnie J. Dorr, Graeme Hirst & Peter D. Turney (2013): *Computing Lexical Contrast*. *Computational Linguistics* 39(3), pp. 555–590, doi:10.1162/COLI_a.00143.
- [29] Mike Oaksford (2002): *Contrast classes and matching bias as explanations of the effects of negation on conditional reasoning*. *Thinking & Reasoning* 8(2), pp. 135–151, doi:10.1080/13546780143000170.
- [30] Mike Oaksford & Keith Stenning (1992): *Reasoning with conditionals containing negated constituents*. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 18(4), p. 835, doi:10.1037/0278-7393.18.4.835.
- [31] Robin Piedeleu, Dimitri Kartsaklis, Bob Coecke & Mehrnoosh Sadrzadeh (2015): *Open system categorical quantum semantics in natural language processing*. arXiv preprint arXiv:1502.00831.
- [32] Jérôme Prado & Ira Noveck (2006): *How reaction time measures elucidate the matching bias and the way negations are processed*. *Thinking & Reasoning - THINK REASONING* 12, pp. 309–328, doi:10.1080/13546780500371241.
- [33] Benjamin Rodatz, Razin Shaikh & Lia Yeh (2021): *Conversational Negation using Worldly Context in Compositional Distributional Semantics*. In: *Proceedings of the 2021 Workshop on Semantic Spaces at the Intersection of NLP, Physics, and Cognitive Science (SemSpace)*, Association for Computational Linguistics, Groningen, The Netherlands, pp. 53–65. Available at <https://arxiv.org/abs/2105.05748>.
- [34] Mehrnoosh Sadrzadeh, Dimitri Kartsaklis & Esmā Balkır (2018): *Sentence entailment in compositional distributional semantics*. *Annals of Mathematics and Artificial Intelligence* 82(4), p. 189–218, doi:10.1007/s10472-017-9570-x.
- [35] Peter Selinger (2007): *Dagger Compact Closed Categories and Completely Positive Maps: (Extended Abstract)*. *Electronic Notes in Theoretical Computer Science* 170, pp. 139–163, doi:10.1016/j.entcs.2006.12.018. Proceedings of the 3rd International Workshop on Quantum Programming Languages (QPL 2005).
- [36] Dominic Widdows & Stanley Peters (2003): *Word vectors and quantum logic: Experiments with negation and disjunction*. *Mathematics of language* 8(141-154). Available at <https://www.puttypeg.net/papers/quantum-senses.pdf>.