

A Model of Anaphoric Ambiguities using Sheaf Theoretic Quantum-like Contextuality and BERT

Kin Ian Lo

Mehrnoosh Sadrzadeh

Shane Mansfield

University College London
London, UK

Quandela
Paris, France

{kin.lo.20,m.sadrzadeh}@ucl.ac.uk

shane.mansfield@quandela.com

Ambiguities of natural language do not preclude us from using it and context helps in getting ideas across. They, nonetheless, pose a key challenge to the development of competent machines to understand natural language and use it as humans do. Contextuality is an unparalleled phenomenon in quantum mechanics, where different mathematical formalisms have been put forwards to understand and reason about it. In this paper, we construct a schema for anaphoric ambiguities that exhibits quantum-like contextuality. We use a recently developed criterion of sheaf-theoretic contextuality that is applicable to signalling models. We then take advantage of the neural word embedding engine BERT to instantiate the schema to natural language examples and extract probability distributions for the instances. As a result, plenty of sheaf-contextual examples were discovered in the natural language corpora BERT utilises. Our hope is that these examples will pave the way for future research and for finding ways to extend applications of quantum computing to natural language processing.

1 Introduction

Context plays a central role in determining meanings of words, as words that often occur in similar contexts have similar meanings. Conjured in the 1950's by Firth [10] and Harris [11], this hypothesis led to the field of *Distributional* semantics. Harris noticed that words such as 'eye doctor' and 'optometrist' occur in the same contexts, e.g. in the neighbourhood of 'eye' and 'glasses'. Firth's infamous quote was that you know a word by the company it keeps. The Distributional hypothesis has been formalised by vector semantics and implemented on large corpora of data. Originally, corpora of text were mined to build term-term co-occurrence matrices. Nowadays, contextualised deep neural network architectures such as **BERT** are used to train the vector statistics. Despite the daily successes of contextualised embeddings in Natural Language Processing tasks, they do not have an explicit notion of grammatical or discourse structure. The statistics learnt by engines such as BERT do indeed take some structure into account when training vector embeddings, but they certainly do not represent the grammatical or discourse structures of a piece of text in a vector in the same as they do for words. Compositional distributional semantics [4, 6] is a field of research introduced in an attempt to address this challenge. A line of research of this field showed that by generalising the notion of vectors to tensors, one can embed both word and grammar. Recent research, has showed how quantum computing devices such as the IBMQ's quantum devices can be used to learn these tensors as quantum states [14].

The links between natural language and quantum physics goes beyond the above. Discovery of scenarios such as EPR [9] and Bell [5], made quantum mechanics the first science to formally deal with the notion of contextuality. Scientists argued that quantum theory should be contextual in order to be sound and different mathematical formalisms were introduced to analyse this. Quantum-like contextuality turned out to be, essentially, the failure of having a global explanation to local observations on a system and the presence of incompatible observables, in the sense that a simultaneous global observation

of all observables are not possible, except in trivial systems. Over the last number of years it has been proved that it is this feature of quantum mechanics that is capable of lifting linear computation to universal computation [3, 17, 15] and that contextuality is necessary for magic state distillation [12], a key component in fault-tolerant quantum computing schemes. Roughly speaking, contextual systems hold additional computational power which is absent in non-contextual systems. It is therefore a reasonable conjecture that quantum computers are better at dealing with contextual systems compared to classical computers. There exist a number of different frameworks for treating contextuality. The sheaf-theoretic framework of [2] is amongst the ones that connects the statistical data collected from quantum experiments to the structures defined by quantum mechanics. One of these laws is the *no-signalling* property. The sheaf-theoretic framework can only formalise contextual scenarios that are no-signalling. However, the examples we are aiming to study are highly unlikely to be non-signalling. We remedy this by using a recent extension of them to realistic experiments [16], where the authors derive a new inequality to check the contextuality of systems of measurements with signalling data. In this setting, some degree of signalling becomes possible. We use this inequality to check the contextuality of our examples.

Quantum-like contextuality has been observed in other fields, e.g. in behavioural sciences [8] and natural language [18, 19]. In [18, 19] Wang et al showed that pairs of ambiguous words can produce contextual systems that resemble the Bell/CHSH quantum measurement scenario. In this paper, we propose a novel linguistic construction that exhibits the contextuality of Coreference ambiguities and exemplify it to anaphoric relations. We use BERT to instantiate the construction and extract probability distributions for the instances. Checking the contextuality fraction for these instances showed that it is possible to discover examples of anaphoric ambiguity that exhibit quantum-like contextuality properties. In fact, we were able to find hundreds of examples after only working with a few pairs of nouns and their corresponding adjectives, verbs, and prepositional phrases. We hope finding contextual schema and instances in natural language data help us devise new quantum algorithms that can handle ambiguities better than classical computers do in natural language processing tasks.

2 Ambiguities in Natural Language

One of the ambiguities of natural languages comes from the fact that words have different meanings. For example, ‘bat’ has an animal meaning and a sport meaning, ‘plant’ can mean a living organism such as a tree or a shrub, or a manufacturing industrial unit, such as a power plant. *Word Sense Disambiguation* is a long standing task and evaluation method in Natural Language Processing. Here the goal is to identify which meaning of a word is being used in a context. Another major ambiguity in natural language comes from the *Coreference Resolution* task: the task of deciding which discourse entity is referring to which expression in a context. This is an important part of language engines such as dialogue systems or question answering. For instance, in an automatic MOT booking system, the NLP engine should know which car the user is referring to when they say ‘I have a Toyota RAV4 and a Toyota Aqua, it is the hybrid one for which I need an MOT today.’

Different Coreference Resolution algorithms focus on different classes of referring expressions. Pronouns are in the class of definite referents and refer to entities that are identifiable from the context, because they have been mentioned before (or after). In the discourse ‘Dawn called the AA. The car had broken down and she had no choice’, the pronoun ‘she’ refers to the definite noun phrase ‘Dawn’ and is an instance of the linguistic phenomena *anaphora*. Despite presence of linguistics properties in the anaphoric relations, such as gender and number agreement and grammatical role and verb preferences, these are ambiguities. In ‘Dawn texted Wendy. Her car had broken down.’, or ‘Dawn phoned Wendy.

She was upset and needed help.’, it is not clear whose car was broken or who was upset. Pronominal coreference relations are many-to-many and the ambiguities arise from them taking complex forms. A pronoun can refer to multiple referents and multiple pronouns can refer to the same referent. In the discourse ‘There is a man carrying a boy. He is tired and worn out. He is snoring.’, the first *He* can refer to both *man* or *boy*, but the second *He* most certainly refers to *boy*.

The different choices that give rise to ambiguities, be it in the choice of the meaning of a word in a Word Sense Disambiguation task, or the choice of the potential referent of an expression in a Coreference Resolution task, give rise to probability distributions. An ambiguous word can be treated as an observable which can have possible outcomes. In case of meaning ambiguities, these outcomes are possibilities over the semantic interpretations of the word. A probability distribution over the outcomes can then be defined using the frequency of occurrences of the possible interpretations in a corpus, e.g. the entire English Wikipedia, or in plausibility judgements of human subjects. A single observable is not sufficient to support contextuality, instead pairs of ambiguous words are needed. A pair of words is thought of as a pair of compatible observables measured simultaneously.

The work of [18, 19] focused on meaning ambiguities. In this paper, we focus on coreference ambiguities and model contextual features of ambiguities arising from anaphoric reference relations. An identical approach can be taken if the relationship is cataphoric. We treat the pronouns as observables, of which the measurement outcomes are the possible referents of each pronoun. In what follows we describe the mathematical setting we used, detail how to use it to model ambiguous anaphoric references, explain how we found contextual examples, and present some of the contextual examples.

3 Sheaf Theoretic Framework

In the sheaf-theoretic framework of contextuality [2], a measurement scenario is a tuple $\langle \mathcal{X}, \mathcal{M}, \mathcal{O} \rangle$ with the data \mathcal{X} , a set of observables, \mathcal{M} , a measurement cover, and \mathcal{O} , a set of measurement outcomes. An observable in \mathcal{X} is a quantity that can be measured to give one of the outcomes in \mathcal{O} . A subset of simultaneously measurable observables of \mathcal{X} is called a measurement context (or simply called a context). The measurement cover \mathcal{M} is a collection of contexts which covers \mathcal{X} , i.e. the union of all contexts in \mathcal{M} is \mathcal{X} .

For every measurement context, we can perform a number of repeated simultaneously measurements on the observables in the context. The gathered statistics can then be used to reconstruct an estimated joint probability distribution. Instead, one can also calculate the joint distribution exactly using an underlying theory of the concerned system, e.g. using Born’s rule in quantum mechanics for a quantum system.

An *empirical model* refers to a collection of such joint probability distribution for each context in the measurement cover \mathcal{M} . By definition, a subset of observables in \mathcal{X} that are not all included in a measurement context in \mathcal{M} cannot be measured simultaneously. Therefore, a joint distribution over the said observables cannot be empirically estimated. The empirical model of a system fully encapsulates what is to be known from the system with empirical measurements.

Contextuality comes from the failure of explaining an empirical model in a classically intuitive way: assuming that all measurements are just revealing deterministic pre-existing values, in other words, the measurement outcomes are already fixed when the system was prepared. Thus the randomness comes entirely from the system preparation. That means that there is a global joint distribution over all the observables in the scenario, which marginalises to every local joint distribution in the empirical model. Given an empirical model, if such a global distribution does not exist, then we call such empirical model

	(0,0)	(0,1)	(1,0)	(1,1)		(0,0)	(0,1)	(1,0)	(1,1)
(a_1, b_1)	1/2	0	0	1/2	(a_1, b_1)	1	0	0	1
(a_1, b_2)	3/8	1/8	1/8	3/8	(a_1, b_2)	1	1	1	1
(a_2, b_1)	3/8	1/8	1/8	3/8	(a_2, b_1)	1	1	1	1
(a_2, b_2)	1/8	3/8	3/8	1/8	(a_2, b_2)	1	1	1	1

Figure 1: Empirical tables of measurement scenarios: Bell/CHSH (left), possibilistic Bell/CHSH (right)

contextual. Note that such global distribution exists only in theory as there are observables in \mathcal{X} that cannot be measured simultaneously, unless in trivial scenarios.

For readers familiar with sheaf theory, the said criterion for contextuality can be formalised using the language of sheaf. Consider the presheaf \mathcal{F} which assigns each subset $U \in \mathcal{P}(\mathcal{X})$ the set of all possible probability distributions on the observables in U . Each set inclusion $U \subseteq U'$, interpreted as an arrow in the category $\mathcal{P}(\mathcal{X})$, is mapped to the marginalisation of distributions on U' to distributions on U . For a measurement cover \mathcal{M} , an empirical model is just a family of compatible distributions $\{D_C\}_{C \in \mathcal{M}}$. The presheaf \mathcal{F} is a sheaf if the gluing property is satisfied:

Fix a cover \mathcal{M} of \mathcal{X} . For each family of compatible sections $\{D_C\}_{C \in \mathcal{M}}$, there is a unique global distribution compatible with every distributions in $\{D_C\}_{C \in \mathcal{M}}$.

Thus a contextual empirical model can only live on a measurement scenario for which the presheaf \mathcal{F} is not a sheaf, i.e. not satisfying the gluing property. To say that there is a contextual model that lives on a measurement scenario is to say that the presheaf \mathcal{F} is not a sheaf on the scenario.

As an example, the Bell/CHSH scenario involves two experimenters, Alice and Bob, who share between them a two-qubit quantum state. Alice is allowed to measure her part of the state with one of two incompatible observables, a_1 and a_2 , which gives either 0 or 1 as the outcome. Similarly Bob can choose to measure his part with observables b_1 and b_2 . Therefore, the Bell/CHSH measurement scenario is fully described with the following data: $\mathcal{X} = \{a_1, b_1, a_2, b_2\}$, $\mathcal{M} = \{\{a_1, b_1\}, \{a_1, b_2\}, \{a_2, b_1\}, \{a_2, b_2\}\}$, and $\mathcal{O} = \{0, 1\}$. Notice that $\{a_1, a_2\}$ and $\{b_1, b_2\}$ are not in \mathcal{M} as they cannot be measured simultaneously due to their quantum mechanical incompatibility.

So far we have specified what measurements are allowed and what outcomes are possible. Suppose now Alice and Bob repeat the experiment many times and have gathered sufficient statistics to estimate the joint probability distribution for each context in \mathcal{M} . Their results can be summarised in a table referred to as an *empirical table*, see Figure 1, where each row in the table represents a joint distribution on the context shown in the leftmost column. For instance, the bottom right entry in the table (1/8) is the probability of both Alice and Bob getting 1 as their measurement outcomes when Alice chooses to measure a_2 and Bob chooses to measure b_2 . Note that the empirical model of the system is entirely described by the empirical table.

One can show that, using elementary linear algebra, there exists no global distribution over $\{a_1, a_2, b_1, b_2\}$ that marginalises to the 4 local distribution shown in the above empirical table. Therefore, the empirical model considered here is indeed contextual.

Instead of probability, one can also consider possibility, i.e. whether an outcome is possible or not. If we use Boolean values to represent possibility, 0 for *impossible* and 1 for *possible*, the passage from probability to possibility is just a mapping of all zero probabilities to 0 and all non-zero probabilities to 1. This (irreversible) mapping is called a *possibilistic collapse* of the model. For the empirical table of the possibilistic version of Bell/CHSH see Figure 1. One can visualise a possibilistic model with a bundle diagram, see Figure 2:




Figure 2: Bundle diagrams of possibilistic CHSH (left), PR box (middle), PR prism (right)

The base (i.e. the bottom part) of the bundle diagram represents the measurement cover \mathcal{M} , where each vertex represents an observable in \mathcal{X} . An edge is drawn between two observables if they can be simultaneously measured, i.e. in the same measurement context. What sits on top of the base represents the possible outcomes. For instance, the presence of the edge connecting the 0 vertex on top of a_1 and the 0 vertex on the observable b_1 means that it is possible to get the joint outcome $(0,0)$ when the context (a_1, b_1) is measured.

A system is logically contextual if the inexistence of a global distribution can already be deduced by looking at the supports of the context-wise distributions – or equivalently if the Boolean distributions obtained by the *possibilistic collapse* of the model [2] is contextual. Such systems are said to be *possibilistically contextual*. Logical contextuality manifests on a bundle diagram as the failure of extending at least one of the edges to a loop that wraps around the base once. For the possibilistic empirical model of a PR box, see Figure 2. Note that none of the edges is extendable to a loop that wraps around the base once. Given the possibilistic collapse of an empirical model, if none of the edges can be extendable to a loop that wraps around the base once, we say that the model is *strongly contextual*.¹

Proposition 1 *The minimal measurement scenario that admits contextuality has the data up to relabelling: $\mathcal{X} = \{x_1, x_2, x_3\}$, $\mathcal{M} = \{\{x_1, x_2\}, \{x_2, x_3\}, \{x_1, x_3\}\}$, and $\mathcal{O} = \{0, 1\}$.*

The proof of the above is routine, and so is that of the following:

Proposition 2 *The only strongly contextual system, up to relabelling, for the minimal measurement scenario is where perfect correlation is observed on two of the contexts and perfect anti-correlation is observed on the other one.*

We call this scenario the *PR prism* as an analogy to the PR boxes. See Figure 2 for its bundle diagram. The pairs of parallel edges over contexts $\{x_2, x_3\}$ and $\{x_3, x_1\}$ correspond to perfect correlation and the pair of crossed edges over context $\{x_1, x_2\}$ corresponds to the perfect anti-correlation.

4 Contextual and Signalling Fractions

The contextual fraction (CF) [1] measures the degree of contextuality of a given non-signalling model. Given an empirical model e , the CF of e is defined as the minimum λ such that the following convex

¹Strictly speaking, this definition of strong contextuality only applies to cyclic scenarios where the base of the bundle diagram forms a loop. Nonetheless, cyclic scenarios are the only scenarios considered in this paper.

decomposition of e works²:

$$e = (1 - \lambda)e^{NC} + \lambda e^C, \quad (1)$$

where e^{NC} is a non-contextual (and non-signalling) empirical model and e^C is a model allowed to be contextual. For non-signalling models, the criterion of contextuality is just

$$CF > 0. \quad (2)$$

As e^{NC} is not allowed to be signalling, the CF of a signalling model must be greater than zero. Thus, interpreting CF as a measure of contextuality for signalling models would lead to erroneous conclusions. However, most models, including the ones considered in this paper, are signalling.

One can try to define a signalling fraction (SF), in the same way CF is defined, to quantify the degree of signalling. Given a model e , the SF of e is defined as the minimum μ such that the following convex decomposition of e works:

$$e = (1 - \mu)e^{NS} + \mu e^S, \quad (3)$$

where e^{NS} is a non-signalling empirical model and e^S is a model allowed to be signalling.

In [16], the signalling fraction (SF) was used to quantify the amount of fictitious contextuality contributing to the contextual fraction due to signalling in a signalling model. The authors derived a criterion of contextuality for signalling models that reads

$$CF > 2|\mathcal{M}|SF, \quad (4)$$

where $|\mathcal{M}|$ denotes the number of measurement contexts. Notice how criterion (4) reduces to the generalised criterion (2) when $SF = 0$.

In the general case, one would need to solve a linear program to calculate the CF or SF of a model. The calculation is much simpler with models that share the same support as the PR prism. We call these model *PR-like*. Such models can always be written as the following empirical table upon relabelling:

	(0,0)	(0,1)	(1,0)	(1,1)
(x_1, x_2)	$(1 + \varepsilon_1)/2$	0	0	$(1 - \varepsilon_1)/2$
(x_2, x_3)	$(1 + \varepsilon_2)/2$	0	0	$(1 - \varepsilon_2)/2$
(x_3, x_1)	0	$(1 + \varepsilon_3)/2$	$(1 - \varepsilon_3)/2$	0

where $-1 \leq \varepsilon_1, \varepsilon_2, \varepsilon_3 \leq 1$. Recall that the model e^{NC} in the convex decomposition (1) is noncontextual and non-signalling. For a PR-like model to be non-signalling, one can check that a PR-like model is non-signalling if and only if it is a PR box, i.e. $\varepsilon_1 = \varepsilon_2 = \varepsilon_3 = 0$. However, the PR box is known to be (strongly) contextual. Thus, there does not exist a model e^{NC} that is noncontextual and non-signalling for a PR-like model. Therefore, the SF of PR-like model is always 1.

The calculation of SF for PR-like models is also simple. As e^{NS} in the convex decomposition (3) can be contextual but not signalling, e^{NS} must be the PR box, the one with $\varepsilon_1 = \varepsilon_2 = \varepsilon_3 = 0$. As we cannot have negative probabilities in e^S in the decomposition, the coefficient $(1 - \mu)$ can at most be double the smallest non-zero value in the table, that is, $\min(1 \pm \varepsilon_i)$. Thus we have

$$SF = 1 - \min_{i=1,2,3} (1 \pm \varepsilon_i) = \max_{i=1,2,3} |\varepsilon_i|$$

for PR-like models. We will use this result to calculate the SF of the PR-like models we constructed in the following section.

²Here, we represent the empirical models as empirical tables. Addition and scalar multiplication are then interpreted as standard matrix operations, where the empirical tables are treated as matrices.




Figure 3: The PR prism schema and its adjective modifier instance.

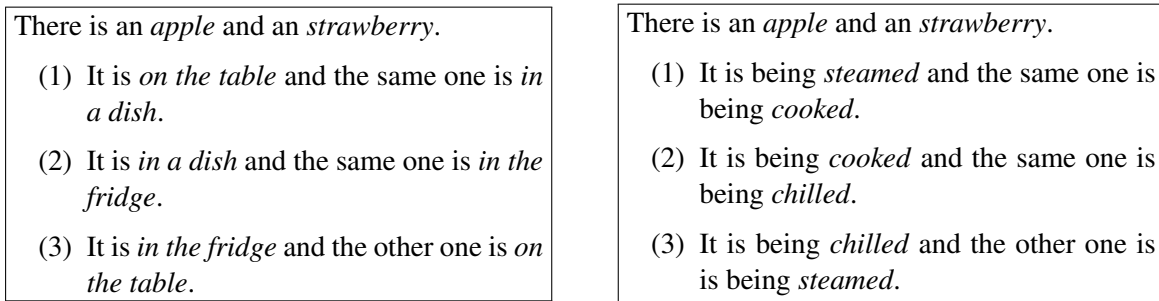


Figure 4: Examples of the PR prism schema with verbs (left) and preposition modifiers (right)

5 Possibilistic Examples

The construction used in a previous work on meaning ambiguities [18] was inspired by the Bell/CHSH scenario in quantum physics. However, the Bell/CHSH scenario is not minimal so we considered the minimal scenario with only 3 observables instead of 4. In our anaphoric setting, the set of possible interpretations is dependent on the ambiguous anaphora, instead of a fixed set of interpretations in the case of meaning ambiguities. This poses a difficulty in obtaining probabilities through a corpus. So we first focus on possibility instead of probability. It is much easier to determine if it makes sense for a word to be the referent of an anaphora than to determine its likelihood. We constructed a schema (Figure 3) that is modelled by the PR prism on the possibilistic level.

In the schema, O_1 and O_2 are two noun phrases as the candidate referents; X_1, X_2, X_3 are three modifiers commonly used to act on O_1, O_2 . The X_i 's are the observables of the scenario.³ Statement (1) and (2) above ensure that the modifiers X_i refer to the same referent, thus resulting in perfect correlation (parallel edges on the bundle diagram). Statement (3) ensures that the modifiers refer to different referents, thus resulting in perfect anti-correlation (crossing edges on the bundle diagram). The schema is constructed such that it is minimal and can immediately be modelled by the PR prism to ensure strong contextuality. For other examples using the same pair of nouns but with instead verbs or prepositional modifiers, see Figure 4. Other types of modifier are dealt with similarly.

6 Probabilistic Examples

We considered possibilistic models in the previous section. In this section, we propose a method for defining probability distributions for schemas such as the one considered in the previous sections.

³The ambiguous anaphoric words in the schema are *it* and *one*. We acknowledge that it is controversial to treat the modifiers X_i , instead of the ambiguous words, as observables. The construction of a more natural sounding schema is left for future research.

We form a probabilistic model through a contextualised language model such as BERT [7], which predicts a masked word (i.e. a blank space) in a sentence. Intuitively speaking, BERT uses the sentence as the context to generate a contextualised word embedding for the masked word, which is then measured for similarity against every word in the vocabulary.

For example, given a sentence: The goal of life is [MASK] ., BERT predicts the most likely word in the place of [MASK] . Moreover, BERT assigns a probability score to every word in the vocabulary.

The top 5 candidate words predicted by BERT and their probability scores are shown below.

<i>probability scores</i>	life	survival	love	freedom	simplicity	...
Probability	0.1093	0.0394	0.0329	0.0300	0.0249	...

We choose to use BERT because it has been providing improved baselines for many NLP tasks. In the following, we will demonstrate how we used BERT to define a probabilistic model for every schema considered in the previous section.

Consider the *apple-strawberry* example of Section 5. To measure a context, we replace the pronoun *It* in the sentence with The [MASK]. In practice, we feed the following 3 sentences separately to BERT:

There is an apple and an strawberry. The [MASK] is red and the same one is round.

There is an apple and an strawberry. The [MASK] is round and the same one is sweet.

There is an apple and an strawberry. The [MASK] is sweet and the other one is red.

BERT will then produce, probabilities $P_i(\text{apple})$ and $P_i(\text{strawberry})$ for the i -th sentence shown above. As BERT gives a probability score to every word in the vocabulary which sum to one, it is almost impossible that $P_i(\text{apple}) + P_i(\text{strawberry}) = 1$. We therefore normalise them by the following map⁴:

$$P_i(\text{apple}) \mapsto P_i(\text{apple}) / (P_i(\text{apple}) + P_i(\text{strawberry}))$$

$$P_i(\text{strawberry}) \mapsto P_i(\text{strawberry}) / (P_i(\text{apple}) + P_i(\text{strawberry}))$$

We will then use the normalised probabilities to construct a PR-like model with empirical table:

	(apple, apple)	(apple, strawberry)	(strawberry, apple)	(strawberry, strawberry)
(red, round)	$P_1(\text{apple})$	0	0	$P_1(\text{strawberry})$
(round, sweet)	$P_2(\text{apple})$	0	0	$P_2(\text{strawberry})$
(sweet, red)	0	$P_3(\text{apple})$	$P_3(\text{strawberry})$	0

It should be obvious how this procedure can be used on other examples of the schemas we considered in the last section. Notice that such an empirical model is non-signalling only if $P_i(\text{apple}) = P_i(\text{strawberry}) = 0.5$ for all i . It is therefore very unlikely that the model is non-signalling. To determine whether a signalling model is contextual, we use the inequality criterion of Equation (4). Recall that the CF of a PR-like model is always 1. Also, all the examples we considered in this paper have 3 contexts, i.e. $|\mathcal{M}| = 3$. Thus, to tell if such a model is contextual, we just need to check if $\text{SF} < \frac{1}{6}$.

As the criterion is actually quite strict, it is unlikely for any model constructed in this way to be contextual. We therefore need to create plenty of examples and to be strategic in the way we construct them. Equation (4) for PR-like models indicates that we need to make the probabilities as balanced as possible to make SF small. For that, we first fix two semantically similar nouns or noun phrases. Then, we ask BERT to associate them with frequently used modifying adjectives, verbs and prepositional phrases. As a result, we can ensure that the probabilities for the masked word given by BERT will be relatively balanced and thus minimising signalling in the model. The examples that produce contextual empirical models are presented in the proceedings subsections.

⁴The normalisation here is equivalent to limiting the vocabulary to just `apple` and `strawberry` when BERT computes the probability scores.




Figure 5: The distribution of signalling fractions of the models constructed with adjective modifiers.

6.1 Adjective Modifiers

We considered 11 pairs of similar noun phrases with between 3 and 18 candidate adjectives respectively. A model is constructed by picking a triple of adjective modifiers as the observables from the list of adjectives shown in the Appendix. This data generated 11,052 empirical models, of which 350 were contextual. Out of the 11 noun pairs considered, (cat, dog), (girl, boy) and (man, woman) produced models that are contextual. See below for the empirical tables of 2 examples of the contextual models we found.

(1)	(cat, cat)	(cat, dog)	(dog, cat)	(dog, dog)
(good, young)	0.4941	0	0	0.5059
(young, small)	0.4536	0	0	0.5464
(small, good)	0	0.5718	0.4282	0

(2)	(girl, girl)	(girl, boy)	(boy, girl)	(boy, boy)
(young, small)	0.5711	0	0	0.4289
(small, little)	0.5655	0	0	0.4345
(little, young)	0	0.5280	0.4720	0

Figure 5 is a histogram of the distribution of signalling fractions of the models constructed using the adjective modifiers considered. One can see that the majority of the model constructed are non-contextual and that the distribution skews towards greater SF.

6.2 Verb Phrases

We considered 2 pairs of similar noun phrases with 8 and 9 candidate verbs respectively, see the table in Appendix for details. This data generated 1,680 empirical models, of which 84 were contextual. For instance, the empirical table of the (apple, strawberry) - (sold, eaten, chilled) contextual model is presented below. The histogram of signalling fractions of the models constructed here is shown in the left panel of Figure 6.

	(strawberry, strawberry)	(strawberry, apple)	(apple, strawberry)	(apple, apple)
(sold, eaten)	0.4587	0	0	0.5413
(eaten, chilled)	0.5621	0	0	0.4379
(chilled, sold)	0	0.4416	0.5584	0




Figure 6: The distributions of signalling fractions for verbs (left) and prepositions (right).

6.3 Prepositional Phrases

We considered 2 pairs of noun phrases with 3 and 6 prepositional phrases respectively, see the Appendix. This data generated 252 empirical models, and we found two contextual models for each noun pair. For an example empirical see below; see the right panel of Figure 6 for the distribution of signalling fractions. Here *strawberry* is abbreviated as *strawb.* in the interest of space.

	(apple, apple)	(apple, strawb.)	(strawb., apple)	(strawb., strawb.)
(on the table, in the fridge)	0.5591	0	0	0.4409
(in the fridge, in a dish)	0.5640	0	0	0.4360
(in a dish, on the table)	0	0.4778	0.5222	0

7 Conclusions and Future Work

Coreference resolution is, amongst other Natural Language Processing tasks, facing the challenge of ambiguities. Instances of this task require extra resources such as context and world knowledge. In this paper, we focused on anaphoric coreference relations and the role of context. We showed how realistic contextual sheaf theoretic models of ambiguous data arising from quantum-inspired scenarios [16] can be used to model these examples. We developed a schema that produces possibilistic contextual models analogous to the PR Box. We mined probabilities for the instances of this schema using the BERT neural language model. Our computations showed that it is possible to find possibilistic as well as probabilistic contextual examples in natural language data, with only a handful of noun phrases and their modifiers. Future works include applying a similar methodology to coreference relations such as indefinite and definite noun phrases, quantifier scope, and situations requiring world knowledge, e.g. the Winograd Schema Challenge [13].

References

- [1] Samson Abramsky, Rui Soares Barbosa & Shane Mansfield (2017): *Contextual Fraction as a Measure of Contextuality*. *Physical Review Letter* 119, p. 050504, doi:10.1103/PhysRevLett.119.050504.
- [2] Samson Abramsky & Adam Brandenburger (2011): *The sheaf-theoretic structure of non-locality and contextuality*. *New Journal of Physics* 13(11), p. 113036, doi:10.1088/1367-2630/13/11/113036.
- [3] Janet Anders & Dan E. Browne (2009): *Computational Power of Correlations*. *Physical Review Letter* 102, p. 050502, doi:10.1103/PhysRevLett.102.050502.

- [4] Marco Baroni, Raffaella Bernardi & Roberto Zamparelli (2014): *Frege in Space: A Program for Compositional Distributional Semantics*. 9, p. 241–346, doi:10.33011/lilt.v9i.1321.
- [5] John S. Bell (1964): *On the Einstein Podolsky Rosen paradox*. *Physics Physique Fizika* 1, pp. 195–200, doi:10.1103/PhysicsPhysiqueFizika.1.195.
- [6] Bob Coecke, Mehrnoosh Sadrzadeh & Stephen Clark (2010): *Mathematical Foundations for a Compositional Distributional Model of Meaning*. doi:10.48550/arXiv.1003.4394.
- [7] Jacob Devlin, Ming-Wei Chang, Kenton Lee & Kristina Toutanova (2019): *BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding*. In: *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, Minneapolis, Minnesota, pp. 4171–4186, doi:10.18653/v1/N19-1423.
- [8] Ehtibar N. Dzhafarov, Janne V. Kujala, Víctor H. Cervantes, Ru Zhang & Matt Jones (2016): *On contextuality in behavioural data*. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 374(2068), p. 20150234, doi:10.1098/rsta.2015.0234.
- [9] Albert Einstein, Boris Podolsky & Nathan Rosen (1935): *Can Quantum-Mechanical Description of Physical Reality Be Considered Complete?* *Phys. Rev.* 47, pp. 777–780, doi:10.1103/PhysRev.47.777.
- [10] John R Firth (1957): *A synopsis of linguistic theory, 1930-1955*. *Studies in linguistic analysis*.
- [11] Zellig S. Harris (1954): *Distributional Structure*. *WORD* 10(2-3), pp. 146–162, doi:10.1080/00437956.1954.11659520.
- [12] M. Howard, J. Wallman, V. Veitch & J. Emerson (2014): *Contextuality supplies the 'magic' for quantum computation*. *Nature* 510(7505), pp. 351–355, doi:10.1038/nature13460. arXiv:1401.4174.
- [13] Hector J. Levesque, Ernest Davis & Leora Morgenstern (2012): *The Winograd Schema Challenge*. In: *Proceedings of the Thirteenth International Conference on Principles of Knowledge Representation and Reasoning, KR'12*, AAAI Press, p. 552–561.
- [14] Robin Lorenz, Anna Pearson, Konstantinos Meichanetzidis, Dimitri Kartsaklis & Bob Coecke (2021): *QNL in Practice: Running Compositional Models of Meaning on a Quantum Computer*, doi:10.48550/arXiv.2102.12846.
- [15] S. Mansfield & E. Kashefi (2018): *Quantum Advantage from Sequential-Transformation Contextuality*. *Physical Review Letters* 121(23), pp. 1–8, doi:10.1103/PhysRevLett.121.230401. arXiv:1801.08150.
- [16] Damian Markham Pierre-Emmanuel Emariau, Shane Mansfield (2022): *Corrected Bell and Non-Contextuality Inequalities for Realistic Experiments*. in preparation.
- [17] R. Raussendorf (2013): *Contextuality in measurement-based quantum computation*. *Physical Review A - Atomic, Molecular, and Optical Physics* 88(2), pp. 1–7, doi:10.1103/PhysRevA.88.022322. arXiv:0907.5449.
- [18] Daphne Wang, Mehrnoosh Sadrzadeh, Samson Abramsky & Victor Cervantes (2021): *On the Quantum-like Contextuality of Ambiguous Phrases*. 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. 42–52.
- [19] Daphne Wang, Mehrnoosh Sadrzadeh, Samson Abramsky & Víctor H. Cervantes (2021): *Analysing Ambiguous Nouns and Verbs with Quantum Contextuality Tools*. *Journal of Cognitive Science* 22(3), pp. 391–420, doi:10.17791/jcs.2021.22.3.391.

8 Appendix

8.1 Data for adjectives

noun pair	adjective modifiers	models	contextual models
cat, dog	cute, furry, lovely, friendly, sweet, big, small, house, young, large, wild, dead, thirsty, hungry, good, gray, black, little	9792	344
girl, boy	little, beautiful, young, pretty, small, baby, teenage	420	1
man, woman	young, dead, little, big, strange, beautiful, tall	420	5
strawberry, apple	round, red, sweet, sour, rotten	120	0
daisy, marigold	yellow, small, beautiful, everywhere	48	0
daisy, sunflower	yellow, small, beautiful	12	0
moth, butterfly	winged, colorful, light, beautiful	48	0
cucumber, courgette	green, long, juicy, tasty	48	0
dolphin, porpoise	grey, wet, slippery, slim	48	0
potato, yam	orange, starchy, healthy, big	48	0
car, bus	fast, sturdy, safe, heavy	48	0

8.2 Data for verbs

noun pair	verbs	models	contextual models
strawberry, apple	sold, bought, washed, eaten, rotten, cooked, chilled, steamed	672	9
cat, dog	fed, chased, watched, held, hunted, touched, pet, bathed, cleaned	1008	75

8.3 Data for prepositional phrases

noun pair	prepositional phrase modifiers	models	contextual models
apple, strawberry	on the table, in a dish, in the fridge	12	1
boy, girl	from the town, at the school, near the shop, on a bus, across the street, in the city	240	1