Highlights

Quantum projections on conceptual subspaces

Alejandro Martínez-Mingo, Guillermo Jorge-Botana, Jose Angel Martinez-Huertas, Ricardo Olmos Albacete

- A new data-driven method is developed and applied for the generation of multidimensional subspaces in Semantic-Vector Space Models.
- We exemplify the violation of symmetry by order effect in a Semantic-Vector Space Model.
- We exemplify the violation of triangular equality in a Semantic-Vector Space Model.
- Two new models are proposed and applied to evaluate the diagnosticity effect in Semantic-Vector Space Models.

Quantum projections on conceptual subspaces

Alejandro Martínez-Mingo[a,*], Guillermo Jorge-Botana[b], Jose Angel Martinez-Huertas[c], Ricardo Olmos Albacete[a]

[a]*Department of Social Psychology and Methodology, Universidad Autonoma de Madrid (UAM), Calle Ivan Pavlov, 6, 28049, Madrid, Spain*
[b]*Department of Psychobiology & Behavioral Sciences Methods, Universidad Complutense de Madrid (UCM), Campus de Somosaguas, Ctra. de Húmera, s/n, 28223, Madrid, Spain*
[c]*Department of Methodology of Behavioral Science (UNED), Calle de Juan del Rosal, 10, 28040, Madrid, Spain*

Abstract

One of the main challenges of cognitive science is to explain the representation of conceptual knowledge and the mechanisms involved in evaluating the similarities between these representations. Theories that attempt to explain this phenomenon should account for the fact that conceptual knowledge is not static. In line with this thinking, many studies suggest that the representation of a concept changes depending on context. Traditionally, concepts have been studied as vectors within a geometric space, sometimes called Semantic-Vector Space Models (S-VSMs). However, S-VSMs have certain limitations in emulating human biases or context effects when the similarity of concepts is judged. Such limitations are related to the use of a classical geometric approach that represents a concept as a point in space. Recently, some theories have proposed the use of sequential projections of subspaces based on Quantum Probability Theory (Busemeyer & Bruza, 2012; Pothos et al., 2013). They argue that this theoretical approach may facilitate accounting for human similarity biases and context effects in a more natural way. More specifically, Pothos & Busemeyer (2011) proposed the Quantum Similarity Model (QSM) to determine expectation in conceptual spaces in a non-monotonic logic frame. To the best of our knowledge, previous data-driven studies have used the QSM subspaces in a unidimensional way. In this paper, we present a data-driven method to generate these conceptual subspaces in a multidimensional manner using a traditional S-VSM. We present an illustration of the method taking Tversky's classical examples to explain the effects of Asymmetry, Triangular Inequality, and the Diagnosticity by means of sequential projections of those conceptual subspaces.

Keywords: Quantum Similarity Model, Semantic-Vector Space Models, Computational Linguistics, Similarity

1. Introduction

Representational theories of concepts have had great importance throughout the history of psychology. Historically, one of the first formalizations of that representation was done by means of semantic networks and ontologies (Quillian, 1967), in which words were defined in a relational way (e.g., a dog is a mammal and in turn a mammal is an animal). In this way, we can make inferences (a dog is an animal) applying logical rules to a set of facts (a dog is a mammal). This type of formalization is successfully used in meaning production systems, but it has some limitations in terms of expressive capacity (both facts and rules must be provided a priori) and plausibility (it becomes difficult to account for context effects in meaning construction processes) (Balkenius & Gärdenfors, 2016; Gärdenfors, 1996). Both limitations have a significant impact on psychological studies. Another early formalization, in addition to the aforementioned one, is the so-called Meaning-Text Theory (MTT) (Žolkovskij & Mel'čuk, 1967a,b; Mel'čuk, I., 1974). This theory posits

*Corresponding author
Email address: alejandro.martinez@uam.es (Alejandro Martínez-Mingo)

semantics as the foundation of the language model, thus rendering its postulates more congruent with the methods proposed in this work. The conceptual representation in the semantic module of the language model of the MTT will serve as the basic unit with which we can later engage within the syntactic, morphological, and phonological modules. In the discussion of this work, we shall assess the compatibility of the proposed data-driven method for generating semantic subspaces with the MTT model.

Another widely adopted trend in modelling is the use of Semantic-Vector Space Models (S-VSMs), which aim to represent the semantics of concepts via vectors in high-dimensional spaces (Deerwester et al., 1990; Blei, Ng, & Jordan, 2003; Mikolov et al., 2013; Pennington, Socher, & Manning, 2014; Joulin et al., 2016; Peters et al., 2018; Devlin, Chang, Lee, & Toutanova, 2018) . One of the key features of these models is their ability to automatically generate a high-dimensional space from the occurrence of words in text corpora. In this approach, concepts can be represented as coordinates in a vector space, enabling the computation of distance and similarity metrics between them. Initially, these models were primarily focused on the acquisition of word meaning as an inductive tool and the perception of similarity through distance measures (Chwilla & Kolk , 2002; Landauer & Dumais, 1997). However, the challenge of context sensibility emerged very early, and increased the interest in the construction of concepts and the generation of inferences (Bhatia, 2017; Jones et al., 2011; Kintsch, 2000; Kintsch & Bowles, 2002; Lemaire & Denhiere, 2006; Mikolov et al., 2013; Millis & Larson, 2008; Yeari & van den Broek, 2015). Moreover, some mechanisms have been used to emulate the construction of meanings, where the retrieval context guides the inhibition or activation of some components of the vectors that represent the words (Baroni & Lenci, 2009; Jorge-Botana et al., 2020; Kintsch & Mangalath, 2011). This is particularly advantageous in the examination of metaphor (Kintsch, 2000; Kintsch & Bowles, 2002; Kintsch & Mangalath, 2011; McGregor et al., 2019), where rule-based models have been demonstrated to be inadequate (Gärdenfors, 1996). Nonetheless, these mechanisms have certain weaknesses because they are highly dependent on their parametrization and have some difficulties to retrieve the right contextualized content (Harati et al., 2021). Space models generating by both Count-based models as LSA or Prediction-based models as word2vec, face similar challenges when simulating semantic tasks, such as the potential asymmetry between similarities of terms or different bias produced by humans (Jorge-Botana et al., 2020). In general, we can say that S-VSMs are able to explain a large number of cognitive processes related with meaning construction, but they need excessive parametrization to handle the retrieval contexts (see, for example, some ad-hoc adaptations like the Construction-Integration model by Kintsch (1988), or the Yeari & van den Broek (2015) online reading model). Thus, these proposals do not provide a completely natural mechanism to account for the context influence, which becomes a limitation when it comes to formalizing a general theory of semantics, since we will not be able to account for a large number of basic semantic phenomena in a naturalistic manner.

Today, neural network models such as recurrent neural networks (RNNs) and Transformers represent a general-purpose mechanism to deal with contextualized representations of concepts. As a generalization, it can be said that neural networks are a theoretical framework for modelling the formation of symbols or concepts themselves in an auto-generative way. The concepts managed within these models are emergent since they come from the interaction of the inputs with the weights of the connections of the net. In the training process, the net learns what constrictions are useful to deal with the inputs and make the concepts arise. Once a net is trained, it can activate concepts in an on-line manner by means of these constrictions. The genesis of such concepts in RNN and Transformers is accounted by a mechanism sensible to the order. They can maintain previous states to constrain further ones. Last words in a sentence are contextualized with previous ones, so a concept is generated by a linked sequence. This is the key to what is called "contextualized embeddings" in this kind of frameworks. Nonetheless, they have some issues that make them cognitively unplausible. They need a huge sample to generalize the rules for concept formation, being unable to model the fact that input received by humans is insufficient to explain their detailed conceptual knowledge they have. This fact is usually called the "poverty of stimulus". The most popular face of this phenomena is the systematic compositionality, where neural networks have still a path to improve because they seem to work by mere statistic (Lake & Baroni, 2018; Pannitto & Herbelot, 2022; McClelland, 2022). The same can be said about S-VSMs mentioned in previous paragraphs (see Jorge-Botana et al. (2020)). For this reasons, rule-based logics remain valid for addressing issues that are difficult to formalize for distributional models and neural networks, for example, the mentioned systematic compositionality. Although rule-based logic treats symbols or concepts as a-priori established entities and the genesis is out of it focus, the approach is considered valid and useful to explain specific aspects of cognition. It manages established entities with rules operating upon them. This framework is inadequate for simulate tasks related with meaning generation with subtle connotations, but adequate for simulating tasks with subtle rule utilization. A joint proposal between rule based and distributional models has been made recently by Gärdenfors (Osta-Vélez & Gärdenfors, 2022) in which conceptual subspaces proposed in previous studies (Gärdenfors, 1996) are combined with rules to formalize expectations as instances of contextualized concepts. We will mention Gärdenfors's subspaces later. In this line, a common ground exists where distributional theories can be integrated with symbolic proposals, yielding mutual benefits. One way in which this rule-based logics framework has impacted NLP models is by providing a foundation for incorporating context-sensitive representations of meaning into these models (Abbott, 1999). For instance, a recurrent neural network (RNN) or transformer-based model could be trained to consider relevant contextual assumptions when generating a semantic representation of an expression (Sundermeyer et al., 2012), thereby improving the accuracy of the model in tasks such as machine translation or information retrieval, where the interpretation of expressions is contingent upon the context in which they are used (Manning et al., 2008).

A more natural way to represent conceptual knowledge comes from a recent proposal based on Quantum Cognition (Batchelder et al., 2018; Busemeyer & Bruza, 2012). The term Quantum in these proposals is not just an analogy. Conversely, it uses some parts of the formalization of the Quantum Probability Theory from a geometric approach to introduce a plausible mechanism by which concepts are represented as subspaces. The fact that a concept is represented in a subspace implies that concepts have potential meanings. A meaning only emerges (or is substantiated) at the retrieval moment and depending upon a given specific context, usually past events. In this vein, Pothos & Busemeyer (2011) developed the Quantum Similarity Model (QSM) based on hypothetical conceptual subspaces and sequential projections of a state vector in such subspaces. Each concept is represented by a subspace, with the most important features of the concept as its basis. These conceptual subspaces are located in a higher order Hilbert space [1] of high dimensionality. All these representations and operations are based on the logic of Quantum Probability Theory. Thus, although the representation continues to be geometric, the entities (concepts) are no longer represented as points or coordinates but as complete subspaces (we deal with this later). Consequently, thanks to the nature of quantum dynamics, it is possible, for example, to naturally capture the order effect in similarity estimations between concepts (e.g., the similarity measure of *China* and *Korea* do not have the same estimation as the similarity of *Korea* and *China*). It is important for the understanding of this work to review the fundamentals of these models, previously formalized by Pothos & Busemeyer (2011); Busemeyer et al. (2011); Busemeyer & Bruza (2012); Pothos et al. (2013); Yearsley et al. (2015). There are four fundamental elements in the quantum similarity model: concepts, mind state, updating by sequential projections and compatibility.

1.1. Concepts

From a quantum point of view, concepts are represented as subspaces within a semantic space $\mathcal{H}_A \subset \mathcal{H}$ (a Hilbert space $\mathcal{H}$) with finite dimensionality N, spanned by an orthonormal set of basis vectors $\mathcal{V} = \{|V_1\rangle, ..., |V_N\rangle\}$. From now on, we will call this space $\mathcal{H}$ as the *container space*, represented by dimensions e1, e2 and e3 in Figure 1. All subspaces are referenced in that container space, so we can imagine that the concepts *China* and *North Korea* are represented in the container basis by the one- and two-dimensional subspaces also depicted in Figure 1. Concretely, *North Korea* subspace $\mathcal{H}_{Korea} \subset \mathcal{H}$ is spanned by $\mathcal{V}_{Korea} \subset$

[1] Hilbert spaces represent a powerful mathematical tool in the study of infinite-dimensional spaces, particularly in the analysis of quantum systems (Von Neumann, 2018). These spaces are complete, normed vector spaces equipped with a notion of length, or norm, and a means of measuring the distance between two points in the space. The defining characteristic of a Hilbert space is the presence of an inner product, which enables the calculation of orthogonality and projections in the space (Von Neumann, 2018). The application of Hilbert space operations in finite Vector Space Models (VSMs) provides a versatile and sophisticated framework for the analysis and manipulation of complex systems. This approach has proven to be particularly useful in the study of quantum systems and has led to significant advancements in our understanding of these systems (Von Neumann, 2018).

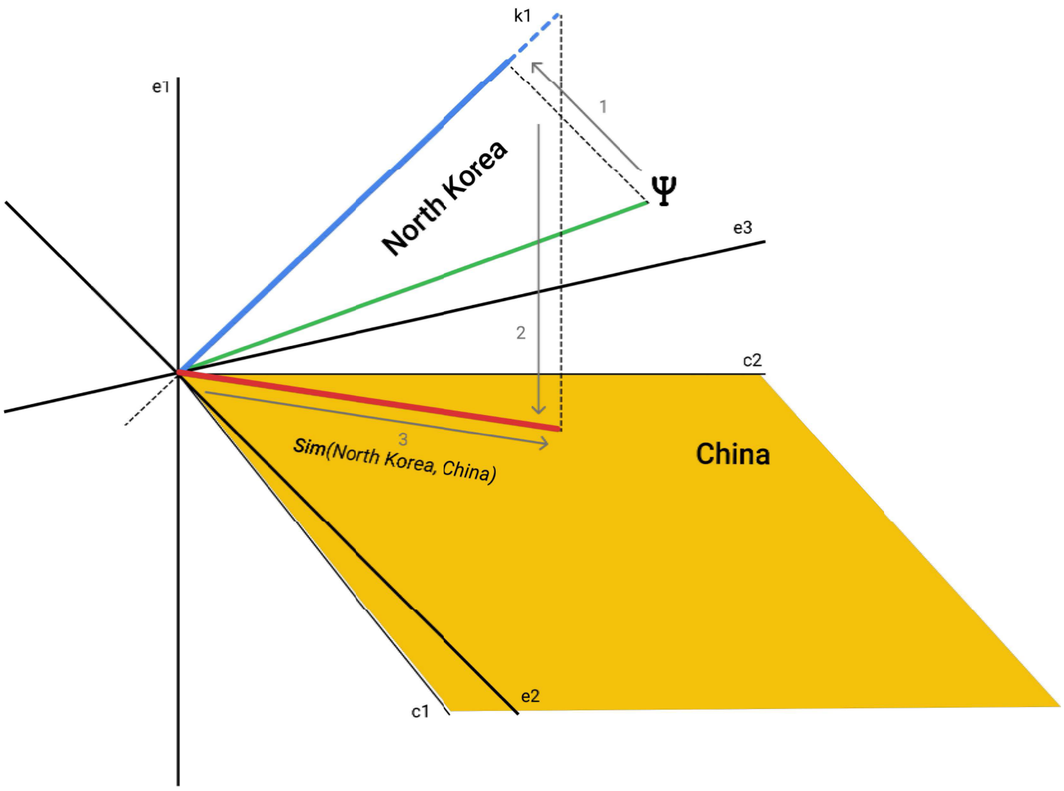

Figure 1: Meaningful subspaces representation (based on Duran et al. (2016)). Assuming that the container space U is defined according to a canonical basis $B = \{e_1, e_2, ..., e_{300}\}$, and that C is a two-dimensional subspace with basis $C = \{c_1, c_2\}$, and K is a one-dimensional subspace with basis $K = \{k_1\}$, the relationship between the *China* and *North Korea* vectors is identified in U, since the set $\{c_1, c_2, k_1\}$ share a common reference, which is the basis of the original latent space U. Note that e_1, e_2 and e_3 represent $B = \{e_1, e_2, ..., e_{300}\}$ in this figure.

$\mathcal{V}$ with only one vector in its basis (k1 in Figure 1), and *China* subspace $\mathcal{H}_{China} \subset \mathcal{H}$ is spanned by $\mathcal{V}_{China} \subset \mathcal{V}$, that has two vectors in its basis (c1 and c2 in Figure 1). This way of representing concepts is a qualitative change with respect to the classical perspective, in which concepts are only points defined by vectors in a vector space, that is, a concept is a set of scores in the components of a vector of the container space.

1.2. Mind State

The mind state is represented by a n-dimensional vector in the container space ($\mathcal{H}$, the Hilbert space), called state vector ψ. Initially, this vector represents a neutral thought with a certain potential of collapse on each of the subspaces representing the concepts, i.e., the state vector can be instantiated into something concrete due to thinking about a concept represented with a subspace. Thus, we will speak about amplitudes between the state vector and each of the subspaces. Technically speaking, collapsing is when the mind state is projected onto a subspace, the probability of collapsing being equal to $q(A) = ||\mathbf{P}_A|S\rangle||^2$, where $\mathbf{P}_A$ is a matrix operator that projects the mind state in the subspace of the concept A. Since a vector-matrix multiplication returns a vector, the squared norm (the squared amplitude) of that vector is a scalar that indicates the probability of collapsing. In Figure 1, we can observe that there is a state vector ψ that has a relative amplitude with respect to *China* and *North Korea* subspaces. To identify the projectors for each subspace, the outer product operation is used. A projector of a subspace is calculated with the sum of the outer products of each basis vector with itself as follows:

$$\mathbf{P}_A = |v_1\rangle\langle v_1| + |v_2\rangle\langle v_2| + ... + |v_n\rangle\langle v_n|, \tag{1}$$

being the subspace of A spanned by a basis with n vectors $\{v_1, v_2, \ldots, v_n\}$. For example, we can obtain the projector of *China* summing up the outer products of the vectors that spanned the *China* subspace. If the subspace of *China* has a basis of two vectors, the vectors $China_1$ and $China_2$, the outer product of the first is noted as $|China_1\rangle\langle China_1|$ and the outer product of the second is $|China_2\rangle\langle China_2|$. The projector is the sum of both outer products: $\mathbf{P}_{China} = |China_1\rangle\langle China_1| + |China_2\rangle\langle China_2|$.

Once we have the projector, we can project the vector state ψ onto the *China's* subspace as $\mathbf{P}_{China}|S\rangle$, and calculate the probability of collapsing as $||\mathbf{P}_{China}|S\rangle||^2$. When the state vector is projected onto *China's* subspace, the state vector is no longer a neutral state, but it is instantiated in *China's* subspace. This vector may be projected in another concept subspace; that is, thinking about other concepts after *China*. This process is explained in the next section.

1.3. Updating by sequential projections

When a state of mind is concrete — that is, the state vector takes the form of a specific vector in a concept subspace — it is susceptible to being updated. State updating happens when the already collapsed state vector is again projected onto another concept subspace. It is, in essence, a sequential projection of the state vector in several concept subspaces, so that the state is taking several forms according to previous projections. A state vector changes its meaning as it is projected from different points of view; that is, is projected in different concepts. For example, we will be able to obtain a sequence of projections thinking first about *Korea* and then about *China* (Equation 2). The resultant state vector is one specific vector in *China's* subspace determined by the previous projection in *Korea's* subspace (Arrows 1 and 2 in Figure 1).

$$\mathbf{P}_{China}\mathbf{P}_{Korea}|S\rangle \tag{2}$$

1.4. Compatibility

The concept of compatibility is key in quantum probability. If we are able to represent two concepts using the same basis, they will be compatible. However, if two concepts cannot be represented using the same basis, they will be considered incompatible, and we will only be able to understand the collapsing probability between the two concepts sequentially. We could consider events as compatible if the one-dimensional subspace of Korea were one of the dimensions of the two-dimensional subspace of China. This way, both concepts could be understood within a framework of compatible dimensions, by using just one dimension to describe Korea and two dimensions to describe China. In our case, we will define concept subspaces as of incompatible bases, each concept being represented by a basis of unitary orthonormal vectors that will be interpreted as the main potential *characteristics* of that concept. In Figure 1, we can see that the state vector can be explained by two different bases (*North Korea* or *China*), each with different dimensionality, so we can represent the state of ambiguity between the two concepts. In addition, each of the vectors of one basis can be explained by the linear combination of all the vectors of the other basis and also by the vectors of the container space. It is important to note that the reciprocity law must be fulfilled between each of the vectors representing the basis of the first concept and each of the vectors representing the basis of the second concept (Busemeyer & Bruza, 2012).

The Quantum Similarity Model (Pothos & Busemeyer, 2011) is based on this general framework. According to this model, we can run successive projections, starting from a neutral state vector, to calculate the probability of the collapse of the system on different concepts sequentially. Taking the example of Figure 1, we can calculate the similarity between *North Korea* and *China* as follows:

$$\begin{aligned} &||\mathbf{P}_{China}\mathbf{P}_{Korea}|\psi\rangle||^2 \\ &= ||\mathbf{P}_{Korea}|\psi\rangle||^2 \cdot ||\mathbf{P}_{China}|\psi_{Korea}\rangle||^2 \\ &= ||\mathbf{P}_{Korea}|\psi\rangle||^2 \cdot ||\mathbf{P}_{China}\frac{\mathbf{P}_{Korea}|\psi\rangle}{||\mathbf{P}_{Korea}|\psi\rangle||}||^2 \end{aligned} \tag{3}$$

In the last expression, $||\mathbf{P}_{China}\mathbf{P}_{Korea}|\psi\rangle||^2$represents the similarity between *North Korea* and *China* (Pothos & Busemeyer, 2011). The dynamics of the QSM is detailed in the studies of Duran et al. (2016); Pothos & Busemeyer (2011); Pothos et al. (2013, 2015); Yearsley et al. (2014, 2015, 2017). In these papers, the authors aim to show the potential of the QSM to successfully model several cognitive biases.

For the QSM, the mind state is only a potential state, and the meaning of a concept is only determined in the presence of a context, being the concept constructed and not statically recovered. According to this, some previous studies have already alluded to the term *aberrant* to define vectors representing non-contextualized terms in a semantic-vector space (Jorge-Botana et al., 2010, 2011). In quantum terms, *degenerated* also alludes to the same purpose. Thus, a polysemous term not represented in any concrete concept subspace will possess *degenerate* meanings (Aerts et al., 2011), and only upon contacting the context of a subspace will it collapse into a concrete meaning. Other studies have alluded to this state as *ground-state*, cognitively defined as the state in which a concept is not evoked; that is, when it does not come into the focus of conscious experience and is not influenced by any context (Gabora et al., 2020). The conclusion is that a concept has infinite possible states depending on the context, and is never experienced in the ground-state. In the same way, in QSM there are no neutral situations (Pothos & Busemeyer, 2011); whenever a person is asked about the similarity of two concepts, they will not be able to judge it in absolute terms, but from a pre-established point of view. For example, if asked about the similarity between *ball* and *moon* with a preeminence of the point of view of the *form*, it is possible that they will be judged as equal but, if there is a preeminence of the point of view of the *game*, they will be judged as different, with neutrality of the point of view not being preeminent (Aerts et al., 2011). All these analogies regarding the construction of meanings already have formal apparatus in Quantum models in physics, so that Gabora & Aerts (2002) suggest a formalism for their treatment, proposing the so-called *contextualization theory* of language. Nonetheless, no data-driven study has been done (that is, there is no empirical evidence using corpus-based conceptual subspaces). As such, it is not clear what the actual vectors are that would generate the basis of the *China* subspace, nor how the degenerate vectors will be represented in a space that determines a context such as the *China* subspace. Both questions need a data-driven approach (corpus-based) to their formalization. Till now, only a tentative approximation has been done with the HAL model, but without a full fit to the quantum similarity model (Busemeyer & Bruza, 2012).The primary objective of this study is to present a method for the identification of both the basis and dimensionality of concepts subspaces with the aim of formalizing the Quantum Similarity Model (QSM) on multidimensional subspaces derived from VSMs. In this case, a textual based model as LSA as a S-VSM.

2. QSM on the rails of a S-VSM

Today, several authors have proposed the use of the mathematical framework provided by Quantum Probability Theory to capture the meaning of words within S-VSMs (Aerts et al., 2013; Blacoe et al., 2013; Gonzalez & Caicedo, 2011; Jaiswal et al., 2018). The challenge now is to define a method that allows to automatically identify concept subspaces in the common reference of a container space. This way, we need to determine which is the container space, and how we can identify the position of the vectors that span each concept subspace in it. The use of a data-driven approach is well suited, as the subspaces will be formed based on the data rather than the researcher's input.

Studies by Aerts & Czachor (2004) and Bruza & Cole (2006) were cornerstones, since they suggested that some S-VSMs such as LSA (Deerwester et al., 1990) or the HAL (Hyperspace Analog to Language model) (Lund & Burgess, 1996) could be considered formal Hilbert's spaces, and thus act as a container space (Blacoe et al., 2013). The use of S-VSMs will be the basis of our study to identify concept subspaces in a data-driven environment. Such a data-driven environment has the advantage that the vectors of the basis of each concept subspace are not generated rationally, as in previous studies (with the exception of the proposal of Busemeyer & Bruza (2012)). Conversely, the locations of those vectors and their distances are extracted by the inferences of a model sensitive to massive cooccurrence.

To this end, we chose to use LSA (Deerwester et al., 1990), as this model has been widely studied in the field of concept representation and has been tested in many cognitive tasks. Another reason is that it provides an orthonormal and well-suited space for QSM. LSA is a method that models the acquisition of meaning from texts by acquiring meanings from concepts (words) presented in different contexts (documents). From the LSA point of view, what happens in our mind when processing language is not different from what happens in any other cognitive process. Thus, we must have a system that facilitates the mapping of information acquired by experience into a format in which our mind can process and store it as language (Landauer & Dumais, 1997). LSA makes use of a linguistic corpus, which ideally contains all the concepts that a human being can acquire, to represent semantic properties given relations between words. Thus, the final product is a system that represents words as vectors whose dimensions map the information acquired by experience with texts. However, the information contained in an entire linguistic corpus is too granular to represent human semantics. Just as a human being does not consider all the stimuli received to generate a semantic representation of concepts, in LSA we must reduce the dimensionality of a linguistic corpus to summarize its information and make it operational. Thanks to this dimensionality reduction, we can represent the information contained in a corpus in a simplified way, which makes it possible to establish relationships between concepts that a priori are not found in the same contexts but are deeply related. Specifically, LSA uses a technique called Singular Value Decomposition (SVD) that allows us to extract N orthonormal dimensions that represent the information contained in a corpus, in a reduced form. It is important that these dimensions are orthonormal and non-interpretable since they act as a symbolic coordinate axis of the semantic properties which allows us to represent each of the concepts contained in the corpus by means of n-dimensional vectors. Hypothetically, the latent dimensions represent a closed semantic world that we can define as a container with all the potential knowledge; that is, the Hilbert's space in QSM terms.

We are aware that LSA has many limitations, and we know that there are other models that currently outperform LSA in various tasks (Jorge-Botana et al., 2020). This is why we do not limit the use of the proposed method to a specific computational model such as LSA, but rather propose a series of algorithms that can be used with any container space in which all concept subspaces can be represented. However, it is important to note that this container space must be composed of a basis of orthogonal vectors in order for it to be considered a Hilbert space (Von Neumann, 2018), and thus, be able to represent the vectors of the different subspaces within it. LSA can ensure this condition as it uses truncated SVD to perform the dimensionality reduction of the term-document matrix. However, other methods such as word2vec, Glove , or even more modern methods based on RNN and/or Transformers such as BERT or GPT, do not prioritize the orthogonality of the space, sacrificing it to find a more accurate representation of each term. From our perspective, we believe that a model that captures the different nuances of a concept from a subspace with base vectors that represent these different "directions" in which we can understand a term, will have a greater explanatory power for the same term in different contexts. This has led us to, for now, only consider the application of our method using an LSA-based container space, although this does not mean that we reject the use of other computational language models, nor the application of subspace composition methods without the consideration of a higher-order container space as proposed in Mu et al. (2017); Li et al. (2018); Yang et al. (2018); Shimomoto et al. (2021), assuming in this case that the subspaces of the different terms do not share a common space, and making it impossible for this reason to apply the QSM as conceived in this work for calculating similarities or distances between concepts.

To generate that space with LSA, we will use journalistic texts containing articles between the years 2018 and 2019 from two Spanish newspapers (*El Mundo* and *El País*) as the corpus. This corpus is composed by 243.320 documents and it has a total vocabulary of 33.503 words. We get a 300-D latent semantic space, after applying the log-enthropy to the term-document matrix, making use of the Gallito Studio tool (Jorge-Botana et al., 2013). With this container space and based on a previous approximation by Martinez-Mingo et al. (2020), we propose a method for the identification of concept subspaces.

3. Multidimensional Subspaces Estimation Method

Busemeyer & Bruza (2012) discussed a method for delimiting subspaces in S-VSMs. In their work, they suggested an ad-hoc method in which a concept subspace can be determined by the weighted sum of different inputs in a HAL semantic vector space, generating groups of inputs for each characteristic of that concept subspace (Busemeyer & Bruza, 2012). Nonetheless, this proposal has limitations concerning the orthonormality of subspaces. The preliminary proposal of Martinez-Mingo et al. (2020) pursued an approximation generating conceptual subspaces via clustering K-means technique. Taking a target word, its closed semantic neighbors (for example its 100 semantic neighbors) were extracted from the LSA semantic space. Then, the K-means clustering technique was applied to group the semantics of that neighborhood into K clusters. Subsequently, the centroid of each cluster was calculated. The subspace was obtained by applying an orthogonalization of the centroids. Nevertheless, K-means solution is relatively unstable, and the bases of the subspaces fluctuate.

The present study aims to settle a robust and more operational implementation of multidimensional subspaces in S-VSMs. As has been said, the current proposal starts from an already estimated LSA semantic space, which is used to estimate the initial coordinate solution of the subspaces of each of the concepts. The proposed method to generate the concept subspaces can be divided into three steps: container space generation, contour extraction, and estimation of the deserved dimensionality of the subspace. This method aims to overcome previous limitations as to the non-orthonormality of subspaces and the instability of preliminary proposals. As we will see, this proposal seems to illustrate some classic cognitive biases of the QSM.

3.1. Container space generation

We begin by generating a semantic vector space, the container space $\mathcal{H} \in R^{300}$, of 300 dimensions (Rehder et al., 1998) using LSA. This container space will be spanned by a set of orthonormal vectors $\mathcal{V} = \{|V_1\rangle, ..., |V_N\rangle\}$ that contains all the knowledge on which concepts can be represented. In our illustration, we are using a journalistic corpus based on samples from two Spanish newspapers. It is important to use a vector space model that guarantees the orthonormality of the dimensions that defines it, and LSA's SVD guarantees that the dimensions of the resulting matrices are orthonormal (Landauer et al., 2013).

3.2. Contour extraction

We define a contour for a specific concept as a set of n-dimensional vectors representing the words most related to the word that labels the target concept (for example, *China*). Therefore, the contour is a set of M vectors $A = \{|A_1\rangle, ..., |A_M\rangle\}$ that are actual points in the container space $\mathcal{H}$. This contour is used to find the basis of the subspace $\mathcal{H}_{\mathcal{A}}$ of concept A. In order to define those points, we are going to extract the nearest neighbors of a specific concept (e.g., the first 10 terms and the last term of the *China* contour can be seen in Table 1), understanding as nearest neighbors those concepts that have a higher similarity with our target concept. Similarity phenomena in LSA and other S-VSMs usually follow a negative logarithmic function, where only a few neighbors have high similarities with each concept (Jorge-Botana & Olmos, 2014). The challenge is to identify the similarity threshold in which that contour ends. In other words, which semantic neighbor should be the last of that contour?

Table 1: China Contour

	Dim. 1	Dim. 2	Dim. 3	Dim. 4	...	Dim. 300
Chinese	2.53	1.20	-2.02	0.65	...	-0.29
Beijing	0.48	0.37	-0.27	0.11	...	-0.03
Jinping	0.13	0.14	-0.01	0.02	...	-0.01
Asian	0.61	0.13	-0.47	0.20	...	-0.05
Taiwan	0.11	0.07	-0.06	-0.01	...	-0.01
Shanghai	0.13	0.02	-0.09	0.08	...	-0.02
Xi	0.26	0.18	-0.04	0.02	...	0.01
Xi Jinping	0.03	0.03	-0.01	0.01	...	0.00
People's Republic	0.02	0.01	-0.01	0.01	...	0.01
Hong Kong	0.28	0.17	-0.15	0.00	...	0.00
...	...	...	...	...	...	...
Geopolitical	0.18	0.14	-0.06	-0.02	...	0.02

In estimating a threshold, we have been inspired by the work of Rekabsaz et al. (2017). Following Rekabsaz et al. (2017) reasoning, we wonder how the semantic neighborhood of a very uncertain word and the form of the function that matches the similarities of its neighbors could be. Answering this question provides a suitable reference to compare the behavior of this uncertain word with a target word and find the similarity threshold of the contour of the latter. The point at which similarities of the neighbors of the uncertain word coincide with similarities of the neighbors of the target word is the threshold of the target's contour. The sense of this method is to find the point at which the similarities of the neighbors of the target word behave as those of the neighbors of the uncertain word. In addition to this, we added an additional correction based on the semantic diversity of the target word. This correction consists in weighting the threshold (similarities of the neighbors) of the uncertain word by a factor based on the semantic diversity of the target. This weighting allows to obtain bigger contours as diversity of target words increases. Formally, the procedure is as follows:

1. The sum of all vector terms in the container space is calculated as in Equation 4. The result of this sum is a vector that represents the uncertain word. We name that word as *superterm*.

$$Superterm = \sum_{i=1}^{K} A_i, \tag{4}$$

where A_i is each word vector in the container space and K is the number of dimensions of the same space.
2. A list of similarities of each word of the space with this *superterm* is calculated using the cosine. This list is defined by each word and its similarity with the *superterm*. We sort that list to draw a distribution of similarities.
3. The Semantic Diversity of the target word is calculated. Semantic Diversity is the degree in which a word appears in different contexts. This index is calculated with the method of Hoffman et al. (2013) as follows:

$$Diversity(w_t) = \left(\sum_{j:w_t \in P_j} \sum_{k:w_t \in P_k} \frac{cos(\vec{P_j}, \vec{P_k})}{m \times m} \right), \tag{5}$$

where w_t is the target word, m is the number of paragraphs in which w_t appears, P_j and P_k are the paragraphs with index j and k respectively, and $\vec{P_j}$ and $\vec{P_k}$ their vector representation with index j

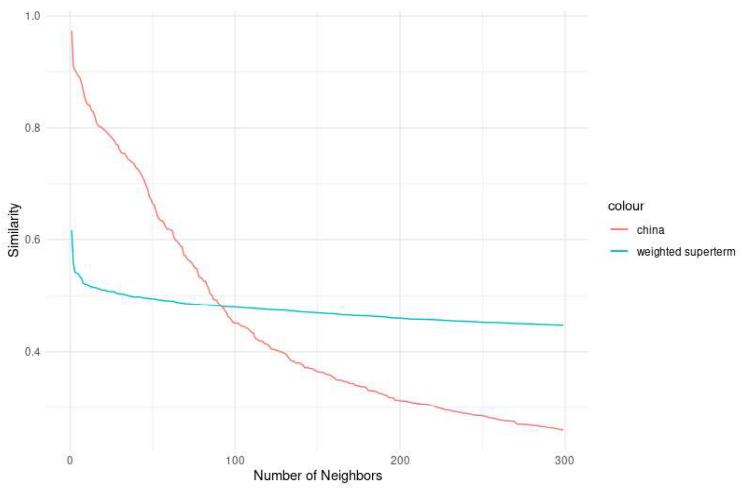

Figure 2: The figure shows the sorted similarity function between the superterm vector and all words in the corpus corrected by the semantic diversity of the target word, and the sorted similarity function between the vector of the target term and all words in the corpus.

and k in the container space. Expressed in natural language, this formula takes all the paragraphs in which the target word appears and calculates a similarity measure (*cosine*) between them. To have a central measure, a mean of all the similarities is calculated and a $-log$ operation is applied to it. Thus, diversity is sensible to the disparity of the paragraphs that the target word occurs in.

4. A list of similarities of each word of the space with the target word is compiled. Again, this list is defined by each word and its similarity with the target word. We also sort that list to draw a distribution of similarities of the target word. Then, a threshold is estimated for this sorted distribution of similarities.
5. Each of the values of the sorted list with the similarities of the *superterm* are modified following Equation 6.

$$S_{weighted} = \frac{S}{1 + D_{target}/10}, \tag{6}$$

where S is the vector of similarity values of the *superterm* (the list of sorted similarities), D_{target} is the decile of the Semantic Diversity corresponding to the target word, and $S_{weighted}$ is the similarity vector of the superterm modified by the Semantic Diversity.

6. Once the previous steps are done, we have two distributions: one with the sorted similarities of the target word with all the words in the space, and another with sorted similarities of the *superterm* with all the words corrected by the semantic diversity of the target. We can now calculate the crossed point of both distributions to define the similarity threshold for the target word contour. All words with a similarity above the threshold to the target word will be included in its contour. Taking the term *China* as an example, as shown in Figure 2, we determine a contour of 91 semantic neighbors.

From a psychological perspective, we can consider this contour as a set of elements that can shape the basic characteristics of the target concept. This interpretation is consistent with many other cognitive models (Busemeyer & Bruza, 2012). Thus, we will use these elements when forming the basis that represents the subspace of a concept; for instance, the concept of *China*.

3.3. Estimation of subspace dimensionality

As stated above, $\mathcal{V}_A = \{|V_{A_1}\rangle, ..., |V_{A_K}\rangle\}$ is the basis of concept A subspace ($\mathcal{H}_A$). This basis contains the main characteristics of the concept in a vectorial form. It is also the reference to judge the world with the point of view of the concept A. Following the formal definition, this basis has K vectors (thus K

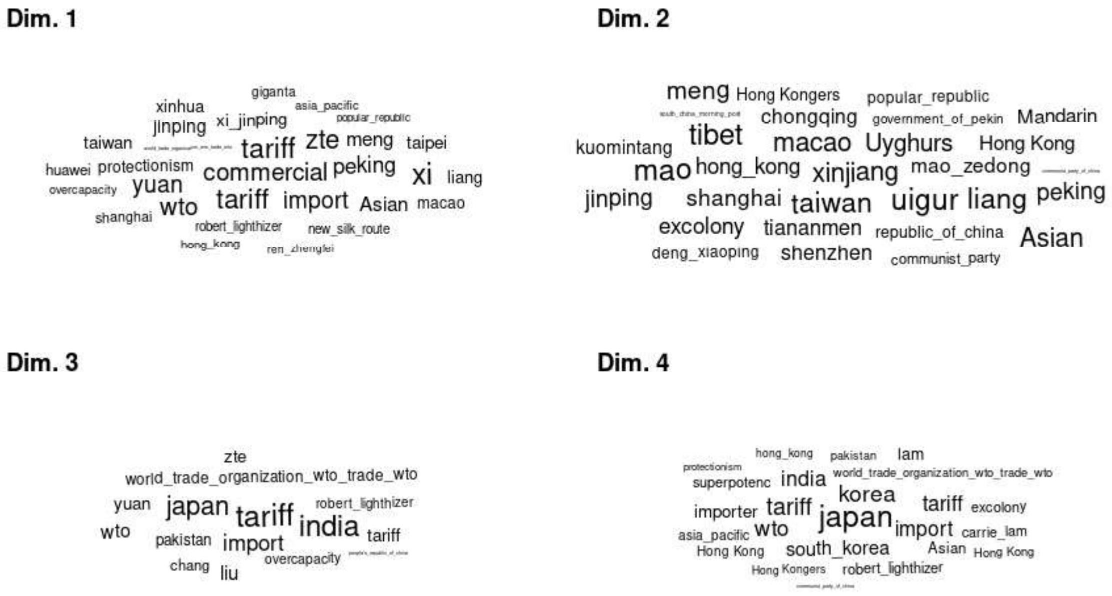

Figure 3: Cloud graphs of the most representative terms of the four dimensions that define the China subspace.

components), K being determined by the number of main directions we can extract from the contour of that concept. Thus, we need to specify K. To identify it, we determine the optimal number of orthogonal directions or main components that we can extract from the contour of a target word. The Parallel Analysis (Horn, 1965) is used to estimate the optimal dimensionality in Principal Component Analysis (PCA). Having a matrix filled with the vectors of the words of the contour (extracted from the container space), we can apply the Parallel Analysis to identify how many principal directions must be preserved in a further PCA decomposition. Then, K would be the number of principal components that the Parallel Analysis indicates to preserve, and the eigenvectors of such decomposition are the vectors of the basis of the concept subspace.

In our *China* example, the Parallel Analysis extracted 4 components. These four directions, which are the eigenvectors extracted in a PCA decomposition of the contour matrix, would be the basis of the *China* subspace, with each of them having a specific latent meaning and being orthogonal to the others.

Then, we can say that the *China* subspace deserves 4 components, having a subspace composed of a basis whose vectors, the 4 eigenvectors preserved from the PCA decomposition, are identified in the container space, so that they share the same reference with the rest of the basis for other concepts. Subspaces identified by this method can now be used to implement the projections of the state vector into them, as proposed by Pothos & Busemeyer (2011). At this point, we can interpret that our example subspace represents the fundamental characteristics of the target concept according to the LSA model. In Figure 3, we can see a visual representation of the main directions of the *China* subspace. Please note that the results obtained in this study are not experimentally validated, so they are merely examples of application of the method presented.

4. Exemplifying Cognitive Biases with Multidimensional Subspaces

In this section, we exemplify that the effects reported by Tversky (1977) can be managed through concept subspaces identified by this method, now applying the QSM (Pothos & Busemeyer, 2011). We want to make clear that we do not pretend to make substantive inferences but to illustrate how our method works using some well-known examples of cognitive biases. Below, we present different cognitive biases, reported by Tversky (1977), which we hypothesize we will be able to manage with our method.

4.1. Asymmetry Effect

Asymmetry is a well-known effect when people must judge the similarity between two concepts. If people are asked for the similarity of *North Korea* and *China*, they estimate it to be bigger than the one between *China* and *North Korea*. To make this asymmetry patent, Tversky (1977) conducted an experiment in which he asked participants: *"Which of the following phrases do you prefer to use?"*, and asked them to select *"Country A is similar to country B"* OR *"Country B is similar to country A"*. In his best-known example (*China* vs. *North Korea*), Tversky observed that 66 of the 69 participants judged the similarity between Korea and *China*, from now on $sim(Korea, China)$, to be greater than $sim(China, Korea)$. Many other examples were used in Tversky's study to demonstrate this effect.

As previously stated, this is a major problem in modeling cognitive biases if S-VSMs' ordinary distances in a space are used because the distance between two concepts will always be the same regardless of the order. Subjects will determine that there is a greater similarity when the less salient concept, or *variant*, is presented first and the more salient object, or *prototype*, second, than when they are presented in the reverse order. Kintsch (2014) noted that fact and adapted the similarity measure in S-VSMs to be sensible to the claim of Tversky (1977); that is, the similarity estimation is conditioned to the amount of information that a subject has about each of the concepts studied, always using n-dimensional LSA vectors to represent concepts, and not conceptual subspaces represented by a set of n-dimensional LSA vectors that compose a new orthogonal basis for the concept. In this section, we show operatively the Asymmetry effects of the classic examples of Tversky (1977) about countries, but using the standard apparatus of QSM with concept subspaces and projections. Importantly, our proposal is data-driven since it uses real sources of texts. We also hypothesize that the Asymmetry effect, given that there is no context effect in this case, would be mainly determined by the deserved dimensionality of the stimuli studied, with the concept with higher dimensionality being more salient.

This way, to make the asymmetry effect operative it is necessary to estimate the different subspaces for all the stimuli. As we have already extracted the *China* subspace, the next step will be to obtain the *North Korea* subspace. Eight dimensions are estimated for the *North Korea* subspace using our method [2]. Thanks to this method we are able to extract new information from the subspaces. For example, we can observe the correlation matrix of the vectors of the bases of both subspaces (Figure 4) and determine the characteristics they share. We can infer that the subspaces of both concepts are incompatible (they clearly do not share the same basis vectors), but there are some main shared features (please note that this correlation matrix is just a first exploration of both subspaces).

Next, we need to estimate the similarities between the two concepts in different orders to account for the Asymmetry effect. To do this we must first calculate the projectors $\mathbf{P}_{China}$ and $\mathbf{P}_{Korea}$ of the two subspaces, using Equation 1. Once we have the two subspaces' projectors, we need to set values to the initial state vector ψ to start with. Our choice is to initialize ψ with no bias toward either subspace, optimizing the values of the ψ initial state by minimizing $||\mathbf{P}_{Korea}|\psi\rangle||^2 - ||\mathbf{P}_{China}|\psi\rangle||^2 = 0$. We use the *stats* library (R Core Team, 2022) for this purpose. In other words, the state vector will be in an intermediate position between the two subspaces (Pothos & Busemeyer, 2011).

Having the state vector ψ, we can now formalize the asymmetry taking similarity values but changing the order. For this purpose, the squared magnitude of the sequential projections in both orders is calculated, obtaining the following results:

$$Sim(China, Korea) = ||\mathbf{P}_{Korea}\mathbf{P}_{China}|\psi\rangle||^2 = 0.277$$
$$Sim(Korea, China) = ||\mathbf{P}_{China}\mathbf{P}_{Korea}|\psi\rangle||^2 = 0.236$$

[2]We must keep in mind that the corpus used is a journalistic corpus from 2018 to 2019, so we cannot assume that this fits human semantics, even less than was the case with the results of Tversky's experiment, since this was carried out in a completely different geopolitical context.

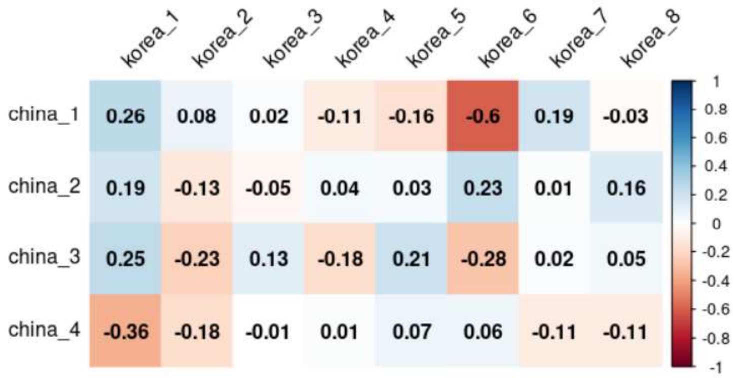

Figure 4: Correlation matrix of the extracted dimensions for *China* and *North Korea*.

As a result, there is an asymmetry between the similarities depending on the order effect. Theoretically, this could be explained by the salience of both concepts given the generation of subspaces (Duran et al., 2016). It seems that *North Korea* deserves a higher dimensionality than *China* according to this journalist corpus, which would cause it to have a higher weight when calculating similarity, with *North Korea* being the prototype and *China* the variant. This effect due to the salience of the concepts is congruent with Tversky (1977) interpretation (although the direction of asymmetry was opposite to the judgment of his experimental subjects, maybe due to the historical context).

4.2. Triangular Inequality

Related to the Asymmetry effect and the violation of the space distances, Tversky (1977) also reported the Triangular Inequality. Triangular Inequality determines that the distance between two points A and B will always be smaller than the sum of the distances between points A and C and points C and B. That is, if we follow space laws, the side of a triangle must be mandatorily smaller than the sum of the other two sides.

$$Distance(A, B) < Distance(A, C) + Distance(C, B)$$

In terms of similarity, Triangular Inequality is explained as follows: $Dissimilarity(A, B)$ will always be smaller than the sum of $Dissimilarity(A, C) + Dissimilarity(C, B)$. This would be the same as saying that $Sim(A, B)$ will always be greater than the sum of $Sim(A, C) + Sim(C, B)$ (Duran et al., 2016).

$$Similarity(A, B) > Similarity(A, C) + Similarity(C, B)$$

Nonetheless, people do not seem to follow classic space laws in their judgments. In his study, Tversky considered $A = Russia$, $B = Jamaica$, and $C = Cuba$. Space laws would predict that $Sim(Russia, Jamaica)$ is always smaller than the sum of $Sim(Russia, Cuba)$ and $Sim(Cuba, Jamaica)$. People contradict space laws in terms of distances and similarities and thus violate Triangular Inequality. However, in S-VSMs we cannot consider that the similarity works as a linear transformation of the distances. For example, working with n-dimensional vectors in LSA to represent words, the violation of Triangular Inequality can be accounted for with cosines, even when it is impossible with Euclidean distances. If we use cosines in our container space, $Sim(Russia, Jamaica) = 0.0252$, and $Sim(Russia, Cuba) + Sim(Cuba, Jamaica) = 0.4829$, which means a violation of Triangular Inequality. Nevertheless, by working with Euclidean distances instead of similarity measures, this violation of the geometric model does not occur, resulting in $Dist(Russia, Jamaica) = 10.5879$, and $Dist(Russia, Cuba) + Dist(Cuba, Jamaica) = 12.0881$. Tversky

(1977) formalized his model by applying theoretical functions to experimental data, but not by using approaches based on Semantic Models. For this reason, there are still subtle issues in Semantic Space Models outside of Tversky's approach. Other authors such as Kintsch (2014) did develop their work on this type of Semantic Models, but there are still some inconsistencies between distance and similarity measurements in demonstrating the violation of Triangular Inequality. We believe that these inconsistencies are due to the use of one-dimensional subspaces to represent the concepts, the results being equivalent to those of the classical model, and therefore we propose the use of our method to estimate multidimensional subspaces to test the Triangular Inequality using quantum distance and similarity measures.

In this way, our objective is to formalize how to deal with Triangular Inequality violation using the formal apparatus of QSM in a data-driven manner, using multidimensional subspaces instead of vectors to represent concepts in order to find violations of this principle with distance and similarity measures. It is important at this point to consider the characteristics of the different quantum distance and similarity measures proposed for semantic spaces and concept representation in the current literature. On the one hand, Gabora & Aerts (2002) propose a distance measure, called *conceptual distance*, as a function of the probability of change from a state p to a state q under the influence of a context e as follows:

$$d_\mu(q,e,p) = \sqrt{2(1-\sqrt{\mu(q,e,p)})}, \tag{7}$$

where $\mu(q,e,p)$ is the probability of this change. In our case, we will understand that the probability of this change is given by the squared length of the projection of the state vector onto the subspace representing the first concept multiplied by the projection of the updated vector on the subspace of the second concept. The formalization is as follows:

$$\mu(q,e,p) = \psi \rightarrow A \rightarrow B = ||\langle A|\psi\rangle \cdot \langle B|\psi_A\rangle||^2, \tag{8}$$

With A representing the event p, ψ the state e, and B the event q. Working with projectors, this will result in $\mu(q,e,p) = ||\mathbf{P}_B\mathbf{P}_A|\psi\rangle||^2$. Although this distance measure perfectly suits the representation of concepts by subspaces that we propose in this paper, it has a dynamic nature, it being possible to represent the order effect by measuring the distance between subspaces, but not to take a simultaneous measurement of the distance between the two subspaces. $\mu(q,e,p)$ implicitly has a sequence of a transition. In other words, it does not allow us to determine a concrete distance between two subspaces independently of the order to measure the Triangular Inequality effect.

On the other hand, we have the distance measure proposed by Zuccon et al. (2009), in which only the maximum dimensionality between the two subspaces compared is taken into account to obtain a distance measure that takes all the vectors of both bases to be calculated (Equation 9), so the result will be the same, regardless of the order in which the concepts are evaluated. This simultaneous measure is calculated as follows:

$$d_s(S_u, S_v) = \sqrt{max(m,n) - \sum_{i=1}^{n}\sum_{j=1}^{m}\langle u_i|v_j\rangle}, \tag{9}$$

where $\{|u_1\rangle, ..., |u_n\rangle\}$ is the basis for S_u, and $\{|v_1\rangle, ..., |v_n\rangle\}$ is the basis for S_v. Thus, we make use of the latter measure of distance in this paper to account for the Triangular Inequality effect.

Considering similarity measures, we have the measures proposed by Pothos & Busemeyer (2011) in the QSM that we have already discussed previously. However, these measures are dynamic too, so in this case we must use a static measure as also proposed by Zuccon et al. (2009). Starting from Equation 9, Zuccon et al. (2009) define a static similarity measure between subspaces as follows:

$$sim(S_u, S_v) = 1 - \frac{d_s(S_u, S_v)}{\sqrt{max(m,n)}}, \tag{10}$$

We use the following concepts in order to show the Triangular Inequality effect: *United States (US)*, *Palestine*, and *Israel* [3]. If we use the vector representation of the concepts from the container space, as expected, the violation of Triangular Equality can be observed with similarity measures such as cosine, being $Sim(US, Palestine) = 0.048$, and $Sim(Israel, US) + Sim(Israel, Palestine) = 1.023$. On the other hand, using Euclidean or subspaces distance measures this violation does not occur, being for the Euclidean distance $Dist(US, Palestine) = 11.178$, and $Dist(US, Israel) + Dist(Israel, Palestine) = 13.3$, and for the subspaces distance $Sdist(US, Palestine) = 0.935$, and $Sdist(US, Israel) + Sdist(Israel, Palestine) = 1.125$.

Let us see what happens if we use multidimensional subspaces to determine the similarity and distance measures. Using our method, we extract sixteen dimensions for the *US*, three dimensions for *Palestine*, and eight dimensions for *Israel*. First we see how the Triangular Inequality assumption is again violated by taking the Zuccon et al. (2009) measure of similarity between subspaces, being $Sim(US, Palestine) = 0.007$, and $Sim(Israel, US) + Sim(Israel, Palestine) = 0.011$. However, for subspace distance measures, Triangular Inequality is not violated, being $Sdist(US, Palestine) = 4.031$, and $Sdist(Israel, US) + Sdist(Israel, Palestine) = 6.799$. In this scenario, the Triangular Inequality remains valid for measuring distances between subspaces. However, it is important to emphasize that, although the Triangular Inequality is a fundamental assumption for distance measures using Euclidean distance, this is not necessarily the case for measuring distances between subspaces or conceptual distances, being theoretically possible to violate this assumption using this framework.

4.3. Diagnosticity Effect

Finally, we study how to formalize the Diagnosticity effect reported by Tversky (1977), using the QSM. The Diagnosticity effect consists of introducing a context effect given by a concrete concept when calculating the similarity between another pair of concepts. Thus, when calculating the similarity between concepts A and B, we must consider the influence of the contextual information determined by concept C. In a famous study proposed by Tversky, subjects were presented with two different experimental tasks. In the first experiment, two set of individuals were asked to group different countries, changing from a sequence $\{a,b,p,c\}$ to a sequence $\{a,b,q,c\}$. We see an example of this task in Figure 5. When the alternatives were *Sweden*, *Poland*, and *Hungary*, the majority of participants grouped *Austria* with *Sweden* (49%), and *Poland* with *Hungary*. However, if presented with the options *Sweden*, *Norway*, and *Hungary*, *Hungary* was grouped with *Austria* in most of the cases (60%) and not *Sweden* (14%). In a second experiment, Tversky asked two different groups of participants with different contextual information (*Poland* or *Norway*) to select the country more related to the *target* (in this case *Austria*), and could demonstrate that the proportion of participants that selected a specific country, like *Sweden*, changed on average a 9% depending on the context. This doesn't mean that the participants changed the grouping strategy as in the first task.

A version of Equation 2, presented above for calculating the similarity between two concepts is used by Duran et al. (2016) in a first proposal to formalize the Diagnosticity effect in terms of similarities given a context. Taking Duran et al. (2016) formalization, to estimate the similarity between *Austria* and *Sweden* in the *Poland* context, we could use the following formula:

$$\begin{aligned} Sim(Austria, Sweden) = ||\mathbf{P}_{Sweden}\mathbf{P}_{Poland}|\psi'\rangle||^2 = \\ ||\mathbf{P}_{Sweden}|\psi'_{Poland}\rangle||^2 \cdot ||\mathbf{P}_{Poland}|\psi'\rangle||^2, \end{aligned} \tag{11}$$

ψ' being the *Austria* vector in the $\mathcal{H}$ container space. Please note that ψ' is no longer considered a neutral vector between *Austria* and *Sweden*. Now, ψ' is considered the vector representing *Austria* in the container space. In sum, to estimate the similarity between *Austria* and *Sweden* in the *Poland* context, *Austria* is taken as a vector from the container space, and it is projected into the subspace representing *Poland* and then projected into the subspace representing *Sweden*. *Austria* is therefore a coordinate, and the two other countries are multidimensional subspaces created using our method.

[3]Since the geopolitical context is not the same as in 1977, different countries have been used to exemplify this effect.

		a Austria	
Set 1	b Sweden 49%	p Poland 15%	c Hungary 36%
		a Austria	
Set 2	b Sweden 14%	q Norway 26%	c Hungary 60%

Figure 5: Figure taken from Tversky (1977) work representing the diagnosticity effect with two sets of countries. The percentage of subjects who selected each country (as grouped with Austria) is presented below the country.

This way, we can perform the procedure using multidimensional subspaces, generating these similarities from the point of view of the QSM to check how the context effect works with this type of space. We will need to extract the subspaces for *Poland* (13 components), *Sweden* (13 components), *Norway* (19 components), and *Hungary* (12 components). Given *Poland* as an intermediate context, the similarity between *Austria* and *Sweden* is 0.692 and the similarity between *Austria* and *Hungary* is 0.724. Given *Norway* as an intermediate context, the similarity between *Austria* and *Sweden* is 0.707 and the similarity between *Austria* and *Hungary* is 0.692. The results indicate that the Diagnosticity effect is not replicated in a vector space model such as LSA with multidimensional spaces using the Duran et al. (2016) proposal. In both cases, the context is not decisive in generating a change in the similarity between *Austria* and the Swedish-Hungarian pair. Indeed, using a context closer to *Sweden* produces the so-called attraction effect (Yearsley et al., 2021). This effect is detailed below.

Yearsley et al. (2021) propose the existence of two differentiated effects that would come into competition in this type of forced-choice task given a contextual stimulus: the *grouping effect* and the *attraction effect*. The grouping effect occurs when, in the presence of a distractor between an initial stimulus and two target stimuli, the subject tends to group the distractor with one of the target stimuli and thus groups the initial stimulus with the remaining target stimulus. The attraction effect occurs when, in the same experimental situation, the subject tends to group the initial stimulus with both the distractor and one of the target stimuli, generating an attraction effect by the distractor. It is worth to mention that these effects are congruent with the behaviour observed in Tversky's second experimental task, since in the first task the grouping strategy is being forced. In this paper we propose two models that try to approach the experimental reality from different perspectives, formalizing the tentative processes that could be operating to make that Diagnosticity effect arise. We again implement these formalizations in a data-driven environment. We name the first process as *contextualization* and the second as *perplexity*.

4.3.1. Contextualized Similarity

This approach, more congruent with the sencond experimental task of Tversky study, consists of representing the degenerate concepts of the container space $\mathcal{H}$ as concepts contextualized in a subspace; that is, to transform the $300-D$ vector that represents a word in $\mathcal{H}$ into the basis of a concept subspace. It is just a change of basis from a $300-D$ to a $k-dimensional$ subspace that represents a context. A degenerate vector of the container space is instantiated from the point of view that is represented by a concept subspace. We determine a context through a concept subspace to represent the previously *degenerate* vectors in the ground state in this context C as:

$$\langle \mathcal{H}_C | A_i \rangle, \tag{12}$$

where $\mathcal{H}_C$ is the context subspace and A_i is the *ith* word label representation of a concept in the container space ($\mathcal{H}$). This allows us to obtain the similarities between contextualized words; that is, similarities between vectors that have a specific meaning due to a context. The contextualized similarity is calculated as follows:

$$||\langle \mathcal{H}_C|A_1\rangle \cdot \langle \mathcal{H}_C|A_2\rangle||^2 \tag{13}$$

The previous formula could be interpreted as the similarity between the concept A_1 within the context C and the concept A_2 within the context C. For example, having the two contexts in the Tversky task, *Poland* and *Norway*, we estimate the concept subspace for both. The resulting subspaces are $\mathcal{H}_P$ and $\mathcal{H}_N$, respectively. Now, we can represent the ***Austria*** vector in the ***Polish*** context as $\langle \mathcal{H}_P|Austria\rangle$, being the similarity between *Austria* and *Sweden* $Sim(Austria, Sweden) = ||\langle \mathcal{H}_P|Austria\rangle \cdot \langle \mathcal{H}_P|Sweden\rangle||^2$. The context of ***Poland*** envelops both ***Sweden*** and ***Austria*** when similarity is calculated. In the ***Polish*** subspace, we obtain that $sim(Austria, Sweden) = 0.765$, and $sim(Austria, Hungary) = 0.805$. On the other hand, in the *Norway* subspace we obtain that $sim(Austria, Sweden) = 0.748$, and $sim(Austria, Hungary) = 0.827$. In these results we can observe the two effects previously explained, with an attraction effect occurring between *Austria* and *Hungary* in the ***Polish*** context, and a grouping effect again between *Austria* and *Hungary* in the *Norwegian* context. In this way we expect that, if the distractor and the main stimulus are related to similar characteristics of the context, this will produce an attraction effect, whereas in case the distractor has a strong relationship with the characteristics of the context and the main stimulus does not have this same type of relation, a grouping effect with the remaining concept will be produced.

4.3.2. Similarity after Perplexity

The second approach to model the Diagnosticity effect is focused on formalizing the grouping behaviour of the first task in Tversky's study. Returning to Tversky's experimental situation, when we are forced to select one of the three countries, *Sweden*, *Poland*, and *Hungary*, to be grouped with *Austria*, people tend to group *Poland* and *Hungary* and choose *Sweden*. The main assumption of this model is that the subject perceives two similar alternatives (*Poland* and *Hungary*) and discards the common properties of these two to make a judgment about the similarity between the initial stimulus (*Austria*) and the third element (*Sweden*). We formalize this fact in three distinct phases:

1. *Grouping*: the subject identifies the most plausible pair of alternatives (*Poland* and *Hungary*) among all possible sequences of concepts presented (*Poland*, *Sweden*, and *Hungary*). Since each country is represented as a concept subspace, in order to identify that pair we use the Zuccon et al. (2009) distance measure between subspaces (Equation 9). The subject identifies the smallest distance between two countries and forms the most plausible pair. This generates additional perplexity between that pair and causes other characteristics to be used to compare the rest of the concepts. In other words, the common features of the pair are inhibited when similarity is estimated. This inhibition takes place in the next phase.
2. *Discarding perplexed properties*: once the most plausible pair is identified, the subject inhibits the common features existing in that pair of subspaces. To estimate the similarity between the initial stimulus *Austria* and the remaining target stimuli *Poland*, *Sweden*, and *Hungary*, perplexed properties are inhibited from *Austria*. To formalize this process, we search for the intersection between the most plausible pair of subspaces using the method of Ben-Israel (2015). This intersection represents the perplexed properties. The Ben-Israel (2015) method takes the singular value decomposition (SVD) of the product of the projectors ($\mathbf{P}_B$ and $\mathbf{P}_C$) of the subspaces B and C (being these two subspaces the most plausible pair) as follows:

$$\mathbf{P}_B\mathbf{P}_C = X\Sigma Y^*, \tag{14}$$

and then computes the projector representing the intersection of the two subspaces as:

$$\mathbf{P}_{\mathbf{B}\cap\mathbf{C}} = \sum_{i=1}^{s} \mathbf{x}_i\mathbf{x}_i^*. \tag{15}$$

Thus, once the perplexed properties of the most plausible pair have been identified by means of the intersection, we subtract the projector of this intersection from the projector of the target stimulus (*Austria*) as follows:

$$\mathbf{P'_A} = \mathbf{P_A} - \mathbf{P_{B\cap C}}. \tag{16}$$

Now, we have inhibited perplexed properties from the *Austria* subspace. There is a dead angle in the subspace of *Austria* represented by $\mathbf{P}'_A$. This fact has consequences in the similarity calculation of the next phase.

3. *Decision*: finally, the subject will analyze the potential relationship between the updated target stimulus (*Austria*) and the rest of the options as:

$$\begin{aligned} &||\mathbf{P}_B\mathbf{P}'_A|\psi\rangle||^2 \\ &||\mathbf{P}_C\mathbf{P}'_A|\psi\rangle||^2 \\ &||\mathbf{P}_D\mathbf{P}'_A|\psi\rangle||^2 \end{aligned} \tag{17}$$

Although, as we have warned on several occasions, the results of applying this approach with the current corpus are largely arbitrary, we present some preliminary results in Table 2 using the above example of the Diagnosticity effect.

Table 2: Diagnosticity effect with Poland as a distractor. In this case the most plausible pair is formed by Poland and Hungary.

Austria	Similarity	Austria Perplexed	Similarity
$\|\|\mathbf{P}_{Poland}\mathbf{P}_{Austria}\|\psi\rangle\|\|^2$	0.714	$\|\|\mathbf{P}_{Poland}\mathbf{P}'_{Austria}\|\psi\rangle\|\|^2$	0.014
$\|\|\mathbf{P}_{Hungary}\mathbf{P}_{Austria}\|\psi\rangle\|\|^2$	0.700	$\|\|\mathbf{P}_{Hungary}\mathbf{P}'_{Austria}\|\psi\rangle\|\|^2$	0.012
$\|\|\mathbf{P}_{Sweden}\mathbf{P}_{Austria}\|\psi\rangle\|\|^2$	0.678	$\|\|\mathbf{P}_{Sweden}\mathbf{P}'_{Austria}\|\psi\rangle\|\|^2$	0.027

As we can see in the results, eliminating the intersection between *Poland* and *Hungary* from the *Austria* projector, the similarity between *Austria* and *Sweden* (0.027) increases with respect to the other two countries (0.014 and 0.012 respectively). Thus, it is possible to formalize a grouping effect by removing information shared by two subspaces from a third subspace. In our example, we have subtracted the information of the conjunction of the *Hungarian* and *Polish* subspaces from the *Austrian* subspace, thus obtaining a higher similarity of *Austria* with *Sweden* after this transformation.

5. Conclusions

We consider that the idea of using subspaces to represent word meanings is broadly consistent with the MTT (Žolkovskij & Mel'čuk, 1967a,b; Mel'čuk, I., 1974), and can be seen as a way of providing a geometric representation of the relationships between words, based on their semantic properties. In MTT, words are seen as carriers of meaning, and the meaning of a word is determined by its relationship with other words in a given context. This is similar to the idea of using the context or "contour" of a word in a LSA space to generate a subspace representation, as the subspaces capture the relationships between different contexts or usage examples of the word. The main advantage of representing contextualized concepts with QSM (with a state vector and a sequential projections) is that it permits to account for the formation of concepts in a natural way, including the separations from prototypical structures (Rosch, 1975). The prototype of China, for example, can be seen as the area of the subspace of China (and the potential vectors within that subspace) that is activated by the most frequently occurring contexts. However, the concept of China can also be instantiated far from this prototype due to varying contexts (e.g. "Chinese companies in the US tech industry"). It is well understood that concepts are not homogeneous and discrete classes, and Rosch's Prototype Theory posits categories as heterogeneous and non-discrete classes, with certain members being more representative of the category than others, known as prototypes. Nevertheless, it is not feasible to define a category by necessary and sufficient conditions. Therefore, a prototype is not the exemplar with the necessary and sufficient conditions, but rather the sum of the typical features of the category's exemplars (Rosch, 1975). The definition of the prototype has been criticized, as it is an assembly of a category's conditions and, taking into account the variability of the exemplars that could belong to it, could result in a very atypical and aberrant prototype that nobody recognizes. This may be the case for "homonymous" categories. For example, the prototype of the "clock" category may include "the tower of London". In contrast, the Quantum Similarity Model is more versatile in dealing with prototypes since conditions are highly contextualized. Each concept is defined in a subspace with a basis that spans all potential categories that it could belong to. The category of "the tower of London" is contextualized because it depends on the previous context in the sequence. The "tower of London" in the context of time is not the same as "the tower of London" in the context of buildings, they are different vectors located in different areas of the subspace. As a result, the prototype of a category never assembles a unique and static "tower of London", this is impossible in terms of the Quantum Similarity models. Therefore, the theory that supports Quantum Similarity is context-centered. In contrast, with the Rosch simple assembly, it is not possible to have an exemplar in an area of the subspace not related to a context. We can consider that Rosch's prototype model arises naturally from the Quantum Similarity Model.

In this same line, "conceptual subspaces theory" proposed by Gärdenfors (2004) argues that concepts can be represented by subspaces in a multidimensional space, where each subspace contains features that are relevant to a particular context. This allows for a more flexible and contextualized representation of concepts, which is in line with our findings that the QSM is capable of accounting for the formation of concepts in a natural way, including separations from prototypical structures. Moreover, Gardenfors' conceptual subspaces theory offers an advance over the prototype theory by providing a more nuanced and precise representation of concepts, allowing for the representation of subtle variations and differences between different instances of a category. This is particularly useful in the context of natural language processing, where words can have multiple meanings and contexts can drastically affect their interpretation. This way, we state that QSM and conceptual subspaces theory offer complementary approaches to understanding the nature of concepts, and their integration can lead to a more comprehensive understanding of how concepts are formed and processed in the human mind.

In the present study, we propose a formal implementation of the Quantum Similarity Model (QSM) within the rails of a semantic-vector space model like the LSA. To achieve this aim, a new method for generating subspaces has been proposed. In addition, different cognitive biases were used to illustrate the usefulness of this method to emulate context influence in a more natural way. We have proposed some ways to formalize the Asymmetry, Triangular Inequality, and Diagnosticity effects to exemplify the possibilities of the QSM over a data-driven model of representation as LSA.

With respect to the Asymmetry effect, the order of the projections and the dimensionality of the concept subspaces seem to be the most important variables for the results. The prototype and the variant of the studied pairs are determined according to the dimensionality obtained in the subspaces, the prototype being always the term with the highest dimensionality. This difference in dimensionality has proved to be determinant in finding an asymmetry between the similarities in pairs of concepts, depending on the order. Regarding the Triangular Inequality violation, we must say that in previous theoretical models, the similarity is interpreted as an inverse function of the distance in geometric models (something extensively discussed in Tversky's work). This is not the case with the Euclidean Distance and the Cosine in S-VSMs. For this reason, Triangular Inequality seems to be violated even with simple cosines, without the need to apply either the quantum model or the generation of subspaces (please note that such was not the case for Euclidean Distances). In any case, we have also shown how to calculate simultaneous distances between subspaces to emulate the Triangular Inequality effect in the framework of the QSM and LSA. Finally, we have proposed some ways to formalize some of the cognitive processes that could be involved in the Diagnosticity effect, using the QSM approach and the generation of concept subspaces; for example, the attraction effect for similarity evaluation or the consequences of grouping a pair of stimuli, as described by Tversky (1977). For this purpose, we present two new similarity measurements, 'contextualized similarity' and 'similarity after perplexity'. We demonstrate how these approaches effectively work with Tversky's experimental tasks. They allow us to better understand the complex relationship between context, perplexity, and similarity.

Overall, this paper is an illustration of how to work with data-driven subspaces in the QSM. The added value is that it offers a data-driven method to manage concepts as potential areas with some conceptual and mathematical tools. In future research, we aim to collect experimental data from human subjects in order to validate the effectiveness of the proposed methods in simulating human biases within the context of natural language processing. Although with some limitations, this work opens up a wide range of possibilities for new research in the field, from the refinement of methods for generating concept subspaces in vector space model contexts to the development of explanatory language models using subspaces to represent concepts.

References

Abbott, B. (1999). Formal semantics and pragmatics for natural language understanding. Journal of Foreign Languages, 23(1), 1-23.

Aerts, D., Broekaert, J., Sozzo, S., and Veloz, T. (2013, July). Meaning-focused and quantum-inspired information retrieval. In International Symposium on Quantum Interaction (pp. 71-83). Springer, Berlin, Heidelberg.

Aerts, D., and Czachor, M. (2004). Quantum aspects of semantic analysis and symbolic artificial intelligence. Journal of Physics A: Mathematical and General, 37(12), L123.

Aerts, S., Kitto, K., & Sitbon, L. (2011, June). Similarity metrics within a point of view. In International Symposium on Quantum Interaction (pp. 13-24). Springer, Berlin, Heidelberg.

Balkenius, C., & Gärdenfors, P. (2016). Spaces in the brain: From neurons to meanings. Frontiers in Psychology, 7, 1-12.

Baroni, M., & Lenci, A. (2009, March). One distributional memory, many semantic spaces. In Proceedings of the Workshop on Geometrical Models of Natural Language Semantics (pp. 1-8).

Ben-Israel, A. (2015). Projectors on intersection of subspaces. Contemporary Mathematics, 636, 41-50.

Batchelder, W. H., Colonius, H., & Dzhafarov, E. N. (Eds.). (2018). New Handbook of Mathematical Psychology: Volume 2, Modeling and Measurement. Cambridge University Press.

Bhatia, S. (2017). Associative judgment and vector space semantics. Psychological Review, 124(1), 1.

Blacoe, W., Kashefi, E., and Lapata, M. (2013, June). 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 (pp. 847-857).

Blei, D. M., Ng, A. Y., & Jordan, M. I. (2003). Latent Dirichlet allocation. Journal of machine Learning research, 3(Jan), 993-1022.

Browne, M. W. (2001). An overview of analytic rotation in exploratory factor analysis. Multivariate behavioral research, 36(1), 111-150.

Bruza, P., Busemeyer, J., and Gabora, L. (2013). Introduction to the special issue on quantum cognition. arXiv preprint arXiv:1309.5673.

Bruza, P. D., and Cole, R. J. (2006). Quantum logic of semantic space: An exploratory investigation of context effects in practical reasoning. arXiv preprint quant-ph/0612178.

Busemeyer, J. R., Pothos, E. M., Franco, R., and Trueblood, J. S. (2011). A quantum theoretical explanation for probability judgment errors. Psychological review, 118(2), 193.

Busemeyer, J. R., and Bruza, P. D. (2012). Quantum models of cognition and decision. Cambridge University Press.

Chwilla, D. J., & Kolk, H. H. (2002). Three-step priming in lexical decision. Memory & Cognition, 30(2), 217-225.

Deerwester, S., Dumais, S. T., Furnas, G. W., Landauer, T. K., & Harshman, R. (1990). Indexing by latent semantic analysis. Journal of the American Society for Information Science, 41(6), 391-407.

Devlin, J., Chang, M. W., Lee, K., & Toutanova, K. (2018). Bert: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint arXiv:1810.04805.

Barque Duran, A., Pothos, E. M., Yearsley, J. M., Hampton, J. A., Busemeyer, J. R., and Trueblood, J. S. (2016). Similarity judgments: from classical to complex vector psychological spaces. In Contextuality from Quantum Physics to Psychology (pp. 415-448).

Gabora, L., & Aerts, D. (2002). Contextualizing concepts using a mathematical generalization of the quantum formalism. Journal of Experimental & Theoretical Artificial Intelligence, 14(4), 327-358.

Gabora, L., Rosch, E., & Aerts, D. (2008). Toward an ecological theory of concepts. Ecological Psychology, 20(1), 84-116.

Gärdenfors, P. (1996). Conceptual spaces as a basis for cognitive semantics. In Philosophy and Cognitive Science: Categories, Consciousness, and Reasoning (pp. 159-180). Springer Netherlands.

Gardenfors, P. (2004). Conceptual spaces: The geometry of thought. MIT press.

Gonzalez, F. A., and Caicedo, J. C. (2011, September). Quantum latent semantic analysis. In Conference on the Theory of Information Retrieval (pp. 52-63). Springer, Berlin, Heidelberg.

Gorsuch, Richard L. "Exploratory factor analysis." Handbook of multivariate experimental psychology. Springer, Boston, MA, 1988. 231-258.

Harati, P., Westbury, C., & Kiaee, M. (2021). Evaluating the predication model of metaphor comprehension: Using word2vec to model best/worst quality judgments of 622 novel metaphors. Behavior Research Methods, 1-12

Hoffman, P., Lambon Ralph, M. A., & Rogers, T. T. (2013). Semantic diversity: A measure of semantic ambiguity based on variability in the contextual usage of words. Behavior Research Methods, 45(3), 718–730.

Horn, J. L. (1965). A rationale and test for the number of factors in factor analysis. Psychometrika, 30(2), 179-185.
Jaiswal, A. K., Holdack, G., Frommholz, I., and Liu, H. (2018, September). Quantum-like generalization of complex word embedding: a lightweight approach for textual classification. CEUR Workshop Proceedings.
Jones, M., Gruenfelder, T., and Recchia, G. (2011). In defense of spatial models of lexical semantics. In Proceedings of the annual meeting of the cognitive science society (Vol. 33, No. 33).
Jorge-Botana, G., León, J. A., Olmos, R., & Escudero, I. (2011). The representation of polysemy through vectors: some building blocks for constructing models and applications with LSA. International Journal of Continuing Engineering Education and Life Long Learning, 21(4), 328-342.
Jorge-Botana, G., León, J. A., Olmos, R., & Hassan-Montero, Y. (2010). Visualizing polysemy using LSA and the predication algorithm. Journal of the Association for Information Science and Technology, 61(8), 1706-1724.
Jorge-Botana, G., and Olmos, R. (2014). How lexical ambiguity distributes activation to semantic neighbors: Some possible consequences within a computational framework. The Mental Lexicon, 9(1), 67-106.
Jorge-Botana, G., Olmos, R., and Barroso, A. (2013, July). Gallito 2.0: A natural language processing tool to support research on discourse. In Proceedings of the 13th Annual Meeting of the Society for Text and Discourse.
Jorge-Botana, G., Olmos, R., and Luzon, J. M. (2020). Bridging the theoretical gap between semantic representation models without the pressure of a ranking: some lessons learnt from LSA. Cognitive processing, 21(1), 1-21.
Joulin, A., Grave, E., Bojanowski, P., & Mikolov, T. (2016). Bag of tricks for efficient text classification. arXiv preprint arXiv:1607.01759.
Kahneman, D., and Frederick, S. (2002). Representativeness revisited: Attribute substitution in intuitive judgment. Heuristics and biases: The psychology of intuitive judgment, 49, 81.
Kahneman, D., and Tversky, A. (1982). Judgments of and by representativeness (pp. pp-84).
Kintsch, W. (1988). The role of knowledge in discourse comprehension: a construction-integration model. Psychological review, 95(2), 163.
Kintsch, W. (2000). Metaphor comprehension: A computational theory. Psychonomic Bulletin & Review, 7(2), 257-266.
Kintsch, W. (2014). Similarity as a function of semantic distance and amount of knowledge. Psychological review, 121(3), 559.
Kintsch, W., & Bowles, A. R. (2002). Metaphor comprehension: What makes a metaphor difficult to understand? Metaphor and Symbol, 17(4), 249-262.
Kintsch, W., & Mangalath, P. (2011). The construction of meaning. Topics in cognitive science, 3(2), 346-370.
Lake, B., & Baroni, M. (2018, July). Generalization without systematicity: On the compositional skills of sequence-to-sequence recurrent networks. In International conference on machine learning (pp. 2873-2882). PMLR.
Landauer, T. K., and Dumais, S. T. (1997). A solution to Plato's problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. Psychological review, 104(2), 211.
Landauer, T. K., McNamara, D. S., Dennis, S., and Kintsch, W. (Eds.). (2013). Handbook of latent semantic analysis. Psychology Press.
Lemaire, B., & Denhiere, G. (2006). Effects of high-order co-occurrences on word semantic similarity. Current Psychology Letters. Behaviour, Brain & Cognition, 18(1), 1–12.
Li, Q., Uprety, S., Wang, B., & Song, D. (2018). Quantum-inspired complex word embedding. arXiv preprint arXiv:1805.11351.
Lund, K., and Burgess, C. (1996). Producing high-dimensional semantic spaces from lexical co-occurrence. Behavior research methods, instruments, and computers, 28(2), 203-208.
Manning, C., Raghavan, P., & Schütze, H. (2008). Introduction to Information Retrieval. Cambridge: Cambridge University Press. doi:10.1017/CBO9780511809071
Martinez-Mingo, A., Jorge-Botana, G., Olmos, R., & Martinez-Huertas, J. A. (2020, July). Quantum approach for similarity evaluation in LSA vector space models [Paper presentation]. 6th Stochastic Modeling Techniques and Data Analysis International Conference and Demographics 2020, Athens, Greece.
McClelland, J. L. (2022). Capturing advanced human cognitive abilities with deep neural networks. Trends in Cognitive Sciences. Vol. 26, No. 12.
McGregor, S., Agres, K., Rataj, K., Purver, M., & Wiggins, G. (2019). Re-representing metaphor: Modeling metaphor perception using dynamically contextual distributional semantics. Frontiers in psychology, 10, 765.
Mel'čuk, I. (1974). Opyt teorii lingvističeskix modelej Smysl- Tekst. Moskva: Nauka. [Towards a Theory of Meaning-Text Linguistic Models; reprinted in 1999, Moskva: Jazyki russkoj kul'tury]
Mikolov, T., Chen, K., Corrado, G., & Dean, J. (2013). Efficient estimation of word representations in vector space. arXiv preprint arXiv:1301.3781.
Millis, K., & Larson, M. (2008). Applying the construction-integration framework to aesthetic responses to representational artworks. Discourse Processes, 45(3), 263-287.
Mu, J., Bhat, S., & Viswanath, P. (2017). Representing sentences as low-rank subspaces. arXiv preprint arXiv:1704.05358.
Nosofsky, R. M. (1991). Stimulus bias, asymmetric similarity, and classification. Cognitive Psychology, 23(1), 94-140.
Olmos, R., Jorge-Botana, G., Luzon, J. M., Martin-Cordero, J. I., and Leon, J. A. (2016). Transforming LSA space dimensions into a rubric for an automatic assessment and feedback system. Information Processing and Management, 52(3), 359-373.
Osta-Vélez, M., & Gärdenfors, P. (2022). Nonmonotonic reasoning, expectations orderings, and conceptual spaces. Journal of Logic, Language and Information, 1-21.
Pannitto, L., & Herbelot, A. (2022). Can recurrent neural networks validate usage-based theories of grammar acquisition? Frontiers in Psychology, 13.
Karl Pearson F.R.S. . (1901). LIII. On lines and planes of closest fit to systems of points in space. The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science, 2(11), 559–572.
Pennington, J., Socher, R., & Manning, C. D. (2014). Glove: Global vectors for word representation. In Proceedings of the

2014 conference on empirical methods in natural language processing (EMNLP) (pp. 1532-1543).
Peters, M. E., Neumann, M., Iyyer, M., Gardner, M., Clark, C., Lee, K., & Zettlemoyer, L. (2018). Deep contextualized word representations. arXiv 2018. arXiv preprint arXiv:1802.05365, 12.
Pothos, E. M., Barque-Duran, A., Yearsley, J. M., Trueblood, J. S., Busemeyer, J. R., and Hampton, J. A. (2015). Progress and current challenges with the quantum similarity model. Frontiers in psychology, 6, 205.
Pothos, E., and Busemeyer, J. (2011). A quantum probability explanation for violations of symmetry in similarity judgments. In Proceedings of the Annual Meeting of the Cognitive Science Society (Vol. 33, No. 33).
Pothos, E. M., Busemeyer, J. R., and Trueblood, J. S. (2013). A quantum geometric model of similarity. Psychological Review, 120(3), 679.
Quillian, M. R. (1967). Word concepts: A theory and simulation of some basic semantic capabilities. Behavioral science, 12(5), 410-430.
R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. URL https://www.R-project.org/.
Rehder, B., Schreiner, M. E., Wolfe, M. B., Laham, D., Landauer, T. K., and Kintsch, W. (1998). Using Latent Semantic Analysis to assess knowledge: Some technical considerations. Discourse Processes, 25(2-3), 337-354.
Rekabsaz, N., Lupu, M., & Hanbury, A. (2017, April). Exploration of a threshold for similarity based on uncertainty in word embedding. In European Conference on Information Retrieval (pp. 396-409). Springer, Cham.
Rosch, Eleanor (1975). "Cognitive representations of semantic categories". Journal of Experimental Psychology: General. 104 (3): 192-233. doi:10.1037/0096-3445.104.3.192. ISSN 0096-3445.
Schnabel, T., Labutov, I., Mimno, D., & Joachims, T. (2015, September). Evaluation methods for unsupervised word embeddings. In Proceedings of the 2015 conference on empirical methods in natural language processing (pp. 298-307).
Shimomoto, E. K., Portet, F., & Fukui, K. (2021). Text classification based on the word subspace representation. Pattern Analysis and Applications, 24(3), 1075-1093.
Steyvers, M., and Tenenbaum, J. B. (2005). The large-scale structure of semantic networks: Statistical analyses and a model of semantic growth. Cognitive science, 29(1), 41-78.
Sundermeyer, M., Schlüter, R., & Ney, H. (2012). LSTM neural networks for language modeling. In Thirteenth annual conference of the international speech communication association.
Tenenbaum, J. B., and Griffiths, T. L. (2001). Generalization, similarity, and Bayesian inference. Behavioral and brain sciences, 24(4), 629.
Tversky, A. (1977). Features of similarity. Psychological review, 84(4), 327.
Tversky, A., and Kahneman, D. (1981). The framing of decisions and the psychology of choice. science, 211(4481), 453-458.
Tversky, A., and Kahneman, D. (1983). Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment. Psychological review, 90(4), 293.
Von Neumann, J. (2018). Mathematical foundations of quantum mechanics: New edition (Vol. 53). Princeton university press.
Yang, Z., Zhu, C., & Chen, W. (2018). Parameter-free sentence embedding via orthogonal basis. arXiv preprint arXiv:1810.00438.
Yeari, M., & van den Broek, P. (2015). The role of textual semantic constraints in knowledge based inference generation during reading comprehension: A computational approach. Memory, 23, 1193–1214.
Yearsley, J. M., Barque-Duran, A., Scerrati, E., Hampton, J. A., and Pothos, E. M. (2017). The triangle inequality constraint in similarity judgments. Progress in Biophysics and Molecular Biology, 130, 26-32.
Yearsley, J. M., Pothos, E. M., Duran, A. B., and Hampton, J. A. (2015). Diagnosticity: Some theoretical and empirical progress. In CogSci.
Yearsley, J. M., Pothos, E. M., Hampton, J. A., and Duran, A. B. (2014, June). Towards a quantum probability theory of similarity judgments. In International Symposium on Quantum Interaction (pp. 132-145). Springer, Cham.
Yearsley, J. M., Pothos, E. M., Barque Duran, A., Trueblood, J. S., & Hampton, J. A. (2021, in press). Context effects in similarity judgements.
Žolkovskij, A. & Mel'čuk, I. (1967a). O sisteme semantičeskogo sinteza. II: Pravila preobrazovanija [On a System of Semantic Synthesis (of texts). II: Paraphrasing Rules]. Naučno-texničeskaja informacija 2, Informacionnye processy i sistemy, 17-27.
Žolkovskij, A. & Mel'čuk, I. (1967b). O semantičeskom sinteze [On Semantic Synthesis (of texts)]. Problemy kibernetiki 19, 177-238. [French translation: Sur la synthèse sémantique (1970). TA Informations 2, 1-85].
Zuccon, G., Azzopardi, L. A., & Van Rijsbergen, C. J. (2009, March). Semantic spaces: Measuring the distance between different subspaces. In International Symposium on Quantum Interaction (pp. 225-236). Springer, Berlin, Heidelberg.