

Quantum projections on conceptual subspaces: A deeper dive into methodological challenges and opportunities

Science Progress

2024, Vol. 107(2) 1–8


© The Author(s) 2024

Article reuse guidelines:

sagepub.com/journals-permissions

DOI: 10.1177/00368504241238773

journals.sagepub.com/home/sci

Alejandro Martínez-Mingo¹ ,
Jose Ángel Martínez-Huertas²,
Ricardo Olmos¹ and
Guillermo Jorge-Botana³

¹Social Psychology and Methodology, Universidad Autonoma de Madrid, Madrid, Spain

²Department of Methodology of Behavioral Science, Universidad Nacional de Educacion a Distancia, Madrid, Spain

³Department of Psychobiology & Behavioral Sciences Methods, Universidad Complutense de Madrid, Madrid, Spain

Abstract

In alignment with the distributional hypothesis of language, the work “Quantum Projections on Conceptual Subspaces” (Martínez-Mingo A, Jorge-Botana G, Martínez-Huertas JA, et al. Quantum projections on conceptual subspaces. *Cogn Syst Res* 2023; 82: 101154) proposed a methodology for generating conceptual subspaces from textual information based on previous work (Martínez-Mingo A, Jorge-Botana G and Olmos R. Quantum approach for similarity evaluation in LSA vector space models. 2020). These subspaces enable the utilization of the quantum model of similarity put forth by Pothos and Busemeyer (Pothos E, Busemeyer J. A quantum probability explanation for violations of symmetry in similarity judgments. In *Proceedings of the annual meeting of the cognitive science society*, 2011, Vol. 33, No. 33), allowing for the empirical examination of the violations of assumptions concerning symmetry and triangular inequality (Tversky A. Features of similarity. *Psychol Rev* 1977; 84: 327–352; Yearsley JM, Barque-Duran A, Scerrati E, et al. The triangle inequality constraint in similarity judgments. *Prog Biophys Mol Biol* 2017; 130: 26–32), as well as the diagnosticity effect (Tversky A. Features of similarity. *Psychol Rev* 1977; 84: 327–352; Yearsley JM, Pothos EM, Barque-Duran A, et al. Context effects in similarity

Corresponding author:

Alejandro Martínez-Mingo, Social Psychology and Methodology, Universidad Autonoma de Madrid, Madrid, Spain.

Email: alejandro.martinez@uam.es



Creative Commons Non Commercial CC BY-NC: This article is distributed under the terms of the Creative Commons Attribution-NonCommercial 4.0 License (<https://creativecommons.org/licenses/by-nc/4.0/>)

which permits non-commercial use, reproduction and distribution of the work without further permission provided the original work is attributed as specified on the SAGE and Open Access page (<https://us.sagepub.com/en-us/nam/open-access-at-sage>).

judgments. *J Exp Psychol Gen* 2022; 151: 711–717), within a *data-driven* environment. These psychological biases, deeply studied by authors such as Tversky and Kahneman, inform us about the limitations of modeling psychological similarity measures using tools from classical geometry. This commentary aims to offer methodological clarifications, discuss theoretical and practical implications, and speculate on future directions in this field of research. Concretely, it aims to propose the use of different contours (conceptual or contextual) to generate the subspaces, which lead to subspaces of terms or contexts. Once these contours are defined, a differentiation is proposed between *Aggregated Terms Subspaces* (ATSs), *Aggregated Contexts Subspaces* (ACSs), and *Aggregated Features Subspaces* (AFSs) depending on whether we define the subspaces by grouping the terms or contexts within the contour, or from the latent dimensions of the semantic space obtained in the contour window. Finally, new data is provided on the violation of the triangular inequality assumption through the application of the quantum similarity model to ATSs.

Keywords

Natural language processing, computational linguistics, cognitive processes, quantum cognition, artificial intelligence

Introduction

Recently, we have witnessed a genuine revolution in computational language studies. Large language models (LLMs) such as the GPT family models have fundamentally altered the landscape, outperforming all prior expectations. Intriguingly, this leap did not result from a significant advance in academic research—although the merit of Transformer⁷ architectures should not be understated. Instead, it has primarily been a matter of scaling up existing pre-trained language models (PLMs), supplemented by a suite of optimization and training methodologies tailored for commercial applications.⁸ This scaling, in terms of both training data and model parameters, has spawned emergent behaviors in LLMs that have led many in the scientific community to reconsider the feasibility of general artificial intelligence in the not-so-distant future. Specifically, *In-Context Learning* (the ability to adapt its responses based on the specific context provided within a single interaction, rather than learning from a historical database of past interactions) and *Chain of Thoughts* (the sequential process by which these models generate text, building upon each step to form coherent and contextually relevant outputs) stand out as unique behaviors that have arisen from the scaling of PLMs.⁸ However, a number of studies to date have served as a sobering counterpoint,^{8,9} urging us to temper our expectations. According to these scholars, we are still far removed from a computational model capable of mirroring our linguistic abilities. To approach that goal, it is imperative to explore the alignment of existing solutions with human cognitive biases.⁹

With this overarching goal, we are keen to provide an academic commentary on the paper “Quantum Projections on Conceptual Subspaces,”¹ which implicitly raises a research question that challenges the very foundations of distributional language models. Our discussion leans on the seminal experiments conducted by Tversky about human cognitive biases like the violation of the assumptions of symmetry and triangle inequality in cognitive similarity assessments.⁴ In his groundbreaking 1977 work, Tversky posed an exceptionally pertinent question: are spatial models adequately suited

for assessing psychological similarity? Given Tversky's findings, the answer appears to be a resounding no. Building upon these insights, researchers like Pothos, Busemeyer, and Yearsley^{3,5,6} have proposed extending the conventional geometric model to a subspace-based model. This extension enables the incorporation of mathematical apparatus derived from the axiomatization of quantum mechanics into the study of similarity.³ It is within this context that an important research question concerning computational language models emerges: Are we employing the correct geometry? This leads us to hypothesize that the classical geometric model will be insufficient if the goal is to emulate, in an ecologically valid manner, the cognitive processes inherent in human language. Studying perspectives that merge logic with conceptual subspaces could inspire better optimization methodologies over the architectures that allow LLMs to work. It could open ways to think new representation and updating mechanisms in order to overcome their sample-based style of learning. Considerations about computational and algorithmic human plausibility in some tasks as understanding metaphors, abstract reasoning, systematic compositionality and contextual adaptability could result in LLMs improvements.

Key points from the article

To the best of our knowledge, there are no existing methodologies that enable the creation of conceptual subspaces within a *container* semantic space in a *data-driven* environment in a manner that allows for the empirical assessment of the quantum similarity model.³ With this in mind, we propose a series of methods for generating such subspaces.^{1,2} Our starting point is a *300-dimensional* semantic space, estimated through *latent semantic analysis* (LSA), using a journalistic corpus extracted from two major Spanish newspapers.¹

Building upon an established semantic space, we introduce an approach for estimating the semantic contour of a target term. This contour serves as a reservoir of key information needed for constructing a subsequent conceptual subspace. The foundational idea for generating such a conceptual subspace is to create it based on terms that are semantically related to this target term, which we define as the semantic contour of the target term.² To ascertain this pertinent information, we suggest utilizing the similarity function between the target term and all terms in the corpus, which is then contrasted with the similarity function of an abstract, intermediary term within the semantic space, coined as the *super-term*—an entity derived from the vectorial summation of all term vectors within the corpus. The intersection point between these two similarity functions informs the selection of the number of neighboring terms that constitute the contour of the target term. Moreover, the super-term's similarity function is weighted by the semantic diversity of the target term, thereby enhancing the information capture for terms that exhibit greater semantic diversity within the corpus.¹

Once the contour of the target concept is created, we delineate the corresponding subspace by undertaking dimensionality reduction. This aims to capture orthogonal components that account for the maximum variance within the contour, thereby ensuring that the semantic terms defining the contour are represented by a reduced number of dimensions that effectively summarize and simplify the semantic contents of the contour, and consequently, the target concept. The selection of the subspace dimensions is guided by the

Parallel Analysis technique. Here, the deserved number of dimensions is chosen by comparing the variance explained by the contour's components against that explained by components in a random matrix of equivalent dimensions. Thus, the resulting representation subspace for a target concept will consist of K normalized, orthogonal base vectors, each of which retains a representation in the original overarching semantic space.¹

In the paper under discussion, the proposed methodology successfully accounts for asymmetry biases and diagnosticity effects. Regarding the violation of the triangle inequality, it holds in terms of similarity within both the classical geometric model and the quantum model. Importantly, no violations of this assumption in terms of distances were observed with the stimuli utilized across either model.¹ This commentary has a multifaceted aim: to provide methodological clarifications, delve into theoretical and practical implications, and offer informed speculation on the trajectory of future research in this domain. More specifically, the article advocates for employing various types of contours—either conceptual or contextual—as a basis for generating semantic subspaces. Once the contours are clearly defined, an analytical distinction is introduced between three types of subspaces: aggregate term subspaces (ATSs), aggregate context subspaces (ACSs), and aggregate feature subspaces (AFSs). Lastly, the commentary introduces new data that disrupts traditional geometric assumptions, specifically the triangular inequality assumption.

A deeper dive into methodological challenges and opportunities

Firstly, it is important to clarify that our choice of latent semantic analysis (LSA) as the technique for constructing the container semantic space is dictated by the resultant vector space's nature, which ensures orthogonality of the generated latent dimensions via singular value decomposition (SVD). This orthogonality constraint was initially adopted to ensure the container space's conformity to a Hilbert space—being this orthogonality a desirable but not essential feature for this type of vector spaces.¹⁰ While this condition does confer interpretability to the system when it collapses into a basis vector, the latent dimensions in LSA are not inherently interpretable, thus adding no additional informative value and merely constraining our container space. Consequently, acknowledging the container space as merely a coordinate system that abstractly represents information (term vectors in the space can be considered *degenerate* states¹¹), it becomes feasible to use any semantically defined vector space as a container space. Alternative approaches could involve predictive models like Word2Vec¹² or even Transformer-based architectures like GPT models or BERT.¹³ In both instances, the embedding space generated when applied to a linguistic corpus will produce a vector space amenable for use as a container space.

Regarding contour generation, the original paper outlines a methodology for composing contours based on the term vectors nearest to the target term (*conceptual contour*). It is also conceivable to define contours based on the vectors of the contexts where our target term appears (*contextual contour*). In doing so, we can construct both *saturated* and *non-saturated* contours depending on whether we include all contextual vectors where the target term appears. For this purpose, representation vectors of contexts (documents or paragraphs) in LSA or Doc2Vec could be utilized, or contextual vectors of the

target term in models like BERT or GPT models. However, employing saturated contextual contours, while exhaustive, is computationally intensive. Hence, for the time being, the use of ad hoc non-saturated contextual contours is recommended. In Table 1 the reader can see an example of the *China* contour defined using BERT embedding as the container space, using 30 Wikipedia pages that contain the target term to obtain the contextual vectors.

In the final stage of the proposed methodology,¹ Principal Component Analysis (PCA) is employed for contour dimensionality reduction. However, we do not elaborate on the nature of the matrix subjected to decomposition. In the article under review, PCA is applied to the LSA latent dimensions in such a way that the retained eigenvectors, as determined by Parallel Analysis, have a direct representation in the container space, obviating the need for subsequent identification and orthogonalization. Yet, as evident from the findings of this study, these obtained dimensions have limited interpretability.

We want to clarify that, within the conceptual window provided by the conceptual contour of a target term, it will be feasible to perform dimensionality reduction either on the terms, or on the latent features. In this light, we suggest that for each conceptual contour, two types of subspaces could be defined: an *Aggregated Terms Subspace* (ATS) or an *Aggregated Features Subspace* (AFS). Thus, while AFS captures latent information from the space within the window defined by the contour (be it conceptual or contextual), ATS informs us about the aggregation of concepts contained in said contour. These could either constitute conceptual subspaces formed by the clustering of terms in the semantic neighborhood of the target term or contextual subspaces shaped by the aggregation of context vectors where the target term is found, resulting this last approach in an *Aggregated Contexts Subspace* (ACS). Figure 1 illustrates word clouds featuring terms from the conceptual contour most similar to each dimension in both the AFS and ATS for *China*, respectively, given an LSA container space. To summarize, when generating a subspace, it can be determined either by conceptual or contextual information, depending on the composition of the contour. For contours of conceptual vectors, we may choose between ATS or AFS, depending on whether we reduce the contour’s dimensionality with respect to the terms or the latent dimensions of the space, respectively. For contours of contextual vectors, we can opt for ACS or again AFS, depending on whether we reduce the contour’s dimensionality, in this case, with respect to the contexts, or once more with respect to the latent dimensions of the space, respectively.

Finally, we would like to augment this commentary with additional results from applying Aggregated Term Subspaces (ATs) to investigate violations of the triangular

Table 1. Contour of China in a 176 dimensions BERT container space.

Context Vec.	Dim. 1	Dim. 2	Dim. 3	Dim. 4	...	Dim. 768
China 1	-0.58	0.55	-0.20	-0.31	...	0.74
China 2	-0.24	0.76	-0.31	-0.47	...	0.45
China 3	-0.86	0.47	0.38	-0.31	...	0.85
...
China 841	-0.34	0.72	-0.12	-0.11	...	0.40

of quantum mechanics. Although this line of research is still in its early stages, it holds considerable potential, thereby creating numerous opportunities for future studies in this area. Moreover, the empirical testing of the effects outlined in the paper—and their potential relation to the results obtained from the computational application of this methodology—is of utmost importance, but it is essential to define metrics capable of evaluating the performance of the various versions of the proposed method, particularly in the context of standard tasks in natural language processing and semantic analysis. Specifically, one of the main challenges of the quantum similarity model proposed in the commented paper is its potential difficulty in evaluating hierarchical relationships between concepts (subspaces) and their features, which may also be concepts (other subspaces) with their own sets of features. This concept-feature duality can be easily integrated into the model in a recursive fashion, yet a method for representing semantic relationships that denote a certain hierarchy has not been established. The research conducted by Martínez-Huertas et al.¹⁵ may provide insights into this issue by exploring possible higher- and lower-order dimensions in conceptual subspaces.

In a more targeted vein, we see considerable value in undertaking a thorough examination of the characteristics of the various subspaces that can be formed (ATS, ACS, or AFS) at different layers of abstraction in transformer-based models. This could potentially demystify the so-called *black-box* nature of these models.¹⁶ Lastly, the use of subspaces for conceptual representation may pave the way for their dynamic updating through the inclusion of new dimensions as fresh information is integrated into the model. Collectively, these efforts could contribute to a two-fold benefit: first, in clarifying the underlying mechanisms of computational language models, and second, in developing computational models that are potentially more ecologically valid in replicating human cognitive reality.

Declaration of conflicting interests

The authors declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The authors received no financial support for the research, authorship, and/or publication of this article.

ORCID iD

Alejandro Martínez-Mingo  <https://orcid.org/0000-0002-8375-0952>

References

1. Martínez-Mingo A, Jorge-Botana G, Martínez-Huertas JÁ, et al. Quantum projections on conceptual subspaces. *Cogn Syst Res* 2023; 82: 101154.
2. Martínez-Mingo A, Jorge-Botana G and Olmos R. *Quantum approach for similarity evaluation in LSA vector space models*. 6th Stochastic Modeling Techniques and Data Analysis International Conference and Demographics. Athens, Greece, 2020.

3. Pothos E and Busemeyer J. A quantum probability explanation for violations of symmetry in similarity judgments. In Proceedings of the annual meeting of the Cognitive Science Society, 2011, Vol. 33, No. 33.
4. Tversky A. Features of similarity. *Psychol Rev* 1977; 84: 327–352.
5. Yearsley JM, Barque-Duran A, Scerrati E, et al. The triangle inequality constraint in similarity judgments. *Prog Biophys Mol Biol* 2017; 130: 26–32.
6. Yearsley JM, Pothos EM, Barque-Duran A, et al. Context effects in similarity judgments. *J Exp Psychol Gen* 2022; 151: 711–717.
7. Zhao WX, Zhou K, Li J, et al. A survey of large language models. *arXiv preprint arXiv:2303.18223*. 2023.
8. Shanahan M. Talking about large language models. *arXiv preprint arXiv:2212.03551*. 2022.
9. Titus LM. Does ChatGPT have semantic understanding? A problem with the statistics-of-occurrence strategy. *Cogn Syst Res* 2023; 83: 101174.
10. Heil C. *A basis theory primer: expanded edition*. Springer Science & Business Media, 2010.
11. Aerts S, Kitto K and Sitbon L. Similarity metrics within a point of view. In *Quantum interaction: 5th International Symposium, QI 2011, Aberdeen, UK, June 26–29, 2011, Revised Selected Papers 5* (pp. 13–24). Springer Berlin Heidelberg. 2011.
12. Mikolov T, Sutskever I, Chen K, et al. Distributed representations of words and phrases and their compositionality. *Adv Neural Inf Process Syst* 2013; 26.
13. Devlin J, Chang MW, Lee K, et al. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*. 2018.
14. Gabora L and Aerts D. Contextualizing concepts using a mathematical generalization of the quantum formalism. *J Exp Theor Artif Intell* 2002; 14: 327–358.
15. Martínez-Huertas JÁ, Olmos R, Jorge-Botana G, et al. Distilling vector space model scores for the assessment of constructed responses with bifactor inbuilt rubric method and latent variables. *Behav Res Methods* 2022; 54: 2579–2601.
16. Ethayarajh K. How contextual are contextualized word representations? Comparing the geometry of BERT, ELMo, and GPT-2 embeddings. *arXiv preprint arXiv:1909.00512*. 2019.

Author biographies

Alejandro Martínez-Mingo holds a PhD in Cognitive Psychology and Language Models and serves as a professor of Natural Language Processing. His research area lies at the intersection of cognition and computational linguistics, exploring how models from quantum physics can be adapted to elucidate cognitive phenomena.

Jose Ángel Martínez-Huertas is a professor in Formal Models of Cognitive Processes. His research focuses on the development of computational language models and statistical models for explaining psychological processes.

Ricardo Olmos is a professor in Psychometrics and Data Analysis in Psychology. His research interests include computational linguistics, psychometrics in the realm of Classical Test Theory (CTT) and Item Response Theory (IRT), and statistics within linear models.

Guillermo Jorge-Botana is a professor in Supervised Learning. His research area concentrates on the development of computational models for the representation of linguistic and cognitive phenomena.