# A Metaphysical Composition

Nicola Pinzani

March 23, 2018

## Introduction

This is going to be a story rather that a more formal and hard core scientific essay. It aims to convey to the reader a feeling of how could a theory of meaning emerge from an apparently sterile branch of mathematics like category theory. Building on the existing literature on the topic, which is now an active and fruitful research field, we want to guide the reader trough the process of developing new approach to model the meaning of certain composition of words.

The quest to find the way to be able to even speak about the emergence of meaning is one of the most ambitious and beautiful tasks of science and it manages to connect and link seemingly distant areas such as mathematics, cognitive science and philosophy. In this context, the mathematical language is one of the best tools to explore the subtle dividing line between the rational mechanism and the mysterious meanders of cognitive processes, doing mathematics in fact is often a study of emergent properties, of the logical interactions between epiphenomena. More than a language, category theory is the formalisation of mathematical serendipity, of the beautiful analogies that stimulate the intellect and the curiosity of mathematicians. Abstract in nature but surprisingly an excellent practical tool for conceptualisation and model building, Category theory has often been employed in various branches of sciences, from biology to the analysis of cognitive processes in neuroscience [2].

It has to be said, that the question of the emergence of meaning is not only a domain of science and philosophy; as in many other cases it has been foreseen and anticipated by art and the artistic exploration. Art, as an intellectual production of the humankind is probably in the first line in the millenary battle of breaking the limit between form and substance. The sparks that made this mini project start, roots in the artistic exploration that took place at the beginning of the 20th century. It is fascinating to observe how sacrificing shape colour, playing with the prospective and composition can preserve, maybe even amplify this strange phenomenon called meaning. I therefore decided to honour this inspiration making use of paintings as concrete examples in the process of drawing up this mini project. In this essay we will take a more discursive approach and describe the principal mechanisms at play in the Compositional Distributional Models of meaning, it is by far not comprehensive but we will point out the sources for useful reference. Everything will rotate around the application of monoidal categories to model the meaning in natural language processing described for the first time in [5].

In the first part we provide a very simple example using vector spaces to show the mechanisms at work. In the second part of the project, the most substantial, we will describe and propose an application of the convex relation to model spatial relationship, the ambition is to be able to provide a simple and flexible way to model the verb *is on*. The aim is to introduce the possibility of using the convex relations framework to address the problem of accessing the meaning of simple sentences related to spatial position. To keep the examples as simple as possible we will always refer to two dimensional and static scenes. In the last section we will briefly explore the possibility of expanding the model using the generalised relations discussed in [9] and provide some indication of possible further work in this direction.

## A (very) brief practical introduction to compositional model of meaning

In this section we will provide a sumnmary of the principles and philosophy behind the compositional models of meaning. In [5], Coecke, Sadrzadeh and Clark provide the mathematical foundations of a categorical approach to Natural Language Processing. This approach aims at unifying the symbolic and distributional theories of meaning; in order to establish the bond between the syntactic and semantic 'understanding' they introduced a way to model both using compact closed monoidal categories, one for the grammatical structure of the sentence and one modelling the meaning of the words. The category that where originally used by the authors were vector spaces, to model word meaning, and the Lambek's pregroup grammar which assigns grammatical structure to sentences [8]. The idea is therefore to individuate domains of meaning for each word and then let them interact between each other by means of a grammatical structure. It should be said that the vector space model of lexical meaning is alone widely used in many applications, it aims to extract meaning from the frequency of a word appearance in a given context. For a in depth literature review of the topic we refer to [3]. It can be relatively easy to observe that although the vector space model strikes with its simplicity and flexibility, alone it lacks some of the fundamental properties that we could expect to be important in constructing a model of meaning. For example, although it may capture the context of a word well, in doesn't provide a specific answer to the problem of making the meaning of words interact.

When constructing a formal model of meaning, the role of the grammar itself may be controversial. It is possible to argue that it doesn't play an explicit role in defining the real meaning of a sentence, however there are circumstances where the grammatical understanding seems to be essential. It resolves ambiguities and help us to place concepts and meaning in a structured, coherent and more or less universal spatial and temporal context; everyone who ever tried to learn a second language unequivocally often appeals directly to grammar to understand to flow of meaning in a sentence. If one doesn't perceive this importance in native speakers, it might be because these structures are eradicated in the subconscious mechanisms inherited from childhood. We will now present an application of the approach using vector spaces and the pregroup grammar. We will only mention the fundamental mechanisms at play, we will soon depart from the vector space model to introduce different more expressive and general spaces that could be used to model concepts and their interactions.
As we have stated above, the approach can be summarised by the following diagram:

$$\begin{tikzcd} & language \arrow[dl, "meaning"'] \arrow[dr, "grammar"] & \\ \mathbf{Fvect} & \mathbf{Fvect} \times \mathbf{P} \arrow[l] \arrow[r] & \mathbf{P} \end{tikzcd}$$

Both meaning and grammar are here thought as compact closed structure that can be projected out of a language. To recover and formalise meaning we can therefore choose to consider the product of the two categories. We now have that every morphisms in $\mathbf{Fvect} \times \mathbf{P}$ exists only if there is a relevant morphism in the structure modelling grammar, it has therefore to be mediated by a chain in the (in our case posetal) category which models grammar. In this way the grammatical structure manages to restrict the unreasonable freedom of the **FVect** category. It is useful to think about it as establishing what is a flow of information between different parts of the sentence. We will not discuss the details of pregroups and compact closed monoidal categories, this can be found respectively in [8] and [4], we start providing a simple example of its direct use.
Assume that we have a simplified context where there are 3 men looking at each other:

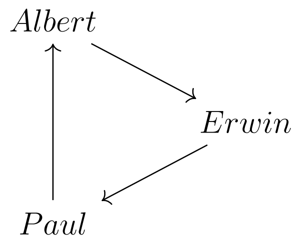


We can now try using the type of analysis exemplified in [5] to provide a model which is able to discern the truth value of sentences of the type:

*Paul is looking at Albert.*

Using the graphical notation for compact monoidal categories, we can translate the sentence using cups and caps thus establishing the way the information between words interacts in a simple sentence with a transitive verb.

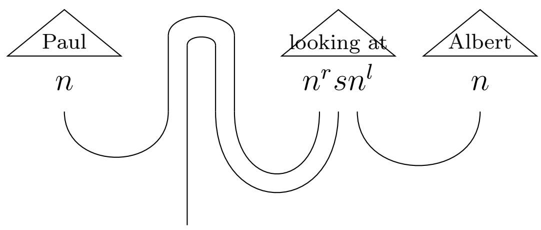


There are a couple of things to be noticed about the diagram, firstly we have decided that the role of the the auxiliary verb *is* is mainly to transfer information from the subject to the verb *look* and it doesn't influence the meaning of the sentence, it is just a grammatical requirement of the present continuous. To simplify the construction we have also assumed that the preposition *at* has also an exclusively structural role. This simplifying assumption however hides the fact that a way to satisfactorily treat preposition has not been yet provided, we will return to this point in later sections. Using the graphical calculus for monoidal closed categories, the relation between the words of the sentence becomes

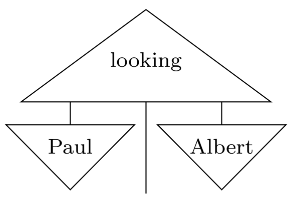


We are now ready to play with this simplified toy example. We express the three individuals as vectors spanned by $\{m_i\}_i$ which is the set of all men. Of course we understand the reader's feeling that expressing each men as a linear combination of every human being seems to be an oversimplification, we stress again that this is only to exemplify the basic technical foundations.

$$\overrightarrow{Paul} = \sum_i \alpha_i \overrightarrow{m_i} = m_6$$

$$\overrightarrow{Erwin} = \sum_i \alpha_i \overrightarrow{m_i} = m_3$$

$$\overrightarrow{Albert} = \sum_i \alpha_i \overrightarrow{m_i} = m_2$$

The three men live in the noun space $N$ which is spanned by the basis vector described above. The verb looking will have to be of the type $N \otimes S \otimes N$, since the type of a transitive verb in a pregroup grammar is $n^l s n^r$. A choice of sentence space $s$ will depend on the kind of answer we want to obtain by the model. The simplest analysis can be performed taking the sentence space to be one dimensional. We take $S = \langle v \rangle$, thus the verb will be expressed as $\overrightarrow{looking} = \sum_{ij} \overrightarrow{m_i} \otimes \alpha_{i,j} \overrightarrow{s} \otimes \overrightarrow{m_j}$. Where we consider

$$\alpha_{ij} = \begin{cases} 0 \text{ if } i \text{ is looking at } j \\ 1 \text{ otherwise} \end{cases}$$

In the particular example given above the verb takes the form of the potential superposition of all related " actions" taking place in the scene:

$$\overrightarrow{looking} = \overrightarrow{m_3} \otimes \overrightarrow{s} \otimes \overrightarrow{m_6} + \overrightarrow{m_6} \otimes \overrightarrow{s} \otimes \overrightarrow{m_2} + \overrightarrow{m_2} \otimes \overrightarrow{s} \otimes \overrightarrow{m_3}$$

Therefore returning to the sentence *Paul is looking at Albert*:

$$\sum_i \overrightarrow{m_6} \otimes (\overrightarrow{m_i} \otimes \alpha_{i,j} \overrightarrow{s} \otimes \overrightarrow{m_j}) \otimes \overrightarrow{m_3}$$

The canonical morphisms defining the *cups* and *caps* in **FHilb** are given by

$$1 \mapsto \sum_i \overrightarrow{e_i} \otimes \overrightarrow{e_i}$$

and

$$\sum_{ij} c_{ij} \overrightarrow{v_i} \otimes \overrightarrow{w_j} \mapsto \sum_{ij} c_{ij} \langle v_i | \, | w_j \rangle$$

Therefore applying the reductions we get that *Paul is looking at Albert* $= \overrightarrow{s}$ and it is therefore true while for example

*Paul is looking at Albert*

returns the 0 vector and may therefore be considered false. Similarly also

*Joe is looking at Richard*

where $\overrightarrow{Joe} = m_{42}$ and $\overrightarrow{Richard} = m_9$ is false. We have briefly shown what is the main idea behind the compositional distributional model of meaning presented in [5]. When we think about our cognitive processes I believe that the reader, like myself, has a certain difficulty to connect our experience of cognition with the structure of a vector space. More general models are needed, using category theory we can try to construct a space which is more intertwined with our traditional cognitive experience and preserves all those aspects that are technically needed to continue our analysis.

## Describing simple spatial relationships using ConvexRel

Using vector spaces is not the best approach if the intent is to capture a more natural way of understanding the composition and interaction of concepts. In fact this is an old field fathomed by cognitive science. Peter Gärdenfors introduced the concept of *conceptual spaces* i.e regions of the cognitive space representing different concepts. In describing the fundamental structure of the conceptual spaces, he identifies as a key property, convexity. If two points in the space are representatives of the color *red*, everything in between should also be a representative of the same colour. Constructing a framework where such a constraint holds is particularly amenable to a categorical analysis, in [1] the authors present a categorical description of generalised concepts spaces. The purpose of the rest of this project is to present a simple application of the models developed in [1] and [9].

### A (meta)phisical interaction

While I was trying to find a way to model the preposition *on*, I started thinking that the easiest way was to start by describing a scene where objects where lying on table, still and powerless, subject to the dictatorial power of gravity. Instead of creating ad hoc examples, I decided to make reference to another exploration of the concepts of shape, colour, spatial composition: the artistic exploration of the beginning of the 20th century. In particular there is a group of Italian metaphysics that wonted to explore the metaphysical essence of composition in a fundamental and abstract way, departing from the typical and naturalistic tradition of still painting. These paintings are in fact a perfect example for the kind of modelling that we want to perform. I will mainly consider a particular oil on canvas painted by Giorgio Morandi in 1920, the picture shows a very simple kitchen table with a couple of bottles and vases, some abstract geometrical objects, among them spheres that we can wish and will with an imaginative leap interpret as fruits.

Figure 1: *Natura Morta*, Giorgio Morandi (1920)

The first aim is to try to conceptually understand what could the prepositional verb *is on* actually mean. Generally there is a lot of literature about the meaning of the preposition [6, 7], Gärdenfors in [6] suggests that the meaning of a preposition doesn't depend exclusively on the spatial definition of the object. In [6, 7] they present a simple picture which captures the inadequacy of a purely spatial definition of the preposition *in*:

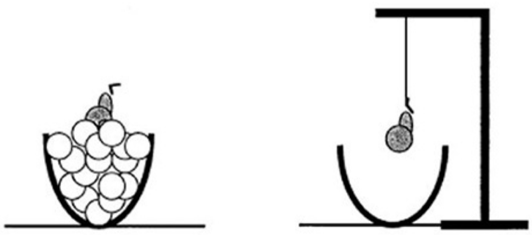

While in the first case we consider the pear to be in the bowl, removing the support between the bowl and the pear it's enough to convince us that the pear is not anymore inside the container. Gärdenfors goes on concluding that a correct exposition of the meaning of the preposition must therefore involve the concept of force. He proposes that the containment is based on a principle of dependance of the object contained exerted by the object that contains.

For example, moving a box containing a sphere certainly forces the sphere itself to change its position in space [6]. Even though this picture is appealing and it seems to fit quite well to our intuition, it seems to be excessive to involve in the picture complex physical formalisms such as vector fields as suggested in [6]. After all, one could Wittgensteinaly argue that our intuition on the interplaying forces is a consequence of experience, a consequence of the language we use to describe concepts and this may lead to a dangerous circularity. However we don't want to complicate the situation invoking physics or philosophical speculations. I believe that the simplest and most fruitful answer is to escape from the debate and decide to provide a simplified approach that may work for static pictures, scenes where we are unable, even in principle, to perform any kind of action or verification, where the viewpoint and the perspective is unique and unchanging. Our domain will therefore be exclusively the spatial location of objects, in which we assume a very simplified and naive force hierarchy. While the definition of the word *on* could be controversial, understanding it in isolation appears to be an unnecessary task, as it's meaning is almost always complemented by a verbal form which helps to specialise it. For example the sentence

*The bottle lies on the table*

implies a stronger condition than

*The bottle is on the table*

the syntagm *lies on* requires an additional constrain: to the sentence is often associated an idea of a further stability implied by horizontality. It is therefore certainly true that the general picture is much more complex then the one presented in this work. Our aim in fact, more then anything else, to provide an example of the flexibility of the approaches based on **ConvexRel** and the generalised relations with the hope of opening up a discussion of its use in modelling a spatial hierarchy.

We will start by doing something controversial, add a couple of objects and abstracting the abstract:

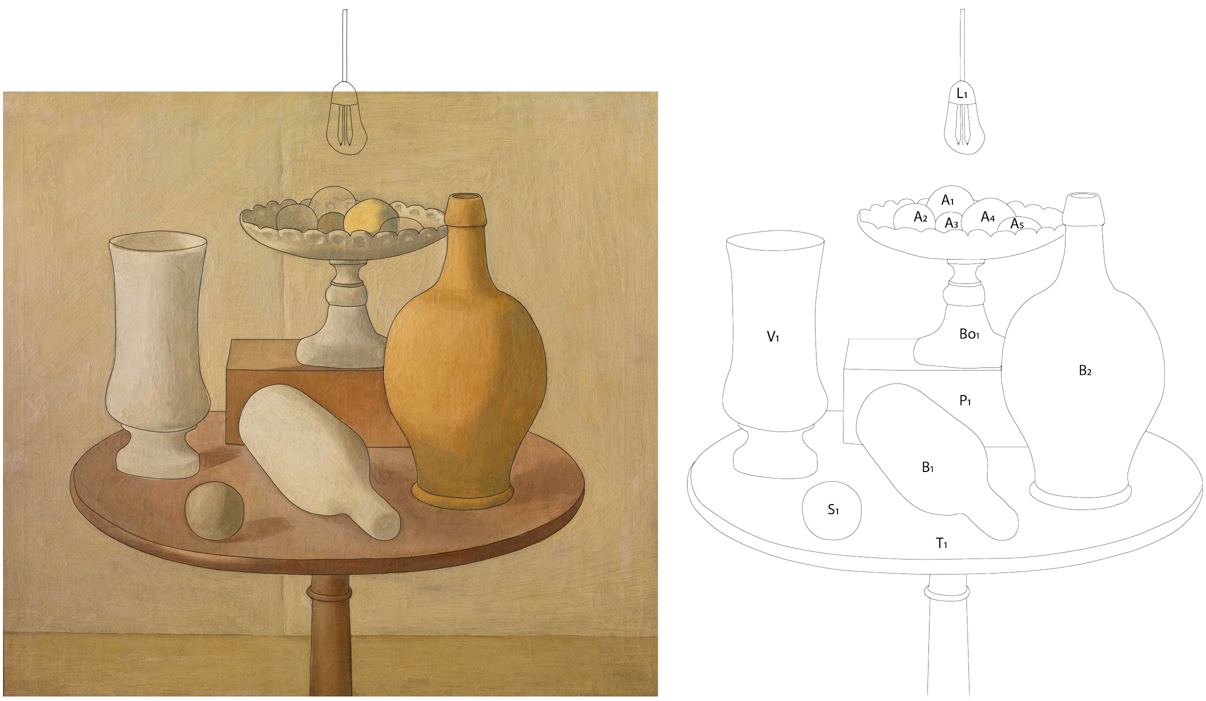


We will first focus on the verbal phrase *is on*. The aim is to find a relation in the context of **ConvexRel** which could be used to model simple instances of the use of that verb. The universe set $\mathcal{U}$ will be constituted of all the elements appearing on the scene. Note that here we have to make a choice about how many and what are going to be the *atomic* words, the fundamental objects composing the scene. This is a big but acceptable degree of ambiguity, it reflects the fact that the interpretation of the contextual phenomena is a subjective experience. A surgeon will certainly have a different set of atomic words and a different depth of understanding of a scene happening in an operating room from a mathematician. Once we have established this fundamental building blocks, we want to be able to describe the relationship between any possible subset of those 'objects', we will therefore define our relationship on the convex algebra induced by the powerset. In fact we show that the powerset structure with the union operation forms a convex algebra.

**Lemma 0.1.** Let $U$ be an arbitrary set, therefore $\mathcal{P}(\mathcal{U})$ has a convex algebra structure given by:

$$\sum_i p_i a_i = \bigcup_i \{a_i \mid p_i > 0\}$$

where $a_i \in \mathcal{P}(\mathcal{U})$.

*Proof.* A convex algebra is a function $\alpha : D(A) \to A$ satisfying the following conditions [1]:

$$\alpha(|a\rangle) = a$$

$$\alpha\left(\sum_{i,j} p_i q_{i,j} |a_i, j\rangle\right) = \alpha\left(\sum_i p_i \middle| \alpha\left(\sum_j q_{i,j} |a_{i,j}\rangle\right)\right\rangle\right)$$

The first condition is trivially satisfied, for the second, note that the coefficients that we use to form the formal sums do not have any meaning, we can therefore assume without loss of generality that all the non zero coefficients $p_i$ and $q_{i,j}$ are equal. Starting from the left hand side we have that

$$\alpha\left(\sum_{i,j} p_i q_{i,j} |a_{i,j}\rangle\right) = \bigcup_{i,j} \{a_{i,j} \mid p_i q_{i,j} > 0\}$$

while for the right hand side:

$$\begin{aligned}\alpha\left(\sum_i p_i \middle| \alpha\left(\sum_j q_{i,j} |a_{i,j}\rangle\right)\right\rangle\right) &= \bigcup_i \left\{\alpha\left(\sum_j q_{i,j} |a_{i,j}\rangle\right) \mid p_i > 0\right\} \\ &= \bigcup_{i,j} \{\alpha(|a_{i,j}\rangle) \mid q_{i,j} p_i > 0\} = \bigcup_{i,j} \{a_{i,j} \mid p_i q_{i,j} > 0\}\end{aligned}$$

□

To define the relation between subsets we define a tree which tries to capture an elementary notion of force relationship between objects. A fair objection would be to say that the knowledge of a tree like this already implies some sort of assumption about a spatial hierarchy, however we are interested in something different, given visual knowledge of the different objects on scene and their relative position, how can we verbally and dinyamically access such information in the context of natural language processing? An artificial intelligence for example sees forms, lines, however it doesn't a priori understand the scene if we don't define a way to unleash and access the meaning of this composition of shades lines and colours. For example we can synthesise the relative position of each object in Figure 1 using a directed tree $\mathcal{T}$. Each edge $e \in E(\mathcal{T})$ will be denoted by the tuple $e = (s(e), t(e))$, where $s(e)$ denotes the starting object and $t(e)$ the terminal object. For example returning to our still life, we can imagine the tree to look somehow similar to this one:

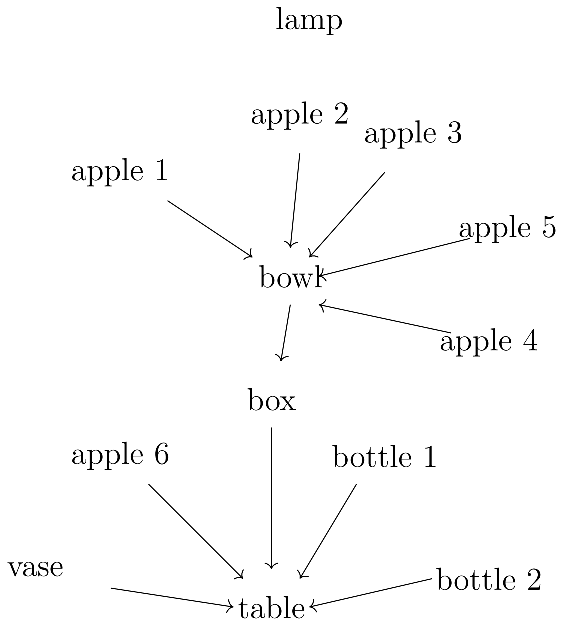


Note that we seek above everything else simplicity and flexibility. As such, we are only defining a notion of vertical dependence. We don't particularly care whether an object is on the left or on the right of another one, or whether the *bottle* is further away from the *table*. We believe this to be a faithful representation of the way our brain abstracts concepts. It seems to me that often independent hierarchies are in play with the aim to abstract some specific attributes. As such, properties that in certain contexts are irrelevant may just get discarded representing an additional degree of freedom, in our case the relative position of each branch of the tree. Now the aim is to assign a relation between the different subsets of the powerset using the tree constructed above. In order to do so, we have to provide some preliminary definitions:

**Definition 0.2.** Let $A \in \mathcal{P}(\mathcal{U})$ then the boundary of $A$ denoted $\delta A$ is given by

$$\delta A = \{t(e) \mid e \in E(A, A^c)\}$$

**Definition 0.3.** Let $A, B \subseteq \mathcal{P}(U)$ then we say that $A$ *dominates* $B$, denoted by $A \gg B$ if have that for every $b \in B$, either $b \in A$ or $b \in \delta A$ and $\delta A \subseteq B$.

The following picture explains the definition described above.

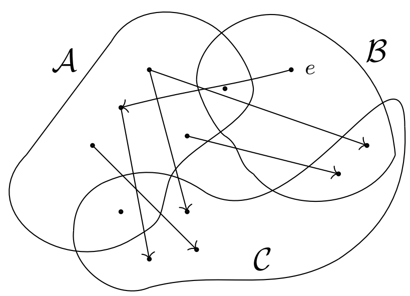


The set $A$ dominates the set $C$, however it does not dominate $B$ as the element denoted by $e$ is not in $\delta A$, morever $B$ doesn't even contain the entire boundary.

**Definition 0.4.** We now define the relation $R$ as follows:

$$R = \{(A, B) \mid A, B \in \mathcal{P}(\mathcal{U}), A \gg B\}$$

**Definition 0.5.** A **convex relation** of type $(A, \alpha) \to (B, \beta)$ is a binary relation $R : A \to B$ between underlying set that commutes with forming convex mixtures

$$(\forall i. R(a_i, b_i)) \implies R\left(\sum_i p_i a_i, \sum p_i b_i\right)$$

**Lemma 0.6.** The relation $R : \mathcal{P}(\mathcal{U}) \twoheadrightarrow \mathcal{P}(\mathcal{U})$ defined as above is a convex relation $(\mathcal{P}(\mathcal{U}), \cup) \twoheadrightarrow (\mathcal{P}(\mathcal{U}), \cup)$

*Proof.* Let $\{A_i\}_i$ and $\{B_i\}_i$ be families of sets such that $A_i \gg B_i$ for all $i$. Let $A = \bigcup_i A_i$ and $B = \bigcup_i B_i$. Consider an arbitrary element $b \in B$ then either $b \in A$ or $b \notin A$. assume $b \notin A$; since $b \in B$, there exists $B_k$ such that $b \in B_k$. We have that $A_k \gg B_k$ and $b \notin A$, it must therefore be the case that $b \in \delta A_k$ and $b \notin A$. Because $\delta A_k \backslash A \subseteq \delta A$, we get that $b \in \delta A$. The relation is therefore closed under taking unions and it respects the convexity of the underlying power set. □

The requirement of preserving convexity reflects the fairly intuitive fact that if have a lamp *on* the beside table and a book laying *on* the table on the other side of the room, then we can say that

*The Lamp and the book are on the tables*

It is shown in [1] that **ConvexRel** is compact closed, we are therefore able to use the full power of the graphical calculus when composing the words of a sentence. We can see the relation $R$ defined above as a morphism that goes from the monoidal unit $I$ to a convex subset of $\mathcal{P}(\mathcal{U}) \times \mathcal{P}(\mathcal{U})$, a perfect candidate for the description of a transitive verb. We can use graphical calculus and represent $R$ as follows:

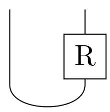


we can therefore start compose words in sentences, let us start with the simplest example, we consider the space of the noun $N$ to simply be the powerset of the universe set $\mathcal{U}$, this is already a big assumption as it supposes that we are potentially able to associate a name to every set in $\mathcal{P}(\mathcal{U})$ but what is the common noun which conveys the meaning of the noun clause *apple and box*? Is it fair to create a noun space which goes beyond the expressive powers of a language? Clearly a noun spaces doesn't exclusively contain grammatical nouns, but also sentences and clauses that behave like a noun. The prepositional verb *is on* is transitive, consistently with [1, 5] we have defined it to be a convex subset of $\mathcal{P}(\mathcal{U}) \times S \times \mathcal{P}(\mathcal{U})$ where $S$, the sentence space is just defined to be a singleton, which is isomorphic to its monoidal unit. Wiring up the sentence and performing the appropriate reductions, we will end up with a scalar in **ConvexRel** which represents a Boolean value related to the truth value of the sentence. This is how it looks graphically:

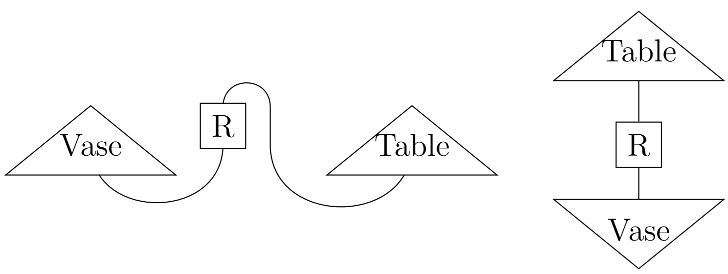

The way we described the relation may seem a bit strange at first sight but our choices are well motivated. For example another possible choice was to consider a subset to be *on something* if and only if its entire boundary is contained in a set, at first it seemed to me to be the most natural choice. A moment of thought however suggests that t is certainly not ideal, first it doesn't preserve convexity, secondly this would imply that we can show that the bottle is on $\{$*table*, *light bulb*$\}$, or that the light bulb hanging from the roof *is on* everything else at the same time. Looking at tree it is in fact clear that the set $\{$*table*$\}$ lies entirely in the boundary of the set $\{$*vase*$\}$. If you for example select all the singletons that are related in this way to *table*, Figure 2 emerges. There are many other subsets of elements which can be considered

Figure 2: A first layer of singletons related to $\{$*table*$\}$

to be on the table and in particular there is nothing particularly significative about being a singleton, except from the fact mentioned earlier that we need to choose some elementary structures in our description. By convexity, as shown in Lemma 0.6, an arbitrary union of such subsets also lies on the table. Moreover, consider the case where we have a "noun" $A =\{$*box*, *bowl*$\}\cup$ *fruits in the bowl*, where we denote by *fruits in the bowl* the set of spheres contained in the bowl. To denote this concept we do not have a specific word, we can use an auxiliary subordinate sentence that specifies the object, as it is however pretty much still 'something' and it is certainly on the table:

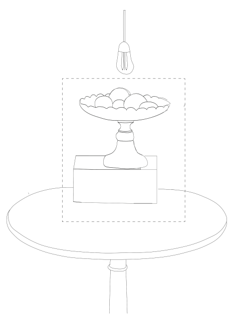

If we now evaluate the meaning of the sentence with respect to the space $\{true,false\}$ we get that the following:

*A is on the table*

is graphically represented by

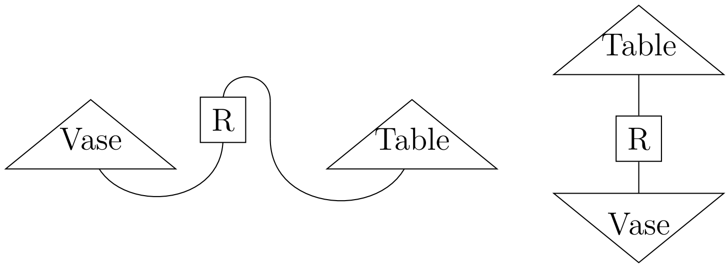


And the case of the composition $A$ shown in the picture it evaluates to true. This is because there is only one edge leaving the component of the graph represented by $A$ and this edge terminates in the only element of the set $\{table\}$.
If we want to express the subsets described above compositionally and access a bigger part of the space $\mathcal{P}(\mathcal{U})$ we could for example use the semantics of subject relative pronouns developed by Sadrzadeh, Clark and Coecke in [10].

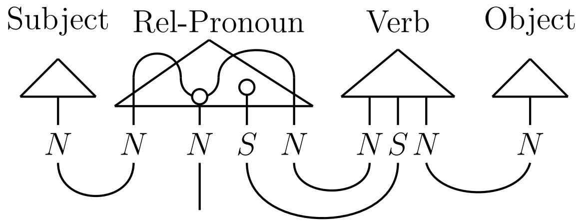


The diagram shows that the role of relative pronouns in the diagram is to 'return' all the generalise elements of *Subject* which are related by the verb to *Object*. The multi-wires, or spider that are here used, generalise the *cups* and *caps* notation allowing us to compare and therefore select the parts from different meaning spaces that overlap. The Frobenius structures in **ConvexRel**, allows to have multiple inputs and multiple outputs. A more detailed description can be found in [4], we provide here the formal definitions:

**Definition 0.7** (multi-wires)**.** A *multi-wire* is a relation $A \otimes A \ldots \otimes A \to A \otimes A \ldots \otimes A$ given by $\{((a, a \ldots a), (a, a \ldots a)) \mid a \in A\}$ we denote the multi-wire as follows:

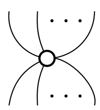

We can now compute the meaning of the sentence:

*Objects that are on the table*

We remind that our approach is contextual to the scene we are describing and we will consider the word 'Objects' to be the universe set $\mathcal{U}$ itself, 'are on' is modelled by the relation $R$. In our case, the morphism $R$ goes from $\mathcal{P}(\mathcal{U}) \times \mathcal{P}(\mathcal{U})$ to the monoidal unit $I$, therefore:

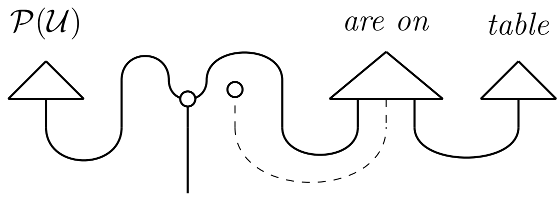

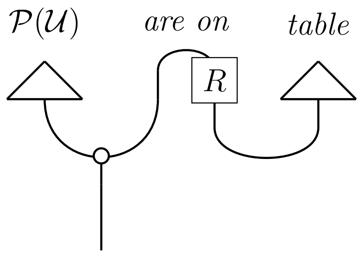


we see that by Definition 0.7 for an element in $A \in \mathcal{P}(\mathcal{U})$ there is a relation given by the white spider if and only if $A$ is related to $\{table\}$. This will result in a convex subset of the powerset, which are all the objects and the different possible combinations of objects that can be found *on* the table. Similarly we could for example represent the sentence:

*Objects that are on the box*

this will clearly contain the *bowl* and all the other subsets which have *box* as their boundary. For example the subset $S = \{$*bowl*,*apple 2*,*apple 3*$\}$ is considered on the box, while $\{$*apple 2*,*apple 3*$\}$ is not. When saying *Objects that are on the table* we may not immediately think about the subset $S$, on the other hand we certainly would if there existed a familiar expression denoting it. If the universe changed and the scene evolved to:

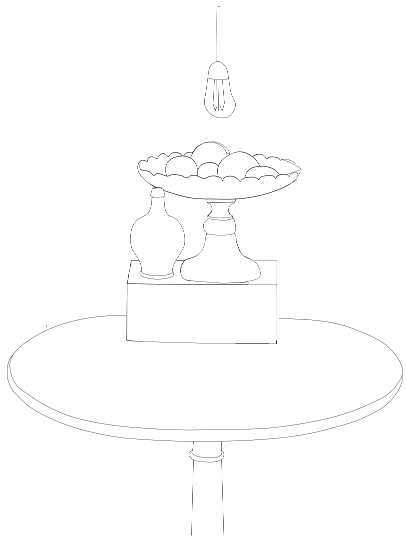

the sentence *Objects that are on the box* would now hold a different meaning, thus including the bottle which was initially on the table. Notice that the grammatically inexact sentence

*Objects that are on the bowl*

will in this context have the same meaning that we would attribute to the more correct sentence

*Objects that are in the bowl*

it is a case when the grammatically inexact sentence naturally refers and carries the meaning of the correct sentence without causing ambiguities. We may therefore consider in this particular example the verb *is in* to be modelled analogously to the verb *is on*. This clearly depends on the type of objects we are considering. In the case of a vase for example the difference appears certainly more striking:

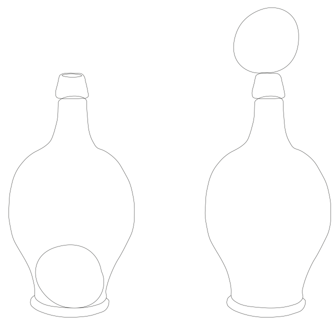

here certainly the two different prepositions reflects two very different ideas, we are not however concerned with solving this particular problem. Moreover, the fact wheter the apple inside the transparent vase is on the table is more of a philosophical question rather that linguistic one. From now onwards we may use the same relation $R$ to model both verbs, justified by the fact that in our context the difference between the two turn out to be merely grammatical. One of the most interesting aspect of using **ConvexRel** is to be able to appreciate the interaction between the different domains, being a monoidal category, we have the possibility via the tensor product to take the composite of different meaning spaces. The next section will extend the setting that we have presented by showing how it relates to another example of convex spaces presented in [1].

### Combination and interaction of spaces

Consider now the original picture by Giorgio Morandi, Figure 1, there are three apples: two green and one yellow shading into green. The sentence:

*The green apple that is on the table*

is specific and exact but in order to model it we have to be able to associate new attributes to the objects of our universe. The spatial relationship described above will therefore be only a part, a subspace of the individual noun spaces. In [1] Bolt et al. show that different convex relations living in the category **ConvexRel** can be composed together taking as the underlying set the cartesian product of the two sets and where the mixing operation is given by:

$$\sum_i p_i \,|(a_i, b_i)\rangle \mapsto \left(\sum_i p_i a_i, \sum_i p_i b_i\right)$$

Our new noun space could therefore be composed by

$$N_s \otimes N_{colour} \dots$$

where $N_s$ is a convex subset of $\mathcal{P}(\mathcal{U})$, a subset closed under union of sets and $N_{colour}$ is the convex RGB domain spanned by the sets $N_s$. For example we have used the word *Objects* to mean every possible combination of $\mathcal{U}$ for any possible colour combination:

$$\mathcal{P}(\mathcal{U}) \otimes \{(R, G, B) \mid (R \geqslant r), (G \geqslant g), (B \geqslant b)\}$$

it is not unreasonable for example to say that $r = g = b = 0$, as in a room with a table, a bunch of fruits, vases and ornaments, it is possible that the colours present in the scene span the entire RGB colour cube. On the other hand, if we want to stick to the Morandi's interpretation of the word, we would only need to consider a much smaller convex subset, the colour are in fact mostly confined inside a spectrum from pale yellow to dark brown. Now we colour the apples, let us say that there are two red apples on the table 3. First we define the subset *apples* to be

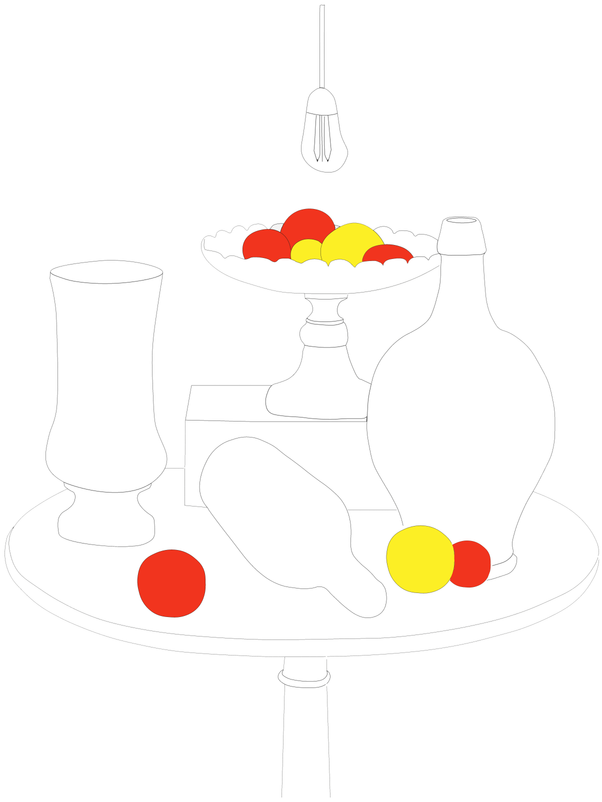

Figure 3: *Slightly more realistic apples*

given by the powerset, there are 8 apples in 3:

$$apples = \mathcal{P}(\{apple1, apple2, apple3, apple4, apple5, apple6, apple7, apple8\})$$

Each apple is itself a region of the noun space and lives in $N_s \otimes N_{colour}$ for example *apple 6* may be represented as:

$$(apple6, \{(R, G, B) \mid (R \geqslant r_0), (G \geqslant g_0), (B \geqslant b_0)\})$$

In *red apples* the word *red* is an adjective in [1] adjectives are of type $N \otimes N$, as they interact with a noun giving back a new noun with changed meaning. In [1] they classify different classes of adjectives which provide different kinds of interaction with the noun. According to this classification, in our case the adjective *red* is a subsective adjective, where the meaning of its interaction with the noun is a subset of the noun. To understand the adjective *red*, we have to understand the noun *red* , the interaction *adj noun* is then given by:

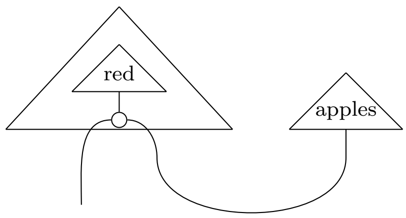


The noun *red* doesn't have any spatial and tangible connotation, it's convex domain in $N_s \otimes N_{colour}$ may therefore be given by:

$$(\{\varnothing\}, \{(R, G, B) \mid (R \geqslant 0.6), (G \leqslant 0.3), (B \leqslant 0.3)\})$$

while the adjective *red* is the morphism $(N_s \otimes N_{colour}) \otimes (N_s \otimes N_{colour}) \to N_s \otimes N_{colour}$ induced by the spider:

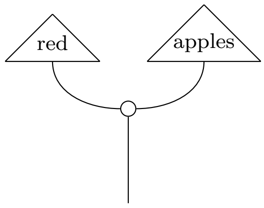


let $a$ be an element of the convex subset defined by the interaction graphically described above. Therefore $a$ is related with $a'$ in *apples* if $a = a'$ and the colour of $a$ lies inside the domain *red*. This shows that the structure of our adjective is

$$\{(x, x) \mid x_{colour} \in red_{colour}\}$$

Therefore:

*Red apples that are on the table*

Can be graphically presented as

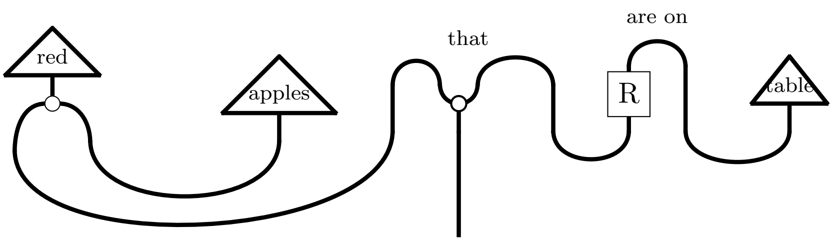


Let us how the meaning of this sentence is therefore calculated: first the adjective picks out all the subsets of the apples that can be considered to be red, we will call this subset $R \subseteq \mathcal{P}(\mathcal{U})$. $R$ is a convex subset of $\mathcal{P}(\mathcal{U})$, in the union of two subsets of red apples is again subset of red apples in *red apples*.

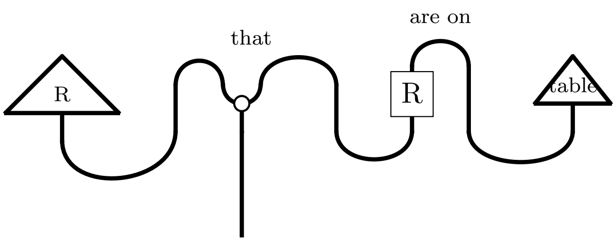


Now the pronoun transfer the information about the subject o the verb, the other white spider returns all the subsets of $\mathcal{P}(\mathcal{U})$ that are both in $R$ and related to *table*. In particular note that the result is a set of subsets, in fact this is a consequence of the choices of defining apples as

$$apples = \mathcal{P}(\{apple1, apple2, apple3, apple4, apple5, apple6, apple7, apple8\})$$

this is in line with the generality of *Red apples that are on the table*, it could refer to any subset of red apples. It is somehow different the meaning of the sentence *The red apples that are on the table*. In this case we only have one specific subset in mind. Does it mean that we could have started defining apples simply as $apples = \{apple1, apple2, apple3, apple4, apple5, \ldots\}$ unfortunately not, and this is because even though the singletons form (trivial) convex subsets of $\mathcal{P}(\mathcal{U})$ we see that we are then unable to separate and isolate elements. It provides a “topology” which is to coarse to be expressive.

### What about the light bulb?

In some of our pictures we shown there was a light bulb disturbingly hanging from an imaginary ceiling. What was that about? It was probably about the fact that meaning doesn't always emerge with perfect clarity either as black or white. We need to take in consideration endless shades which cast a worrying shadow on anything that has the aim to explain and describe the emergence of meaning. In this section we will try to expose a possible answer to this problem, integrating the approach given above with the generalised relations proposed in [9]. There is a strange period in the artistic production of Giorgio De Chirico, contemporary of Morandi, when he started to draw bananas piled up on straight surface, for example in *Le reve transforme*, The Transformed Dream, 1913:

Figure 4: *Le Rêve Transformé*, Giorgio De Chirico (1913)

No one would have any kind of trouble saying that the Bananas are on the floor, in fact so they are. What about the topmost banana? Is it still lying on the floor? Probably not, its lying on other bananas, however it is still possible to say that *Topmost banana lies on the floor* is much more true than the sentence *The light bulb lies on the table*. How could this fundamental difference be formally captured? This problem could be already addressed using the standard

Figure 5: *L'Incertitude du poéte*, Giorgio De Chirico (1913)

relational model provided before. Let us see what happens if we iterate the application of the relation $R$, consider $R^2 = R \circ R$:

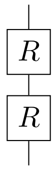


now we see that the subset $\{table\}$ becames related to the subset $\{bowl\}$. Iterating $R$ we eventually relate to the object *table* more and more objects but we will never manage to relate the light bulb. This captures the idea, that the concept of being on the table somehow carries trough objects or subsets of objects that are *on* top of each other and that there is a natural notion of weak transitivity of the truth value. In the case of our bananas in the oneiric world of Giorgio De Chirico, we might describe the relationship between the different fruits as an inextricable net, something like (note that we have ignored arrow for simplicity, we are always talking of directed graphs):

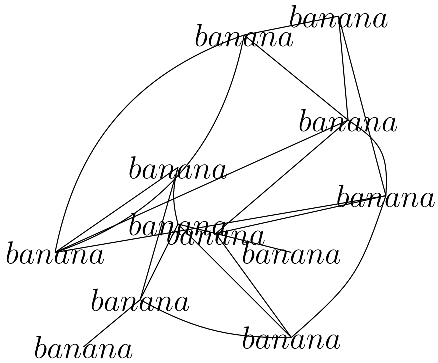


But now we can probably trace at least one path which witnesses the relation with the table for every fruit (even though this strictly speaking still depends on the direction of the edges, but this is simply a mathematical subtlety). The setting of the generalised relations has been expanded in [9], we will now provide a brief introduction to this generalisation, mostly referring to [9]. While in normal relations each element is either related or not related to any other, generalised relations allow us to describe the degree or quality of these relatedness. In [9] the authors show that this can be done without leaving the world of compact closed categories and therefore having the access to an elegant graphical calculus and therefore without compromising the flexibility that we learn to appreciate in the previous section. Normal relations assign to each couple of elements a boolean value, we will have the freedom to assign a value from an arbitrary *quantale*

**Definition 0.8** (Quantale [9])**.** A *quantale* is a join complete partial order $Q$ with a monoid structure $(\otimes, k)$ satisfying the following distributivity axioms, for all $a, b \in Q$ and $A, B \subseteq Q$:

$$a \otimes \left[\bigvee B\right] = \bigvee \{a \otimes b \mid b \in B\}$$

$$\left[\bigvee A\right] \otimes b = \bigvee \{a \otimes b \mid a \in A\}$$

A $Q$ relation is therefore a function of the form $A \times B \to \mathcal{Q}$. Then they define the compact closed category $\mathbf{Rel}(\mathcal{Q})$ with composition of morphisms and identities given by

$$(S \circ R)(a, c) = \bigvee_b R(a, b) \otimes S(b, c)$$

and

$$1_A(a, b) = \bigvee \{k \mid a = b\}$$

Let us try for example to generalise our relations using the Lewvere quantale $\boldsymbol{L}$. The Lwevere quantale is given by the extended positive reals with reverse ordering, where the monoid operation is the addition and the joins are given by the minima. The setting we constructed before naturally fits inside this framework as well, we can define a new $Q$-relation $R'$, we as follows:

$$R' = \begin{cases} 1 \text{ if A >> B} \\ \infty \text{ otherwise} \end{cases}$$

Considering the quantale $\boldsymbol{L}$ to be our sentence space, we can now evaluate the value of the sentence

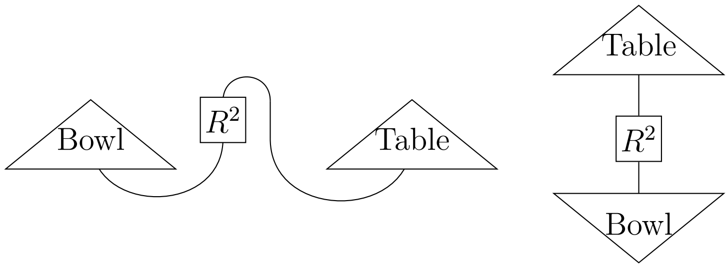


In this case:

$$R^2(\{bowl\}, \{table\}) = \bigvee_b R(\{bowl\}, b) + R(b, \{table\})$$

there are some possible intermediate subsets $b$, one of them is clearly the set $\{box\}$. Therefore

$$R^2(\{bowl\}, \{table\}) = 2$$

Clearly

$$R^2(\{light\ bulb\}, \{table\}) = \infty$$

In fact if we set $R^n = R^{n-1} \circ R$, we see that $R^n = \infty$ for all the possible value of $n$. Showing that it is in fact quite difficult to consider the light bulb to be in anyway *on* the table. It is clear however that connectedness it's not the only requirement, for example any of the apples in the bowl will never be connected to the singleton $\{vase\}$, independently on the number of iteration of $R$ that we are interested in taking, there are not subsets of $\mathcal{U}$ (except for the set itself) which dominate $\{vase\}$. There is nothing be considered to be on the vase. The existence of a directed path is also not a requirement, this is normally consistent with our intuition, there are however some problematic consequences: consider for example the subset $A = \{bowl, box\}$, there is clearly a directed path between $A$ and $\{apple\ 4\}$, however the reader will promptly realise that it is not possible to find a chain of dominating subsets starting from $\{apple\ 4\}$ and finishing in $A$. This is a problem which traces back to the requirements that relations are identified as convex spaces.
Our example is similar to the example provided in [9] about the description of relationships in an academic family tree. In fact, inspired by that paper we may also try to see what can be represented if we compose in a different way the relation $R$. So for example, consider the analogue of the sibling relationship described in [9]:

$$N = (R^{-1} \circ R)$$

two subsets $A$ and $B$ in the powerset are related by this morphism if there is a common set which both dominate, using our notation, it there exists $C \in \mathcal{P}(U)$ such that

$$A >> C, B >> C$$

for example in the following picture

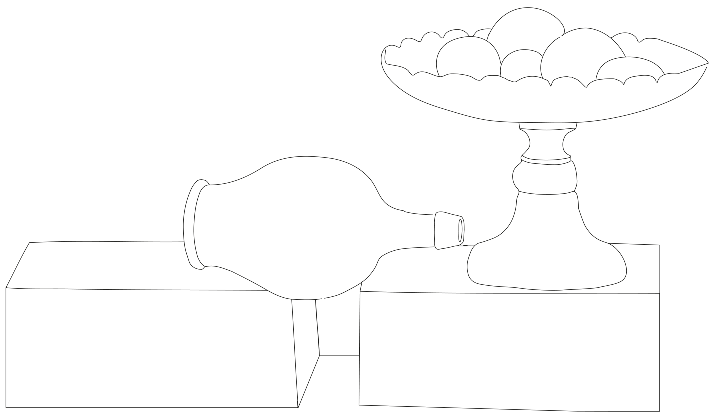

we have that any subset of the apples 'in' the bowl is related to any other subset of the apples. The bottle is for example related to the bowl.

$$\{bottle\} >> \{box\}$$

$$\{bowl\} >> \{box\}$$

we are therefore able to capture to certain extent the meaning of

*The bottle is next to the vase*

We could model the verbal form *is next to* considering the fact that two things lie next to each other if there is a common support and they both dominate a common subset. We could therefore define the verb+preposition *is next to* as the 'sibling' relation.

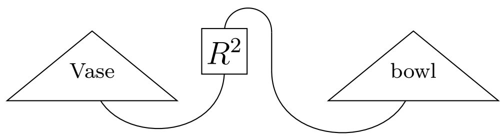


The flexibility of the model obtain using only a very simple relation is striking. Using generalised relation may seem at first just unecessary. However the notion that we have described before, how iteration witnesses the degree of truth of a sentence cannot be normally formalised categorically. The real gain over the normal concept of relations is that we can define, as done in [9], the closure of the relation $R$, which will still be a well defined relation with respect to the quantale:

$$\overline{R}(A,B) = \bigvee\{R^n(A,B) \mid n \geqslant 1\}$$

the sequence $R_n(A,B) = \bigvee\{R^i(A,B) \mid i \leqslant n\}$ is in fact increasing in the one-point compactification of $\mathbb{R}$, as such $\overline{R}(A,B)$ is certainly well defined. For example we have that:

$$\overline{R}(\{lamp\},\{table\}) = \infty$$

$$\overline{R}(\{apple\ 4\},\{table\}) = 3$$

$$\overline{R}(\{bottle\ 1\},\{table\}) = 1$$

$$\overline{R}(\{apple\ 6\},\{apple\ 1\}) = \overline{R}(\{apple\ 1\},\{apple\ 6\}) = \infty$$

$$\overline{R}(\text{ apples on the bowl },\{box\}) = 2$$

We could think about using different quantales or defining differently the relation. One idea that may be worth exploring, is to relate subsets using the same quantale, by defining the relation so that it defines an idea of distance between the subsets. However this somehow defies the simplicity of the model and making it maybe too artificial.

This work could be expanded in several directions, firstly it would be interesting to be able to computationally simulate the model to produce some more formal and concrete examples. It would be in fact useful make a computer program suitable to deal with a scene with a small number of objects, able for example to translate sentences into specific subsets of objects and decide on the degree of truth and falsity of sentences in a given sentence space.
Another interesting direction would be to explore more broadly the different choices of possible Quantales and the potential to create richer sentence spaces, allowing us to decide if there are other operations on subsets that could still be allowed in the categorical framework.
Unfortunately powersets have several weak points, the most obvious is given by the exponential growth of the space itself which blows up, and this is even more worrying considering the impracticality of the tensor products. Another interesting perspective may be to take this idea about neighbouring sets one step further and suggest a connection with topology rather than the geometrical approach taken by Gardenfors. I was always of the idea that actual distances play little role and that most of our spatial intuition is in fact more topological in flavour.

# Bibliography


[1] Joe Bolt, Bob Coecke, Fabrizio Genovese, Martha Lewis, Dan Marsden, and Robin Piedeleu. Interacting conceptual spaces I : Grammatical composition of concepts. *CoRR*, abs/1703.08314, 2017.

[2] R. Brown and T. Porter. Category theory and higher dimensional algebra: potential descriptive tools in neuroscience. 2003.

[3] Stephen Clark. *Vector Space Models of Lexical Meaning*, chapter 16, pages 493–522. Wiley-Blackwell, 2015.

[4] Bob Coecke and Eric O. Paquette. *Categories for the practising physicist.* 2009.

[5] Bob Coecke, Mehrnoosh Sadrzadeh, and Stephen Clark. Mathematical foundations for a compositional distributional model of meaning. *CoRR*, abs/1003.4394, 2010.

[6] Peter Gardenfors. The geometry of preposition meanings. *Baltic International Yearbook of Cognition, Logic and Communication: Vol. 10*, 2015.

[7] S Garrod, G Ferrier, and S Campbell. In and on: investigating the functional geometry of spatial prepositions. *Cognition*, 72(2):167?189, September 1999.

[8] J. Lambek. *From Word to Sentence: A Computational Algebraic Approach to Grammar.* Open access publications. Polimetrica, 2008.

[9] Daniel Marsden and Fabrizio Genovese. Custom hypergraph categories via generalized relations. *CoRR*, abs/1703.01204, 2017.

[10] Mehrnoosh Sadrzadeh, Stephen Clark, and Bob Coecke. The frobenius anatomy of word meanings I: subject and object relative pronouns. *CoRR*, abs/1404.5278, 2014.