



Structured representations in a quantum probability model of similarity



Emmanuel M. Pothos^{a,*}, Jennifer S. Trueblood^b

^a Department of Psychology, City University London, London, EC1V 0HB, UK

^b Department of Cognitive Sciences, 3151 Social Sciences Plaza, University of California, Irvine, CA 92697-5100, USA

HIGHLIGHTS

- We consider a model of basic similarity judgments, based on quantum probability principles.
- We augment this model with Smolensky et al. (2014) ideas for structure in representations.
- We show that this proposal can accommodate the main insights regarding structure in similarity judgments.
- We consider the formal properties of our model and discuss the placement of this work in the similarity, analogy literature.

ARTICLE INFO

Article history:

Received 13 February 2014

Received in revised form

15 September 2014

Available online 29 December 2014

Keywords:

Quantum probability theory

Similarity

Representation

ABSTRACT

Recently, Busemeyer et al. (2011) presented a model for how the conjunction fallacy (Tversky & Kahneman, 1983) emerges, based on the principles of quantum probability (QP) theory. Pothos et al. (2013) extended this model to account for the main similarity findings of Tversky (1977), which have served as a golden standard for testing novel theories of similarity. However, Tversky's (1977) empirical findings did not address the now established insight that, in comparing two objects, overlap in matching parts of the objects tends to have a greater impact on their similarity, than overlap in non-matching parts. We show how the QP similarity model can be directly extended to accommodate structure in similarity comparisons. Smolensky's et al.'s (2014) proposal for modeling structure in linguistic representations, with tensor products, can be adapted 'as is' with the QP similarity model. The formal properties of the extended QP similarity model are analyzed, some indicative fits are presented, and, finally, a novel prediction is developed.

© 2014 Elsevier Inc. All rights reserved.

1. Introduction

We call quantum probability (QP) theory the rules for assigning probability to events, without any of the physics (Hughes, 1989; Isham, 1989). QP theory is a framework for probabilistic inference alternative to that of classical probability (CP) theory. A case for adopting QP theory, instead of classical probability (CP) theory, in cognitive modeling has been made when human behavior appears at odds with the prescription from CP theory (e.g., Aerts, 2009; Aerts & Gabora, 2005; Blutner, Pothos, & Bruza, 2013; Nelson, Kitto, Galea, McEvoy, & Bruza, 2013; for overviews see Busemeyer & Bruza, 2012; Haven & Khrennikov, 2013; Khrennikov, 2010; Pothos

& Busemeyer, 2013; Wang, Busemeyer, Atmanspacher, & Pothos, 2013; Yearsley & Pothos, 2014). Recently, Busemeyer, Pothos, Franco, and Trueblood (2011) presented a model of decision making, based on QP principles, with an emphasis on how the conjunction fallacy (Tversky & Kahneman, 1983), and related findings, emerge. For example, in the conjunction fallacy experiment, participants were told of a hypothetical person, Linda, described very much as a feminist (F) and not at all as a bank teller (BT). Participant responses indicated that $Prob(F \wedge BT) > Prob(BT)$, which is impossible classically. In Busemeyer et al.'s (2011) QP model for this conjunction fallacy, if one assumes that the BT, F possibilities are incompatible, then it can emerge that the quantum probability of $F \wedge$ then BT is higher than that of B.

Pothos, Busemeyer, and Trueblood (2013) considered whether the QP decision model could be extended to account for basic similarity judgments. Their motivation was that QP theory is formalized in multidimensional, vector spaces, called Hilbert spaces. The most common, standard way to model basic similarity involves

* Corresponding author.

E-mail addresses: emmanuel.pothos.1@city.ac.uk (E.M. Pothos), jstruebl@uci.edu (J.S. Trueblood).

multidimensional representations (e.g., Shepard, 1987) and the conceptualization of similarity as a function of distance. For example, such models have been employed in the predominant approaches to categorization (e.g., Goldstone, 1994a; Nosofsky, 1984; Wills & Pothos, 2012). Therefore, since QP representations are also geometric (that is, involve elements in some multidimensional vector space), perhaps QP theory can provide some interesting generalization to the standard distance-based similarity models?

Note first that, by basic similarity judgments, we imply ones that are nonanalytic (in the psychological sense), direct, and immediate. If we accept the view that basic similarity judgments can be modeled as some function of distance, then they have to be consistent with the metric axioms—mathematical requirements that all (simple) functions of distances need to obey. These axioms are intuitively appealing. For example, the symmetry axiom, requires that $distance(A, B) = distance(B, A)$, implying that $similarity(A, B) = similarity(B, A)$. In one of the most influential studies in the basic similarity literature, Tversky (1977) showed that *all* metric axioms can be violated in similarity judgments of naïve observers. Tversky's (1977) findings have become a golden standard of empirical results that should be accounted for by any basic similarity model and, indeed, have been the focus of theoretical effort in related research ever since (e.g., Ashby & Perrin, 1988; Krumhansl, 1978). It is worth noting that basic, distance-based similarity metrics can be made to violate the metric axioms. For example, symmetry can be violated if $similarity(A, B) = p_{AB} \cdot distance(A, B)$, where p_{AB} is just a directionality parameter, that is a parameter which can have a different value depending on whether the similarity evaluated is between A and B or between B and A (Nosofsky, 1991). However, the real challenge has been to explore how consistency with Tversky's (1977) findings can emerge from the structure of a basic similarity model.

Pothos et al. (2013) showed how the QP decision model can indeed accommodate Tversky's (1977) key findings, with fairly minor modifications. The objects to be compared in a similarity judgment are represented as subspaces, whose dimensionality depends on the extent of knowledge we have about the objects. Then, similarity judgments are modeled just as conjunctive probabilities of thinking of the first compared object and then the second (one also needs to assume a relevant mental state, that is neutral between the compared objects), that is, $Sim(A, B) = Prob(A \wedge then B)$ (see also Section 2). For example, Tversky's (1977) famous example of violations of symmetry in similarity judgments was the finding that $Sim(Korea, China) > Sim(China, Korea)$, given that participants have more extensive knowledge for China, than Korea (note, actually Red China and North Korea). In the QP model similarity model, this asymmetry can emerge, as long as the dimensionality for the China subspace is greater than that for the Korea subspace.

The application of the QP decision model onto similarity indicates that the formalism can encompass findings from both decision making and basic similarity. Clearly, such a statement needs to be qualified, since the decision QP model addresses only certain kinds of decision making results and, likewise, the similarity one, only certain kinds of basic similarity results. Nevertheless, Pothos et al.'s theory (2013) is encouraging and fits with an overall prerogative to explain as wide a range of empirical findings as possible, with as few explanatory principles as possible. For example, supporting the same QP model in both decision making and basic similarity makes it plausible that the same principles underlie both kinds of cognitive processes.

While our focus until now has been on basic similarity, a somewhat more recent line of work concerns analogical similarity, which partly concerns the study of analogy formation. Simplifying, analogy formation is about how a naïve observer can establish associations between the elements of two representations (e.g., the atom and the solar system). A key focus for models of analogical

similarity has been the correspondence between the constituent elements of two compared objects; how do they develop and what is their role in the overall similarity judgment, between the compared objects (Gentner, 1983; Goldstone, 1994b; Goldstone & Son, 2005; Larkey & Markman, 2005; Taylor & Hummel, 2009)? For example, suppose we are comparing two persons, Sue and Linda, with black hair. Surely, this fact would contribute more to the overall similarity judgment between Sue and Linda, compared to an alternative situation, where Sue has black hair and Linda has black shoes. That is, there is intuition and supporting evidence that human similarity judgments are sensitive to the structure of the compared objects. An influential idea in modeling structure in similarity judgments is that feature matches can be aligned or not aligned (Goldstone, 1994b; cf. Markman & Gentner, 1993). That is, parts of one object can be placed in correspondence with the parts of another or not. The implication is that matching aligned parts have a greater impact on similarity judgments than matching unaligned parts, but the latter can increase similarity too.

We note that the distinction we make between basic and analogical similarity is partly one of convenience, as it allows us to easily refer to models of similarity not emphasizing structure (e.g., Ashby & Perrin, 1988; Krumhansl, 1978) and ones that do (Gentner, 1983; Goldstone, 1994b). Cognitively, it is possible that there are differences between judgments of basic similarity and analogy formation (e.g., the latter has been claimed to be sometimes analytic, Casale, Roeder, & Ashby, 2012), but these issues do not concern us presently and the distinction between basic and analogical similarity we employ refers to the objectives and scope of corresponding models.

The purpose of this work is to examine whether the QP basic similarity model can be further extended to cover some key requirements for analogical similarity, notably the way similarity computations are affected by correspondences between representation parts (we do not consider the mechanisms that lead to the discovery of which features align or not; this is an important aspect of research in analogical similarity, but beyond the scope of this work). An extension of this sort cannot be expected to perform as well on analogical similarity results, as thoroughbred models of analogical similarity (e.g., Goldstone, 1994b; Hahn, Close, & Graf, 2009; Larkey & Markman, 2005). Nevertheless, attempting the extension is important: if successful, it will show that the mathematical mechanisms for basic similarity judgments are (plausibly and to some extent) the same as the ones for analogical similarity judgments. Equally, if the general model leads to inferior fits in the novel domain, perhaps a restricted scope is more appropriate. This question of possible equivalence of mathematical mechanisms is separate from the one concerning brain systems (cf. Casale et al., 2012). Note also that most basic similarity models cannot account for structure in similarity judgments (Ashby & Perrin, 1988; Krumhansl, 1978; Tversky, 1977; we return to this point later). Likewise, current attempts to extend analogical similarity models to cover Tversky's (1977) results have difficulties. This separatedness of the literatures on basic similarity and analogical similarity further motivate the present effort to develop a QP model of basic/analogical similarity.

As it turns out, there is a very straightforward way to extend the QP model of basic similarity into one of analogical similarity, using Smolensky, Goldrick, and Mathis's (2014) and Smolensky (1990) ideas for modeling structure in linguistic representations. They were interested in the similarity between linguistic representations, where the role in one representation was compared to the role in the other, and likewise for the fillers (e.g., in relation to phonology, a role could be 'syllable-onset' and a corresponding filler could be 'r', for the word 'rat', as in Smolensky et al.'s, 2014, example in p. 5). They derived their similarity method by employing a tensor product representation, which effectively separated

out information of different kinds (roles, fillers), in separate vector spaces. Smolensky et al.'s ideas can be applied to the QP similarity model more or less as they are (since tensor products are standard in QP representations), thus revealing an interesting convergence between QP modeling and computational tools in prior, totally unrelated work. Does the resulting QP model of basic/analogical similarity has any interesting or unique properties? We take up this challenge in relation to possible interference in the contributions to similarity, from different parts of the compared objects. We also provide some indicative fits of the QP model, showing that, with a reasonable parameterization, the performance of the model is acceptable.

2. Quantum probability theory in cognitive modeling and the quantum similarity model

We briefly outline the QP model of basic similarity (for a more complete presentation and the application of the model to Tversky's, 1977 findings see Pothos et al., 2013). The QP basic similarity model is based on a unit length state vector representing mental states, which exists in a Hilbert space. Subspaces of varying dimensionality represent the instances, objects, or concepts, which are to be compared. Each subspace is associated with a linear operator, called a projector. Projectors are fundamental in QP theory. Their function is to take a vector and 'lay it down' onto a corresponding subspace. Suppose that we are interested in the similarity between objects A and B , such that the projectors to the corresponding subspaces are P_A and P_B . The mental state vector just prior to the comparison is denoted as $|\psi\rangle$, in Dirac bra-ket notation, whereby $|x\rangle$ denotes a column vector and $\langle x|$ the complex conjugate of the transpose of $|x\rangle$ (and so a row vector). Then, according to the quantum similarity model, just prior to a comparison, the state vector is assumed to be neutral between the compared objects. For example, when comparing objects A and B , we assume $\|P_B \cdot |\psi\rangle\|^2 = \|P_A \cdot |\psi\rangle\|^2$.¹ This condition $|P_A \cdot |\psi\rangle|^2 = |P_B \cdot |\psi\rangle|^2$ on the state vector is equivalent to that of an uninformed Bayesian prior and basically means that the mental state prior to comparing objects A and B is such that it favors neither A nor B . As noted, $\text{Sim}(A, B) = \text{Prob}(A \wedge \text{then } B)$, which is given by $\text{Sim}(A, B) = |P_B \cdot P_A \cdot |\psi\rangle|^2$; this is the squared length of the projection of the state vector, from the subspace of one compared entity, to the next. The specific interpretation of probability is that the state vector is consistent with the first possibility and then the second (note that, a remarkable theorem by Gleason, 1957, shows that the QP probability rule, for how to go from projections to probabilities, is unique).

As a simple illustration of the quantum similarity model, consider the case where A and B both correspond to unidimensional subspaces in a two-dimensional space. Let $|A\rangle$ and $|B\rangle$ be normalized column vectors along the A and B subspaces respectively. Then, the complex conjugate transpose vectors would be denoted as $\langle A|$ and $\langle B|$. The corresponding projector operators are easily defined as, e.g., $|A\rangle\langle A|$. If $|A\rangle = \begin{bmatrix} a_1 \\ a_2 \end{bmatrix}$, then $P_A = |A\rangle\langle A| = \begin{bmatrix} a_1^* & a_2^* \\ a_2 & a_1 \end{bmatrix}$, which is a 2×2 matrix. Also, $\langle A|A\rangle = \begin{bmatrix} a_1^* & a_2^* \end{bmatrix} \begin{bmatrix} a_1 \\ a_2 \end{bmatrix} = a_1^* \cdot a_1 + a_2^* \cdot a_2$, that is, $\langle A|A\rangle$ is a dot product. With the above definition of similarity we have $\text{Sim}(A, B) = |P_B \cdot P_A \cdot |\psi\rangle|^2 = ||B\rangle\langle B| \cdot |A\rangle\langle A| \cdot |\psi\rangle|^2 = |\langle A|B\rangle|^2 |\langle A|\psi\rangle|^2$ and $\text{Sim}(B, A) = |\langle A|B\rangle|^2 |\langle B|\psi\rangle|^2$, noting that the norm of $|A\rangle$ and

$|B\rangle$ is 1 and that $|\langle A|B\rangle|^2 = |\langle B|A\rangle|^2$. With the assumption that $|\langle B|\psi\rangle|^2 = |\langle A|\psi\rangle|^2$, the relative similarity between A and B is effectively based on $|\langle A|B\rangle|^2$. In a real space, $\langle A|B\rangle = |A| \cdot |B| \cdot \cos \theta$, where θ is the angle between the $|A\rangle$ and $|B\rangle$ vectors (normalization further means that $|A| = |B| = 1$). Thus, overall, $\text{Sim}(A, B) \propto (\cos \theta)^2$, that is, at least for objects represented with one-dimensional subspaces, similarity is just a decreasing function of the angle between the subspaces (up to $\theta = \frac{\pi}{2}$). In brief, this explains how the quantum similarity model works and the basic intuition that drives its formulation.

3. Extending to the QP model of basic similarity to analogical similarity

We introduce the analogical extension to the QP model of basic similarity, using a standard example of analogical similarity, from Goldstone (1994b). In one of his experiments, participants were asked to compare drawings of birds with four distinct parts, left wing, right wing, body, and head (Fig. 1). Each part would include a specific feature (one of 21 random symbols). The features for the four parts can be denoted with four letters, such as ABCD. Thus, a comparison in Goldstone's (1994b) experiment can be denoted as ABCD vs. BADC. In this particular comparison, the two birds share four features, but the shared features do not correspond to the same parts of the birds. Goldstone (1994b) called such matches MOPs, that is, matches out of place. By analogy, MIPs, matches in place, are shared features for the same parts of the bird. Thus, the two representations ABCD and ABXZ would have two MIPs. The distinction between MIPs and MOPs captures a fundamental aspect of structure in representations, namely the fact that particular features could be shared between two objects, without corresponding to the same part of the objects. Goldstone (1994b) found that MIPs have more of an impact on similarity than MOPs, but both MIPs and MOPs increase similarity. He also reported that having both a MIP and a MOP for the same feature (e.g., AACD compared to ABXZ) would actually lead to a lower rating, than just a corresponding MIP (e.g., AYZX). Specifically, in his Experiment 2 (stimuli were schematic birds with symbols), the similarity result for comparing objects with three MIPs was higher than the result with three MIPs and one MOP (Fig. 7, p. 11). Goldstone (1994b) considered this result important enough to highlight the ability of his model, (Similarity as Interactive Activation and Mapping; SIAM) to capture it, against competitors. But, the corresponding difference was not significant and not replicated in Larkey and Markman (2005). Finally, note that there is a related distinction in the literature, between aligned and unaligned differences (e.g., Markman & Gentner, 1996). The consideration of similarity using aligned and unaligned differences is not quite equivalent to that of MIPs and MOPs but, as far as we can see, there are no relevant implications in the discussion below.

Smolensky, Goldrick, and Mathis (2014) solved the problem of introducing structure into geometric representations in the context of combinatorial structure in connectionist models of language processing. Specifically, their starting point (Smolensky et al., 2014) was two structured representations S and S' , which were composed of roles (denoted by r) and fillers for each role (denoted by f). In other words, $S = \{f_j/r_j\}_j$ and $S' = \{f'_k/r'_k\}_k$, whereby $\{f_j/r_j\}_j$ denotes a set of roles and corresponding fillers. For example, S and S' might be two different syllables, r different roles, such as 'onset' or 'coda', and f possible fillers, i.e., phonemes. Then, Smolensky et al. suggested that the similarity between these structured representations should be given by a scheme which computes $\text{sim}(S, S') = \sum_j \sum_k \text{sim}(f_j, f'_k) \cdot \text{sim}(r_j, r'_k)$, that is, a scheme which computes the similarity between the corresponding fillers, for all combinations of roles. Smolensky (1990) and Smolensky et al. (2014) noted that one way to implement this similarity measure is

¹ Note, the double lines indicate the norm of the projected vector and single lines the modulus of a complex number. As in all cases we are dealing with normalized vectors, norms such as the above resolve to computing the modulus of dot products. So, to simplify notation, henceforth we will just write e.g. $|P_A \cdot |\psi\rangle|^2$.




Fig. 1. A figure reproduced from Goldstone (1994b, Fig. 6), illustrating the stimuli employed in his Experiment 2.

by considering the tensor product between roles and fillers and it is this approach which we adopt in the QP model of basic/analogical similarity. Note that other researchers have employed Smolensky's (1990) tensor product idea in analogy formation (e.g., Wiles, Halford, Stewart, Humphreys, & Wilson, 1994; see also Humphreys, Bain, & Pike, 1989), but placed less emphasis on the distinction between MIPs and MOPs. In Wiles et al. (1994) work, tensors are employed to examine analogies between e.g. "mother: baby" and "mare: foal". Humphreys et al.'s (1989) model is a model of memory, whereby tensor products are employed to bind words with their contexts. Also, some researchers have pointed out that a tensor product scheme for combining roles and fillers may involve a combinatorial explosion of overall dimensionality. It is unclear how much a problem this is in practice. Smolensky et al. (2014) argued that it is not. Also, in similarity applications, a problem of overall dimensionality is likely to be less severe than in language.

What is the subset of possible role combinations allowed? This depends entirely on how the vectors for roles are set up. In the most general form for the similarity measure, $\text{sim}(S, S') = \sum_j \sum_k \text{sim}(f_j, f'_k) \cdot \text{sim}(r_j, r'_k)$, it should be clear that all combinations of roles are allowed, so that the similarity measure is consistent with both MIPs and MOPs.

We first illustrate our approach just for MIPs, for simple objects having three distinct parts, top, middle, and bottom (a part is like a role in Smolensky et al.'s example above). Each part can be any of two shapes, a square or a circle. Moreover, any shape can be any of two colors, say blue or red (the shapes and colors are possible fillers, for the roles above). Following Smolensky (1990) and Smolensky et al. (2014), we suggest that the set of all such objects would be represented by vectors of the form $|top\ part\rangle \otimes |shape\rangle \otimes |color\rangle + |middle\ part\rangle \otimes |shape\rangle \otimes |color\rangle + |bottom\ part\rangle \otimes |shape\rangle \otimes |color\rangle$. Note that in our approach the rank of the tensors depends on the number of elementary features, which together characterize an object part (for an alternative approach see Wiles et al., 1994). The role vectors indicate the part under consideration and, since we are assuming only MIPs, we can set these to a canonical basis to write

$$\begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} \otimes |shape\rangle \otimes |color\rangle + \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix} \otimes |shape\rangle \otimes |color\rangle + \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} \otimes |shape\rangle \otimes |color\rangle. \quad (1)$$

The tensor product of the part vector, with the vector corresponding to shape, and the vector corresponding to color, allows all combinations of shapes and colors for each part.

In the QP basic similarity model, the similarity between two objects, A and B , is $\text{Sim}(A, B) = |P_B \cdot P_A \cdot |\psi\rangle|^2$. In the extended

QP model of basic/analogical similarity, the definition of similarity is identical. The similarity computation is straightforward to carry out. Note first that we can write $\text{Sim}(A, B) = |\langle A | \langle A | B \rangle \langle B | \psi \rangle|^2$. The first dot product, $\langle B | \psi \rangle$, can be ignored, since we require the state vector, $|\psi\rangle$, to be neutral with respect to the A and B rays (i.e., $|\langle B | \psi \rangle|^2 = |\langle A | \psi \rangle|^2$). The $\langle A |$ term can be ignored as well, since $|\langle A | \psi \rangle|^2 = 1$. Therefore, we can write $\text{Sim}(A, B) = |\langle A | B \rangle|^2$. The dot product is

$$\begin{aligned} \langle A | B \rangle = & \left\langle \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}_A \otimes (shape)_A \otimes (color)_A \right. \\ & + \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix}_A \otimes (shape)_A \otimes (color)_A \\ & + \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}_A \otimes (shape)_A \otimes (color)_A \\ & \times \left| \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}_B \otimes (shape)_B \otimes (color)_B \right. \\ & + \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix}_B \otimes (shape)_B \otimes (color)_B \\ & \left. + \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}_B \otimes (shape)_B \otimes (color)_B \right\rangle. \quad (2) \end{aligned}$$

This leads to

$$\begin{aligned} & \langle shape_A | shape_B \rangle_{part\ 1} \langle color_A | color_B \rangle_{part\ 1} \\ & + \langle shape_A | shape_B \rangle_{part\ 2} \langle color_A | color_B \rangle_{part\ 2} \\ & + \langle shape_A | shape_B \rangle_{part\ 3} \langle color_A | color_B \rangle_{part\ 3}. \quad (3) \end{aligned}$$

The above computation is simpler than perhaps how it looks. All terms combine with each other and each individual dot product is of the form $\langle x \otimes y \otimes z | x' \otimes y' \otimes z' \rangle = \langle x | x' \rangle \langle y | y' \rangle \langle z | z' \rangle$. In other words, the dot product between two tensor products decomposes into dot products for each 'slot' (space) separately. The first of these dot products involves the vector for the particular parts under consideration in the two objects. But, these part vectors have been set up in such a way that, unless they are the same, their dot product is zero. For example, consider the dot product involving the first part of object A and the second part of object B . This is

$$\begin{aligned} & \left\langle \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}_A \otimes (shape)_A \otimes (color)_A \right| \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix}_B \otimes (shape)_B \otimes (color)_B \rangle \\ & = (1\ 0\ 0)_A \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix}_B \dots = 0. \quad (4) \end{aligned}$$

So, overall, this particular way of setting up the part vectors readily allows for similarity comparisons to be faithful to structure, in the sense that shape equivalence adds to the similarity computation, only when the corresponding parts are the same between the two objects. In brief, this is a scheme for MIPs.

The QP model of basic/analogical similarity can easily allow for contributions to similarity from MOPs as well, in a way identical to that of Smolensky et al. (2014), by allowing non-orthogonal part vectors (i.e., non-orthogonal roles). Note that, in the QP model, MOPs can only arise from similarities between roles; if roles have no similarity structure, then the QP model must predict no MOPs. We are also now led to a key distinction between the QP model of basic/analogical similarity and Smolensky et al.'s (2014) approach.

The native space of QP representations is complex. In QP theory, complex numbers are employed as a technical convenience and all outputs of the theory are guaranteed to be real. But, in part because of the use of complex vectors, QP theory computations sometimes give rise to interference effects, not possible with real vectors. Therefore, with non-orthogonal, complex vectors for roles, we are led to the possibility of not only MOPs, but of interference effects between MOPs and MIPs.

To illustrate, consider a simplified example of two objects created of two possible features in two possible roles, so that $A = r_1 \otimes f_1 + r_2 \otimes f_2$ and $B = r_1 \otimes f'_1 + r_2 \otimes f'_2$, noting that r_1 is the role vector for the first part and r_2 the role vector for the second part. The actual features f_1, f_2 etc. could be the same or not, this is not relevant presently. Then, $\text{Sim}(A, B) \propto |\langle A|B \rangle|^2$. The inner product is computed as $\langle A|B \rangle = \langle r_1 \otimes f_1 + r_2 \otimes f_2 | r_1 \otimes f'_1 + r_2 \otimes f'_2 \rangle = \langle f_1 | f'_1 \rangle + \langle f_2 | f'_2 \rangle + \langle r_1 | r_2 \rangle \langle f_1 | f'_2 \rangle + \langle r_2 | r_1 \rangle \langle f_2 | f'_1 \rangle$. The first two terms are contributions from MIPs; the cross-terms are MOPs. Here, we are only concerned with the cross-terms, since the MIPs' contribution cannot change by altering the role vectors. We can rewrite the cross-terms as: $\text{cross} = \langle r_1 | r_2 \rangle G + \langle r_2 | r_1 \rangle H$, where $G = \langle f_1 | f'_2 \rangle$ and $H = \langle f_2 | f'_1 \rangle$. Depending on the form of r_1, r_2 , there is potential for the cross-terms to reduce the overall value of $\langle A|B \rangle$, when the contributions to similarity from MOPs can be negative. This is simply illustrated in two dimensions. Letting x, y, a, b , be all positive, real numbers, consider four possibilities:

$$r_1 = \begin{pmatrix} x \\ y \end{pmatrix}, \quad r_2 = \begin{pmatrix} a \\ b \end{pmatrix}, \quad \text{then } \langle r_1 | r_2 \rangle G + \langle r_2 | r_1 \rangle H = (xa + yb)(G + H) \quad (5a)$$

$$r_1 = \begin{pmatrix} x \\ iy \end{pmatrix}, \quad r_2 = \begin{pmatrix} a \\ b \end{pmatrix}, \quad \text{then } \langle r_1 | r_2 \rangle G + \langle r_2 | r_1 \rangle H = xa(G + H) + iby(H - G) \quad (5b)$$

$$r_1 = \begin{pmatrix} x \\ iy \end{pmatrix}, \quad r_2 = \begin{pmatrix} ia \\ b \end{pmatrix}, \quad \text{then } \langle r_1 | r_2 \rangle G + \langle r_2 | r_1 \rangle H = ixa(G - H) + iby(H - G) \quad (5c)$$

$$r_1 = \begin{pmatrix} ix \\ iy \end{pmatrix}, \quad r_2 = \begin{pmatrix} -ia \\ -ib \end{pmatrix}, \quad \text{then } \langle r_1 | r_2 \rangle G + \langle r_2 | r_1 \rangle H = -(xa + yb)(G + H). \quad (5d)$$

In the first case, we have a baseline, additive, positive contribution from MOPs to similarity. In cases (5b) and (5c) the contribution from MOPs is lower, relative to that in the baseline case (5a). Finally, in case (5d) the contribution to similarity from MOPs is negative, meaning that the similarity we would expect from MIPs would be reduced, because of MOPs. Note that all the modifications to the role vectors do not alter their norms (i.e., the computations for MIPs are unaffected).

Thus, with complex role vectors, the QP model of basic/analogical similarity predicts a possibility that MOPs will interfere with MIPs, which is a novel prediction from the QP model of basic/analogical similarity. Goldstone's (1994b) SIAM model can also predict such an interference. The QP model and SIAM are not necessarily competing models, as they emphasize different aspects of psychological explanation (the QP model is concerned with the guiding mathematical principles, SIAM is a network model aimed at capturing underlying processes). As noted, Goldstone (1994b) indeed reported a nonsignificant trend where MOPs reduce the overall similarity between two objects. Exploring the reality of such interference effects and identifying the circumstances which make them likely is an exciting (and challenging) direction for further work. For example, perhaps more complex stimuli, for which the different roles have complex relationships with each other (such as those employed in Goldstone, 1994b), make interference effects more likely, as opposed to stimuli with simpler structure

and more schematic form (e.g., Larkey & Markman, 2005). Note, Goldstone's (1994b) own view was that such interference effects arise when the same feature is involved in both a MIP and a MOP.

4. Further exploring the QP model of basic/analogical similarity

4.1. Analytical examination

In Section 3 we provided an example for how the QP model of basic/analogical similarity can account for MIPs and MOPs (and possible interference effects between the two). Here, first, we consider the formal properties of the model more systematically, with an emphasis on three main findings from Goldstone (1994b): (1) MIPs have more of an impact on similarity than MOPs, (2) MIPs increase similarity, and (3) MOPs increase similarity. Second, having analytically examined the formal sensitivity of the model to these aspects of structural similarity, we provide a computational demonstration of its behavior.

We use the terminology of Smolensky et al. (2014), to make it easier to appreciate the links between our work and theirs. Consider objects composed of n roles (these would be analogous to stimulus components or parts, in the language of basic similarity models) and a set of fillers for each role (a filler is a feature, that is, the way in which a particular stimulus component can appear). Let $|r_1\rangle, |r_2\rangle, \dots, |r_n\rangle$ be n -dimensional positive real vectors with unit length representing the roles. Let $|f_i\rangle$ be a real unit length vector representing the filler associated with role i . A standard simplifying assumption in similarity research is that fillers are orthonormal, so that there is no similarity between the different e.g. features which make up a schematic stimulus. Although we assume filler orthonormality in this section, the QP model does not require it; we discuss the implications of violating this assumption and, also, do not require it in the computational demonstration (note, Goldstone's, 1994b, model incorporates violations of filler orthonormality in terms of a global parameter). Note that, unless stated otherwise, role orthonormality is *not* assumed (MOPs arise when roles are not orthonormal).

In the quantum similarity model, an object S is represented by the tensor product $r_1 \otimes f_1 + r_2 \otimes f_2 + \dots + r_n \otimes f_n$. The similarity between two objects S and S' can be written as the following, up to a normalizing constant:

$$\begin{aligned} \text{Sim}(S, S') &= |\langle S | S' \rangle|^2 \\ &= |\langle r_1 \otimes f_1 + r_2 \otimes f_2 + \dots + r_n \otimes f_n | r_1 \otimes f'_1 \\ &\quad + r_2 \otimes f'_2 + \dots + r_n \otimes f'_n \rangle|^2 \\ &= \left| \sum_{i=1}^n \sum_{j=1}^n \langle r_i | r_j \rangle \langle f_i | f'_j \rangle \right|^2. \end{aligned} \quad (6)$$

A MIP between S and S' implies that there is some role i , such that the fillers for the role match, $\langle f_i | f'_i \rangle = 1$. A MOP between S and S' implies that there are roles i and j , which are not orthonormal (i.e., $0 < |\langle r_i | r_j \rangle| \leq 1$; note that $\langle r_i | r_j \rangle$ could be negative, hence the absolute value is necessary), such that the fillers match, $\langle f_i | f'_j \rangle = 1$. In other words, the existence of a MOP implies the conjunction of two things: $|\langle r_i | r_j \rangle| > 0$ and $\langle f_i | f'_j \rangle = 1$. Also, notice that $|\langle r_i | r_j \rangle| \leq |\langle r_i | r_i \rangle| = 1$, because of the assumption that role vectors have unit length. That is, matching roles will have more impact on similarity judgments, than non-matching roles.

Our first claim is that MIPs have more of an impact on similarity than MOPs. Let S, S' , and S'' have n roles. Let S and S' have $k \leq n$ MIPs (and no MOPs) and let S and S'' have $k \leq n$ MOPs (and no MIPs). If different fillers are orthonormal, that is, $\langle f_i | f_i \rangle = 1$ and $\langle f_i | f'_i \rangle = 0$ for $f_i \neq f'_i$, then $\text{Sim}(S, S') \geq \text{Sim}(S, S'')$. The proof is straightforward. Since there are no MOPs in the comparison of S

and S' , all of the fillers for mismatching roles are different. Then, $\langle f_i | f'_j \rangle = 0$ for all $i \neq j$ because of filler orthonormality. The presence of MIPs implies that there are exactly $k \leq n$ nonzero terms of the form $\langle r_i | r_i \rangle \langle f_i | f'_i \rangle$ where $\langle r_i | r_i \rangle = 1$ and $\langle f_i | f'_i \rangle = 1$. Since there are no MIPs in the comparison of S and S'' , all of the fillers for matching roles are different. That is, $\langle f_i | f''_i \rangle = 0$ for all i , because of filler orthonormality. The presence of n MOPs implies that there are exactly $k \leq n$ nonzero terms of the form $\langle r_i | r_j \rangle \langle f_i | f''_j \rangle$ where $0 < |\langle r_i | r_j \rangle| \leq 1$ and $\langle f_i | f''_j \rangle = 1$. Because $|\langle r_i | r_i \rangle| \geq |\langle r_i | r_j \rangle|$ for all i and j , $\text{Sim}(S, S') \geq \text{Sim}(S, S'')$.

Without filler orthonormality, we could have situations where two objects have a high similarity without either MIPs or MOPs. Suppose $\langle r_i | r_j \rangle \neq 0$ (no MIPs) and we do not have the same feature in the roles i, j (no MOPs). Then, without filler orthonormality, there could be instances when $\langle f_i | f'_j \rangle > 0$. Depending on the similarity structure of the features, a large number of (relatively high value) cross-terms of this form could exceed the contribution of a small (or zero) number of actual MIPs or MOPs, in a different comparison. This partly illustrates why filler orthonormality is a useful assumption in practical research, since ignoring cross-terms (i.e., ignoring contributions to similarity which are neither MIPs nor MOPs and, instead, arise from similarity between non-identical features) makes similarity computations and modeling more tractable.

The second claim concerns the impact of MIPs on similarity judgments. We show that increasing the number of MIPs increases similarity. Specifically, suppose S and S' have more MIPs than S and S'' and there are no MOPs. Then, if different fillers are orthonormal, that is, $\langle f_i | f_i \rangle = 1$ and $\langle f_i | f'_i \rangle = 0$ for $f_i \neq f'_i$, it will be the case that $\text{Sim}(S, S') > \text{Sim}(S, S'')$. To prove this, note first that, since there are no MOPs, all of the fillers for mismatching roles are different. That is, $\langle f_i | f'_j \rangle = 0$ and $\langle f_i | f''_j \rangle = 0$ for all $i \neq j$ because of filler orthonormality. The remaining terms correspond to matching roles and are of the form $\langle r_i | r_i \rangle \langle f_i | f'_i \rangle$ for $\text{Sim}(S, S')$ and $\langle r_i | r_i \rangle \langle f_i | f''_i \rangle$ for $\text{Sim}(S, S'')$, where $\langle r_i | r_i \rangle = 1$ for all i . A MIP implies the same filler in matching roles. Because we assume that the fillers are orthonormal, $\langle f_i | f'_i \rangle = 1$ if $f_i = f'_i$ and zero otherwise, likewise for $\langle f_i | f''_i \rangle$. Because S and S' have more MIPs than S and S'' , there exist more fillers f'_i such that $f_i = f'_i$ than fillers f''_i such that $f_i = f''_i$. Thus, there are more nonzero terms equal to 1 in the calculation of $\text{Sim}(S, S')$ than $\text{Sim}(S, S'')$.

Eliminating the assumption of filler orthonormality potentially confounds this conclusion. Suppose that there are fewer MIPs in comparing S with S'' , than in comparing S with S' , and no MOPs. In comparing S with S'' , it could be the case that there is a large contribution from feature pairs where the roles are identical, but $\langle f_i | f'_i \rangle \neq 1$ (i.e., no MIP) and $\langle f_i | f'_i \rangle > 0$.

By analogy with the second claim, the third claim is that MOPs increase similarity. That is, suppose S and S'' have a subset of the MOPs of S and S' and there are no MIPs (e.g., ABCD vs. BCXZ compared to ABCD vs. BCDZ). If different fillers are orthonormal, then $\text{Sim}(S, S') > \text{Sim}(S, S'')$. For the proof, consider that, if there are no MIPs, then all of the fillers for matching roles are different. That is, $\langle f_i | f'_i \rangle = 0$ and $\langle f_i | f''_i \rangle = 0$ for all i because of filler orthonormality. The remaining terms corresponding to mismatching roles are of the form $\langle r_i | r_j \rangle \langle f_i | f'_j \rangle$ for $\text{Sim}(S, S')$ and $\langle r_i | r_j \rangle \langle f_i | f''_j \rangle$ for $\text{Sim}(S, S'')$ where $\langle r_i | r_j \rangle > 0$. A MOP implies the same filler in different roles. Because we assume that the fillers are orthonormal, $\langle f_i | f'_j \rangle = 1$ if $f_i = f'_j$ and 0 otherwise, likewise for $\langle f_i | f''_j \rangle$. Because S and S'' have a subset of the MOPs of S and S' , there exist some i and j such that $f_i = f'_j$ and $f_i \neq f''_j$. (Note that the MOPs of S and S'' involve the same role vectors as the MOPs of S and S' , due to the fact that the MOPs of S and S'' are a subset of those of S and S' .) Therefore, there are more nonzero terms of the form $\langle r_i | r_j \rangle \langle f_i | f'_j \rangle$ for $\text{Sim}(S, S')$ than terms of the form $\langle r_i | r_j \rangle \langle f_i | f''_j \rangle$ for $\text{Sim}(S, S'')$.

The implications for violating filler orthonormality are as above. Here, additionally, we have to take into account the similarity structure between roles, between the compared stimuli. All the MOPs concern terms for which $\langle f_i | f'_j \rangle = 1$. In addition, for a MOP, we have $\langle r_i | r_j \rangle > 0$. But the exact magnitude of $\langle r_i | r_j \rangle$ will depend on role similarity and this may vary from comparison to comparison. It could be the case that a smaller number of MOPs could contribute more to similarity than a greater number of MOPs, if the role similarities for the former are suitably higher. This is why the 'subset' assumption is needed to prove this third claim. For simple, schematic stimuli, it is reasonable to assume that role similarity is flat, that is, all terms of the form $\langle r_i | r_j \rangle$, which are nonzero, have the same magnitude. In such a case, the 'subset' assumption is not needed and this claim can be proved just on the basis of filler orthonormality.

4.2. Computational examination

The analytical demonstration shows that broad level predictions from the QP model are sensible. But how well does the QP model perform against relevant models? We considered the results from Hodgetts et al.'s (2009) Experiment 1. Hodgetts, Hahn, and Chater (2009) built on previous empirical studies on analogical similarity (e.g., Larkey & Markman, 2005) and, moreover, they reported results from a range of relevant models. Their Experiment 1 is the simplest one in their paper and so suitable for this first computational examination of the QP model of basic/analogical similarity.

The results from Hodgetts et al.'s (2009) Experiment 1 consisted of 13 object pairs, with information on the proportion of times each object pair was selected, amongst alternative object pairs, in a forced choice task. Selection was based on the perceived similarity between the objects making up an object pair. Each object was composed of two features (out of three available, denoted A, B, and C), each placed in one of two positions (i.e., two roles). Thus, an object pair could be represented as AA/AA, indicating that the same features appeared in both roles, for both objects. Unsurprisingly, this object pair was selected most frequently (91.56% of all times). By contrast, object pair AB/BB was selected only about half the time. The least selected object pair was AB/CC (23.57%). Hodgetts et al. (2009) manipulated the physical form of the three features, so that, depending on the condition, A could be any of a circle, a square, or a triangle.

To motivate the parametric choices for the QP model, consider a simple example of computing the similarity between objects, such that $\text{Object}_1 = r_1 \otimes f_1 + r_2 \otimes f_2$ and $\text{Object}_2 = r_1 f_3 + r_2 f_4$ (for the second object and henceforth we omit the tensor product symbol). Then, $\text{Sim}(\text{object 1}, \text{object 2}) = |\langle r_1 f_1 + r_2 f_2 | r_1 f_3 + r_2 f_4 \rangle|^2 = |\langle f_1 | f_3 \rangle + \langle f_2 | f_4 \rangle + \langle r_1 | r_2 \rangle \langle f_1 | f_4 \rangle + \langle r_2 | r_1 \rangle \langle f_2 | f_3 \rangle|^2$, whereby the first two terms are contributions from MIPs and the latter two from MOPs. Casually inspecting Hodgetts et al.'s (2009) Experiment 1 data, it is clear that there are contributions from MOPs (e.g., the object pair AB/BA was chosen quite frequently; without MOPs, the two objects in the pair would have no similarity). Therefore, we have to allow for non-orthogonal role vectors. Also, the data indicate that the assumption of feature orthonormality is not tenable as well (e.g., AB/AA was not chosen as equally often as AB/BB). Indeed, intuitively, it is unlikely that participants considered a square, circle, and triangle as maximally dissimilar and equally similar to each other.

Based on these observations, we chose to represent roles as non-orthogonal vectors in a two dimensional real space and features as non-orthogonal vectors in a separate two dimensional real space. The restriction to real spaces implies no possibility of interference between the contribution to similarity from MIPs and MOPs, which is reasonable, given our preliminary intuitions




Fig. 2. Comparison of the observed choice data from Experiment 1 (Hodgetts et al., 2009) and the best fit of the QP model of basic/analogical similarity, with a regression line.

for when to expect such effects (Section 3). Note, employing a two dimensional real space for three features implies that it is impossible for the three features to be orthonormal, but limits the overall number of parameters. Overall, the QP model had three parameters, one to characterize feature A , one for feature B , and one for the first role. Since all vectors are normalized and two-dimensional, a single parameter is needed to characterize any vector. The vector for feature C was fixed to $\frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ 1 \end{bmatrix}$, since what matters is the relative similarity structure of the three feature vectors and we allowed A and B to vary freely; likewise, the vector for the second role was fixed to $\begin{bmatrix} 0 \\ 1 \end{bmatrix}$. We allowed the free parameters to vary, so as to optimize the correlation between the similarities between object pairs, computed by the QP model, and the empirically observed choice proportion for each pair.

The best identified correlation was 0.81 and the results are illustrated in Fig. 2 (the optimization was carried in Matlab and the code is available from the authors). The ‘representational distortion’ model, favored by Hodgetts et al. (2009), achieved a correlation of -0.95 . Hodgetts et al. (2009) also report that the correlation for the Structure Mapping Engine model (SME; Falkenhainer, Forbus, & Gentner, 1989) was 0.74, the correlation for the SIAM model (Goldstone, 1994b) 0.76, and the correlation which can be achieved by the best linear weighting of MIPs and MOPs 0.77 (this value is inferred from Fig. 9 in Hodgetts et al., 2009).

We note the limitations of the QP model. First, achieving the correlation of 0.81 relied on post hoc fitting of three free parameters. By contrast, the representational distortion model had no free parameters (the predictions are based on the specification of a coding scheme), regarding SIAM (Hodgetts et al., 2009) employed a recycled set of parameter values from previous work, and SME’s application involved no parameters anyway. Thus, the QP model failed to achieve the best fit, even though it was aided by the greatest number of parameters. We can only consider the performance of the QP model as promising, given of course that this is the very first application of the QP model in this domain.

5. A brief consideration of alternative approaches

The similarity literature is vast and it is beyond our objective to summarize it in detail (for a review, see Goldstone & Son, 2005).

We consider a limited range of basic similarity or analogy models, with a view to illustrate some of the points made throughout this paper.

One important motivation for the present paper has been our argument that basic similarity models are ill equipped to handle analogical similarity. This can be illustrated fairly simply. Distance-based models of similarity and feature-based models (such as Tversky’s, 1977, model) could, in principle, account for some analogical aspects of similarity (e.g., through the introduction of conjunctive features or dimensions). In simple distance-based models of similarity, the similarity between two entities A and B is given by $\text{Sim}(A, B) = e^{-c \cdot \text{distance}(A, B)}$, where c is a constant. In Tversky’s (1977) proposal, $\text{Similarity}(A, B) = \theta f(A \cap B) - \alpha f(A - B) - \beta f(B - A)$, where θ , α , β are parameters, $A \cap B$ denotes the common features between A and B , $A - B$ the features of A which B does not have and $B - A$ the features of B which A does not have. In both cases structure could be modeled by conjunctive dimensions or features (e.g., dimension 1 is shape for feature 1 etc.; cf. Goldstone & Son, 2005), but such a solution is not elegant and appears demanding, in terms of representational resources. Moreover, a distance-based model of similarity could be augmented with a tensor product approach, analogous to Smolensky et al.’s (2014) method. However, without further refinements, such a distance-based model would still have difficulty with Tversky’s (1977) empirical findings.

Two important approaches in basic similarity are Krumhansl’s (1978, 1988) distance-density model and Ashby and Perrin’s (1988) general recognition theory. The former model assumes that alternatives lying within dense subregions of psychological space are subject to finer discrimination than alternatives lying in less dense subregions. This implies that the distance (dissimilarity) between two points A and B in psychological space should be affected by the local density around each point, $D(A)$ and $D(B)$. Thus, $d'(A, B) = d(A, B) + aD(A) + bD(B)$, where $d(A, B)$ is the standard geometric distance, a and b are parameters, and $d'(A, B)$ is the modified distance measure. These ideas are interesting, but it is unclear how to generalize them in the case of structured representations. Ashby and Perrin’s (1988) general recognition theory is a probabilistic approach to similarity for perceptual stimuli. Each stimulus (e.g., presented in different trials) can correspond to different points in psychological space, according to a particular probability distribution. Psychological space is divided into response regions, such that within each response region it is optimal to make a particular response. Thus, similarity between two stimuli depends on the overlap between the distribution of perceptual effects for the first stimulus and the optimal response region for the second stimulus. As with Krumhansl’s (1978) model, it is unclear how to extend Ashby and Perrin’s (1988) approach to take into account the structure of particular stimuli.

We now turn to the converse point: can models of analogical similarity cover Tversky’s (1977) key results? Some models of analogical similarity do not consider these results in a detailed way (e.g., Gentner, 1983; Larkey & Love, 2003 and Markman & Gentner, 1996). An interesting proposal is Taylor and Hummel’s (2009) LISA model, which involves a representation scheme analogous to that of Smolensky (1990), but for the crucial difference that roles and fillers are combined additively, rather than multiplicatively. This means that the impact of role similarity is independent of the impact of filler similarity, in establishing an analogy. Regarding computations of similarity, the similarity between two objects A and B reflects a term from MIPs and a term from MOPs. Consider trying to map the structure of a source, S , to a target, T . Then, the similarity between the two is given by $\sigma_{\text{total}}(T, S) = \sigma_{\text{MIP}}(T, S) + \sigma_{\text{MOP}}(T, S)$, where $\sigma_{\text{MIP}}(T, S)$ and $\sigma_{\text{MOP}}(T, S)$ is a contribution from MIPs and MOPs, respectively. The first term provides a measure of the extent to which units in S uniquely map onto units in T . The second term, $\sigma_{\text{MOP}}(T, S)$, is a measure of similarity across all objects and roles

in T and S , with no regard of whether they map or not. How well does LISA fair with Tversky's (1977) main findings? A key result is that $Sim(Korea, China) > Sim(China, Korea)$, originally explained as arising, in part, because participants' representation of China would involve more features, than that of Korea. The MIPs term in LISA can be shown to have the form $\frac{\text{features of } A \text{ uniquely matched to } B}{\text{all features of } A}$ and it can lead to similarity asymmetries, when there are differences in the number of features for A and B . Thus, the LISA model can explain the Korea, China asymmetry. Another one of Tversky's (1977) key results, a violation of the triangle inequality, involved similarities amongst three countries, Cuba, Jamaica, Russia. Here, for LISA to accommodate the result in the observed direction, it had to be assumed that Russia was represented in terms of government, Cuba in terms of government and climate, and Jamaica in terms of climate, so introducing a somewhat arbitrary difference in the extent of representations. Notwithstanding the fact that, as far as we can judge, LISA is a very competent model of analogy, we think that its coverage of Tversky's (1977) main results is not fully satisfactory. The specification of LISA involves many architectural choices, which are not always possible to fully constrain. In this way we illustrate what we think is a general feature of models of analogical similarity: the emphasis is on the computational procedures which can lead to the identification of MIPs and MOPs (something currently unspecified in the QP model), not on the algebraic properties of similarity judgments; it is the latter that bears on whether satisfactory coverage with Tversky's (1977) results can be achieved or not.

Finally, another model of analogical similarity, the representational distortion model (Hahn, Chater, & Richardson, 2003; Hodgetts et al., 2009) could account for violations of symmetry in similarity judgments if, for example, removing elements from a representation can be assumed to be easier (carry a lower information cost) than adding elements (Hahn et al., 2009). Thus, when comparing an informationally simpler object with a more complex one, we should obtain $Sim(Complex, Simple) > Sim(Simple, Complex)$. The model also appears consistent with violations of the triangle inequality, if, for example, one assumes that in e.g. the case of Russia and Cuba the common political system facilitates the 'transformation' of one object to the other. Thus, the representational distortion model potentially encompasses both basic similarity and analogical similarity results (see also, Hodgetts & Hahn, 2012).

Where we note shortcomings of the above models, we do not imply criticism, as we examined their application in non-native empirical domains. We simply wish to argue that developing a model encompassing both basic similarity and analogical similarity, as we have attempted to do with the QP model, is a non-trivial task.

6. Concluding comments

It is uncontroversial to state that it is desirable to extend specific cognitive models as much as possible. Such extensions inform our understanding of commonalities regarding the underlying cognitive processes, across perhaps seemingly disparate aspects of behavior. They also enable us to understand boundaries in the applicability of particular formalisms and hence the motivation for domain-specific mechanisms or principles. It was in this spirit that we set out to explore whether Pothos et al.'s (2013) QP model of basic similarity could be extended to analogical similarity. We took this to be an interesting challenge, especially because geometric models of basic similarity (such as the QP model) have been thought to be incapable of accommodating structure in similarity judgments (cf. Goldstone & Son, 2005). As things stand, the QP model of basic/analogical similarity we reported can cover certain aspects of decision making (as in Bussemeyer et al., 2011),

basic similarity (Pothos et al., 2013), and analogical similarity (the present paper).

We think that the QP model of basic/analogical similarity is, at the very least, a promising candidate for further consideration in the analogical similarity literature. Its formal properties are consistent with expectation, e.g., in relation to the relative impact of MIPs and MOPs. The computational analysis we carried out, based on the results from Experiment 1 of Hodgetts et al. (2009), revealed a decent performance from the QP model, in relation to some standard models of analogical similarity. This was an encouraging preliminary result and recommends further similar computational examinations, e.g., against alternative datasets. Moreover, we identified an interesting novel prediction from the QP model. In standard QP theory, representations are defined over complex, not real, vector spaces. In analogical similarity, employing complex vectors for roles can lead to interference effects between the contributions for MIPs and MOPs. In other words, the QP model predicts (fairly naturally) that there will be some cases when MOPs will conflict with MIPs and so reduce overall similarity (Goldstone, 1994b states such an implication from his SIAM model too). Developing this prediction in a more specific way is an ambitious objective, though it may be possible to borrow insights from other QP modeling work, for when interference is expected (e.g., Bussemeyer & Bruza, 2012).

In this work we focused on the key challenge for models of analogical similarity of representing structure in a way that it is consistent with MIPs and MOPs, in a certain way. Another key challenge is discovering the appropriate mappings between two objects in the first place (e.g., Goldstone, 1994b and Taylor & Hummel, 2009). This is an aspect of the QP model of basic/analogical similarity that still needs to be developed. Perhaps some mechanism based on alignment of features (which could be easily implemented in the QP model) may provide a suitable way forward. Somewhat relatedly, we have not considered possible cognitive limitations as the rank of relevant tensors increases (which would happen if the compared objects increase in complexity; cf. Simon, 1955). This has been noted in an alternative approach to analogical similarity (and memory) based on Smolensky et al.'s (2014) ideas (Wiles et al., 1994). An interesting difference between the quantum tensor product approach and traditional ones is that, in the former, the overall required dimensionality can be moderated through the use of incompatible representations. For example, with three binary features, classically we need 2^3 dimensions. In a quantum model, if the features are incompatible, the needed dimensionality could be as low as two (Pothos & Bussemeyer, 2013; Trueblood & Bussemeyer, 2011; cf. Emruli, Gayler, & Sandin, 2013). Exploring more carefully the specific computational resources implied by different modeling approaches is an important priority for further work.

In sum, we think that the QP model of basic/analogical similarity is a promising model of analogical similarity. This encourages us in thinking that the QP principles may provide a firm foundation for understanding aspects of cognition including decision making, basic similarity, and analogical similarity.

Acknowledgments

We would like to thank Jerome Bussemeyer and James Hampton for their helpful comments. EMP was supported by Leverhulme Trust grant RPG-2013-004 and by Air Force Office of Scientific Research (AFOSR), Air Force Material Command, USAF, grant FA 8655-13-1-3044. JST was supported by NSF grant SES-1326275. The US Government is authorized to reproduce and distribute reprints for Governmental purpose notwithstanding any copyright notation thereon.

References

- Aerts, D. (2009). Quantum structure in cognition. *Journal of Mathematical Psychology*, 53, 314–348.
- Aerts, D., & Gabora, L. (2005). A theory of concepts and their combinations II: a Hilbert space representation. *Kybernetes*, 34, 192–221.
- Ashby, G. F., & Perrin, N. A. (1988). Towards a unified theory of similarity and recognition. *Psychological Review*, 95, 124–150.
- Blutner, R., Pothos, E. M., & Bruza, P. (2013). A quantum probability perspective on borderline vagueness. *Topics in Cognitive Science*, 5, 1–26.
- Busemeyer, J. R., & Bruza, P. (2012). *Quantum models of cognition and decision making*. Cambridge, UK: Cambridge University Press.
- Busemeyer, J. R., Pothos, E., Franco, R., & Trueblood, J. S. (2011). A quantum theoretical explanation for probability judgment 'errors'. *Psychological Review*, 118(2), 193–218.
- Casale, M. B., Roeder, J. L., & Ashby, F. B. (2012). Analogical transfer in perceptual categorization. *Memory and Cognition*, 40, 434–449.
- Emruli, B., Gayler, R. W., & Sandin, F. (2013). Analogical mapping and inference with binary spatter codes and sparse distributed memory. In *Proceedings of international joint conference on neural networks (IJCNN)* (pp. 1–88). Dallas, Texas: IEEE Press.
- Falkenhainer, B., Forbus, K., & Gentner, D. (1989). The structure-mapping engine: algorithm and examples. *Artificial Intelligence*, 41, 1–63.
- Gentner, D. (1983). Structure-mapping: a theoretical framework for analogy. *Cognitive Science*, 7, 155–170.
- Gleason, A. (1957). Measures on the closed subspaces of a Hilbert space. *Journal of Mathematics and Mechanics*, 6, 885–893.
- Goldstone, R. L. (1994a). The role of similarity in categorization: providing a groundwork. *Cognition*, 52, 125–157.
- Goldstone, R. L. (1994b). Similarity, interactive activation, and mapping. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 20, 3–28.
- Goldstone, R. L., & Son, J. (2005). Similarity. In K. Holyoak, & R. Morrison (Eds.), *Cambridge handbook of thinking and reasoning* (pp. 13–36). Cambridge: Cambridge University Press.
- Hahn, U., Chater, N., & Richardson, L. B. (2003). Similarity as transformation. *Cognition*, 87, 1–32.
- Hahn, U., Close, J., & Graf, M. (2009). Transformation direction influences shape-similarity judgments. *Psychological Science*, 20, 447–454.
- Haven, E., & Khrennikov, A. (2013). *Quantum social science*. Cambridge, UK: Cambridge University Press.
- Hodgetts, C. J., & Hahn, U. (2012). Similarity-based asymmetries in perceptual matching. *Acta Psychologica*, 139, 1–9.
- Hodgetts, C. J., Hahn, U., & Chater, N. (2009). Transformation and alignment in similarity. *Cognition*, 113, 62–79.
- Hughes, R. I. G. (1989). *The structure and interpretation of quantum mechanics*. Cambridge, MA: Harvard University Press.
- Humphreys, M. S., Bain, J. D., & Pike, R. (1989). Different ways to cue a coherent memory system: a theory of episodic, semantic and procedural tasks. *Psychological Review*, 96, 208–233.
- Isham, C. J. (1989). *Lectures on quantum theory*. Singapore: World Scientific.
- Khrennikov, A. Y. (2010). *Ubiquitous quantum structure: from psychology to finance*. Springer.
- Krumhansl, C. L. (1978). Concerning the applicability of geometric models to similarity data: the interrelationship between similarity and spatial density. *Psychological Review*, 85, 445–463.
- Krumhansl, C. L. (1988). Testing the density hypothesis: comment on Corter. *Journal of Experimental Psychology: General*, 117, 101–104.
- Larkey, L. B., & Love, B. C. (2003). CAB: connectionist analogy builder. *Cognitive Science*, 27, 781–794.
- Larkey, L. B., & Markman, A. B. (2005). Processes of similarity judgment. *Cognitive Science*, 29, 1061–1076.
- Markman, A. B., & Gentner, D. (1993). Splitting the differences: a structural alignment view of similarity. *Journal of Memory and Language*, 32, 517–535.
- Markman, A. B., & Gentner, D. (1996). Commonalities and differences in similarity comparisons. *Memory and Cognition*, 24, 235–249.
- Nelson, D. L., Kitto, K., Galea, D., McEvoy, C. L., & Bruza, P. D. (2013). How activation, entanglement, and searching a semantic network contribute to event memory. *Memory and Cognition*, 41, 797–819.
- Nosofsky, R. M. (1984). Choice, similarity, and the context theory of classification. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 10, 104–114.
- Nosofsky, R. M. (1991). Stimulus bias, asymmetric similarity, and classification. *Cognitive Psychology*, 23, 94–140.
- Pothos, E. M., & Busemeyer, J. R. (2013). Can quantum probability provide a new direction for cognitive modeling? *Behavioral and Brain Sciences*, 36, 255–327.
- Pothos, E. M., Busemeyer, J. R., & Trueblood, J. S. (2013). A quantum geometric model of similarity. *Psychological Review*, 120, 679–696.
- Shepard, R. N. (1987). Toward a universal law of generalization for psychological science. *Science*, 237, 1317–1323.
- Simon, H. A. (1955). A behavioral model of rational choice. *The Quarterly Journal of Economics*, 69, 99–118.
- Smolensky, P. (1990). Tensor product variable binding and the representation of symbolic structures in connectionist networks. *Artificial Intelligence*, 46, 159–216.
- Smolensky, P., Goldrick, M., & Mathis, D. (2014). Optimization and quantization in gradient symbol systems: A framework for integrating the continuous and the discrete in cognition. *Cognitive Science*, 38, 1102–1138.
- Taylor, E. G., & Hummel, J. E. (2009). Finding similarity in a model of relational reasoning. *Cognitive Systems Research*, 10, 229–239.
- Trueblood, J. S., & Busemeyer, J. R. (2011). A quantum probability account of order effects in inference. *Cognitive Science*, 35, 1518–1552.
- Tversky, A. (1977). Features of similarity. *Psychological Review*, 84, 327–352.
- Tversky, A., & Kahneman, D. (1983). Extensional versus intuitive reasoning: the conjunctive fallacy in probability judgment. *Psychological Review*, 90, 293–315.
- Wang, Z., Busemeyer, J. R., Atmanspacher, H., & Pothos, E. M. (2013). The potential of using quantum theory to build models of cognition. *Topics in Cognitive Science*, 1–17.
- Wiles, J., Halford, G. S., Stewart, J. E. M., Humphreys, M. S., & Wilson, W. H. (1994). Tensor models: a creative basis for memory retrieval and analogical mapping. In T. Dartnall (Ed.), *Artificial intelligence and creativity* (pp. 145–159). Kluwer Academic Publishers.
- Wills, A. J., & Pothos, E. M. (2012). On the adequacy of current empirical evaluations of formal models of categorization. *Psychological Bulletin*, 138, 102–125.
- Yearsley, J. M., & Pothos, E. M. (2014). Challenging the classical notion of time in cognition: a quantum perspective. *Proceedings of the Royal Society B*, 281, 1471–1479.