

Similarity: A Transformational Approach

Ulrike Hahn (HahnU@Cardiff.ac.uk)

School of Psychology, Cardiff University,
PO Box 901, Cardiff, CF10 3YG, UK.

Lucy B. C. Richardson (RichardsonLB@Cardiff.ac.uk)

School of Psychology, Cardiff University,
PO Box 901, Cardiff, CF10 3YG, UK.

Nick Chater (Nick.Chater@Warwick.ac.uk)

Institute for Applied Cognitive Science, Department of Psychology,
University of Warwick, Coventry, CV4 7AL, UK.

Abstract

Representational Distortion is a new account of similarity in which the transformation distance between representations determines similarity: entities that are perceived to be similar have representations that are readily transformed into one another, whereas dissimilar entities require numerous transformations. Here we present experimental evidence in favour of this viewpoint.

Introduction

The breadth of cognitive and social contexts in which similarity is invoked as an explanatory construct is vast. Similarity forms part of the explanations of memory retrieval, categorization, visual search, problem solving, learning, linguistic knowledge and processing, reasoning, and social judgment. The two classical approaches to similarity are the spatial account (e.g., Nosofsky, 1986), which represents similarity in terms of distance in psychological space, and Tversky's (1977) contrast model which views similarity as a function of common and distinctive features of the entities under comparison. Both of these accounts have been used successfully in cognitive modeling, however both also suffer from fundamental limitations (see Hahn & Chater, 1997, 1998a, 1998b). These models are restricted in scope by the fact that they define similarity over very specific and simple types of representation: points in space or feature sets. However, the representation of complex real-world stimuli, from faces to auditory scenes, is typically assumed to require *structured* representations, that can explicitly describe objects, their parts, properties and the relations between them. Relational information of this kind cannot readily be encoded using lists of features or dimensional values (Hahn & Chater, 1998a).

The present paper considers a recent theoretical approach to similarity, Representational Distortion (henceforth, RD; Hahn & Chater, 1997; Chater & Hahn, 1997), which aims to provide a theoretical

framework applying to similarity judgements. According to RD, the similarity between two entities is a function of the "complexity" required to "distort" or "transform" the representation of one into the representation of the other. The simpler the necessary transformation, the more similar they are assumed to be.

How can the complexity of the transformation between two representations be measured? At a theoretical level, Hahn and Chater draw on a branch of mathematics, Kolmogorov complexity theory (Li & Vitanyi, 1997) that provides a rigorous and general way of measuring the complexity of representations and transformations between them. In intuitive terms, according to Kolmogorov complexity theory, the complexity of a representation is the length of the shortest computer program that can generate that representation. The idea is that representations that can be generated by a short program are simple; those that require longer programs are complex. We will not consider the virtues of this measure of complexity here, except to note that it supports substantial applications in the cognitive and computing sciences (Chater, 1999).

Kolmogorov complexity has a natural application as a measure of similarity between representations. The simplest measure is the length of the shortest program that "distorts" one representation into the other. According to this viewpoint, the degree to which two representations are similar is determined by how many instructions must be followed to transform one into the other. For example, the conditional Kolmogorov complexity between the sequences 1 2 3 4 5 and 2 3 4 5 6 is small, because the simple instructions add 1 to each digit and subtract 1 from each digit suffice to transform one into the other. In the same way, 1 2 3 4 5 and 2 4 6 8 10 (multiply/divide each digit by 2) are presumed to be similar. On the other hand, 1 2 3 4 5 and 3 5 7 9 11 are viewed as less similar, because two operations are required to transform one into the other (e.g., multiply by 2 and add 1). Finally, two entirely unrelated representations will be maximally dissimilar because there will be no efficient way of transforming

one representation into the other. In this case, the most efficient transformation will involve deleting the first representation, and reconstructing the second from scratch, because there is no shared information between the objects that can be exploited. RD should be viewed as a general framework for understanding similarity, rather than as a specific cognitive account in competition with the spatial or featural views. RD can capture these accounts as special cases (see Chater & Hahn, 1997 for derivations) and thus does not contrast but rather subsumes these accounts. Another motivation for RD is that it aims to provide an explanation for the utility of similarity in inference, for example, to categorize items on the basis of the categories of similar items. To build a concrete psychological account of similarity we need to consider (i) the nature of the mental representations that are relevant to making a similarity judgement; (ii) the set of transformations or instructions that can be used to distort one representation into another; (iii) any constraints on the ability of the cognitive system to discover simple transformations between mental representations.

Despite its generality, RD makes clear empirical predictions. First and foremost, is the prediction that transformations are relevant to similarity. It is this prediction for which the current paper provides empirical support. Crucially, though our own interest in establishing the relevance of transformations to similarity judgments is driven by our research program on RD, the relevance of the general issue of similarity and transformations, and thus of our results, extends beyond our particular theory. As we will demonstrate, the experimental evidence presented here raises substantial problems for classical theories of similarity and raises novel issues for any future work on similarity.

Previous work

The central claim of RD, that similarity is based on transformation distance, has several tentative precursors in the experimental and computational literature. The two most directly relevant experimental studies are by Imai (1977) and Franks and Bransford (1975). Imai proposed that pattern similarity between strings of either filled or unfilled circles was based on transformational structure. He found support for this claim in terms of a qualitative relationship between the number of transformations between two patterns and their judged similarity. However, no statistical analysis was performed. Franks and Bransford (1975) sought to extend Posner and Keele's (1968) work on prototype abstraction, replacing the original random dot patterns with simple geometric figures. Underlying the stimulus set was a prototype that was not shown during training; all other items in the stimulus set were derived from this prototype through the application of one or more simple transformations. Recognition ratings were directly related to transformational distance to the

prototype, with the prototype itself receiving the highest rating. Finally, the account has some resonance in the perception literature where transformational explanations have been used to explain figural regularity or "goodness", as well as figuring in theories of object recognition.

In summary, there is currently no clear experimental evidence for the importance of transformations in the context of similarity, despite previous research hinting at this idea. Three experiments were designed to address this issue. Each of the experiments shared the same basic correlational design and differed only in their stimulus materials. The aim was to establish transformational distance as a predictor of perceived similarity, while at the same time providing evidence for the limitations of featural (and spatial) accounts.

Experiment 1

Experiment 1 was based on Imai (1977), and uses sequences of filled or unfilled circles. Transformational distance was manipulated in terms of the number of operations such as mirror imaging, reversal, phase shift, insertion and deletion that were necessary to convert the test stimulus into the target. This is best illustrated with an example stimulus pair, shown below.

○●●●○○○●●●●○ ○●●○●○○○●●●○

The two rows of "blob" patterns can be transformed into one another through the application of a single operation that "reflects" one row to create the other row in a mirror image. This prediction contrasts with that of the featural account (Tversky, 1977), if we adopt the natural assumption that features correspond to individual blobs. According to this viewpoint, blob patterns should be similar to the extent that they overlap.


Participants 35 undergraduates psychology students.

Materials The stimuli consisted of strings of black and white blobs, presented in pairs. The key transformations were phasic, reversal and mirror, with the addition of insertions and deletions. Any one of these in isolation constituted a single transformation. There were 16 examples of single transformations in this experiment, (4 each of phasic, reversal, mirror and deletion). Multiple transformations were achieved by combining two or more of the above operations. The total set of 56 comparison pairs consisted of 16 examples of two transformational changes (four each of reversal & mirror, reversal & phasic, deletion & mirror and insertion & phasic) and 16 examples of three transformational changes (four each of deletion, reversal & mirror, deletion, reversal & phasic, insertion, reversal & mirror and insertion, reversal & phasic). As a control, there were also 8 pairs of stimuli that were unrelated (or so multiply transformed as to make the transformations unperceivable). Each pair of stimuli was printed horizontally onto a single sheet of paper

together with brief instructions and a rating scale from 1 (very dissimilar) to 7 (very similar). These sheets were then placed into a different random order for each participant and bound into a booklet.

Results

Bivariate correlations between number of transformations and mean similarity rating of each item were found to be highly significant with Spearman's rho = -.69, $p < .005$. The comparison featural model which left aligned the two rows of blobs and counted the number of (mis)matching features fared considerably worse: Spearman's rho = -.28, $p < .05$. Analysis of individual subject ratings confirmed these findings, revealing great conformity across participants: 25 of 35 participants showed significant correlations as predicted. Such consistency was not found using the featural model, with only 8 of the 35 participants showing a significant correlation. The general relationship between number of transformations and mean similarity ratings is graphed below. The results suggest, somewhat surprisingly (see e.g. Shepard, 1987), an approximately linear relationship.




Discussion

Experiment 1 provides evidence of a statistically significant relationship between transformation distance and perceived similarity to complement Imai's (1977) more qualitative data. The results of Experiment 1 thus confirm both Imai's original intuition and the predictions of the RD framework. In direct comparison, the featural model fares considerably worse in predicting participants' ratings. Consequently the results also provide evidence for limitations in featural approaches. It is, of course, possible that more powerful featural descriptions of the data could be found, but, at present, none are available. Crucially, any putative featural explanation of this kind requires an independent motivation of the features adopted, that is, the postulated features must themselves not be motivated exclusively by salient transformations. Otherwise, the featural description becomes an entirely redundant

mimicry. The materials of Experiment 2 make this point more clearly.

Experiment 2

This experiment used simple geometric shapes related by different transformations, e.g., the pair of items shown below has a transformation distance of two as they can be made identical through rotation and color change of one object part. Here, there is no obvious way to apply a featural model for contrast purposes. Furthermore, many of the "features" such as the orientation of an item in a pair where one has been rotated are salient only because of the relevant transformations. This means that central object "features" will be derivative on the transformations present: e.g., orientation is unlikely to have cognitive salience in a comparison until orientation is manipulated through rotations. Consequently, though it might, in principle, be possible to derive featural descriptions for our stimulus items, these descriptions would be likely to implicitly underscore the importance of transformations, rather than providing an alternative to relying on transformations. As in Experiment 1, the prediction is that number of transformations will be negatively correlated with degree of perceived similarity.




Participants 21 psychology undergraduates.

Materials There were three alternative "target" geometric shape line drawings. The construction of the stimulus sets began with the target shapes to which a single transformation was then applied. Examples of single adjustments were stretching the whole object or changing all striped areas to filled-in black areas. Multiple transformations were constructed by using a combination of such techniques, one at a time. For each of the three target stimuli there four examples each of one, two, three, four and five transformations. This made a set of 20 pairs of pictures for each base target, yielding 60 in total. Each transformed geometric shape was printed onto a separate page together with its corresponding "target", which was always placed to the right of the transformed item. At the top of each page were a set of instructions and a rating scale from 1 (very dissimilar) to 7 (very similar).

Results

Bivariate correlations between number of transformations and mean similarity rating of each item were highly significant with Spearman's rho = -.89, $p = .000$. Analysis of individual participants' ratings again revealed great consistency across subjects, with 19 of the 21 participants showing a significant correlation as predicted. The relationship between number of transformations and mean similarity ratings as graphed below closely matches that found in Experiment 1. The relationship between transformation distance and similarity is again approximately linear.



Discussion

Experiment 2 provide further supports the role of transformations in the context of similarity, mirroring the results of Experiment 1 with very different stimulus materials. The materials of Experiment 2 also illustrate how even natural "features", such as orientation are influenced by transformations. Many of the very "features" that a featural account might posit seem salient due to the transformational relationships between the two compared objects. This is indicative of the general bi-directional relationship posited by RD theory between object representation and transformation, with perceived transformations influencing which aspects of an object become salient and vice versa.

Experiment 3

Experiment 3 sought to take the argument against featural and spatial representations one step further, by using materials for which such representation schemes are obviously inadequate, because they depend on relational structure. We used 3D objects assembled

from (typically) three Lego bricks: one large brick, colored blue; a medium size yellow brick; and a small red brick. Each similarity comparison comprised two objects assembled from these three bricks, albeit in different spatial arrangements. Despite the extreme simplicity of these stimuli, relational information (i.e., information about relative position, such as, for example, that the red brick is on top of the yellow brick) is paramount to the representation of the composite objects. However, the appeal of these materials is not limited to the difficulties they pose for featural or spatial accounts. From a transformational perspective, the Lego brick objects are of interest for two reasons. First, they allow an initial examination of the role of transformations in the similarity assessment of real-world objects, albeit maximally simple ones. Second, these materials support a whole new range of transformations to complement those investigated in Experiments 1 and 2. Our assumption, here, was that the judged similarity between pairs of objects would be determined primarily by the physical manipulations required to turn a target object into the comparison object.


Participants 27 psychology undergraduates.

Materials The stimuli were based upon an initial "target" array of three Lego bricks, (a four-point square red brick, upon a six-point oblong yellow brick, on top of an eight-point oblong blue brick) arranged into a particular three-dimensional structure. Apart from the two examples that were chosen to be totally unrelated to the target, all of the Lego stimuli were constructed by transforming the original target object a set number of times, prior to the experiment. The researcher began with the target arrangement and made adjustments to it, each of which constituted one transformation. For example, one adjustment (or transformation) could involve moving an object within the arrangement, substituting a brick for one of a different size or colour, adding an additional brick, subtracting an existing brick, or changing the order of the bricks within the arrangement. The transformations required to create each arrangement were counted and the entire set of stimuli constructed so that it comprised of 42 items (10 examples of one and 10 of two transformations, and five examples each of three, four, five and six transformations away from the "target". In addition there were two items that were unrelated--multiply transformed--stimuli). Once these arrangements of Lego bricks had been constructed, each was glued into a permanent structure.

Procedure Participants were shown the "target" Lego brick object. They rated how similar they perceived each stimulus to be to the target, on a scale of 1 (very dissimilar) to 7 (very similar). Every participant rated all the Lego stimuli within the set of 42 items.

Results

Bivariate correlations between number of transformations and mean similarity rating of each item were again highly significant, Spearman's $\rho = -.76$, $p < .005$. Analysis of individual participants' ratings again revealed great consistency across subjects, with all 35 exhibiting a significant correlation between number of transformations and rated similarity. The general relationship between number of transformations and mean similarity ratings (shown below) is again very similar to that found in Experiments 1 and 2.



Discussion

Despite the very different materials and set of relevant transformations, the results of Experiment 3 closely match those of Experiments 1 and 2. Again, they provide evidence for the importance of transformations in explaining similarity judgements, and are difficult to account for in terms of the featural or spatial views of similarity, which cannot easily handle relational information. Thus, these results are in line with the predictions of a central tenet of the RD account: that the transformational relationship between representations of two objects determines their judged similarity. But the results of Experiment 3 also have broader implications. The inherently relational nature of the materials in Experiment 3 poses a problem for any representational scheme which does not allow structured representations; conversely it lends support to any account such as structural alignment theories to which such structured representations are central. Similarly, the result that number of transformations is a significant predictor of perceived similarity lends support to the general notion

of an influence of transformations on similarity, whether or not the particular framework of RD theory is adopted.

General Discussion

The results of all three experiments provide robust evidence as to the importance of transformations in explaining similarity judgments, across a variety of different stimulus types. These findings support the central tenet of the RD theory of similarity, that similarity is based on the complexity of the transformation between the representation of two items. These results also provide new evidence for the limitations of classical accounts of similarity. All three experiments provide evidence against purely featural views of similarity. Experiment 1 provides a direct test. The version of the featural model we tested (assuming that features correspond to blobs) is not the most sophisticated featural description possible, given that, in principle, any property including all higher-order regularities such as "symmetry" etc. could be posited as features (Tversky, 1977). Crucially, however, a more sophisticated featural account which succeeds in providing comparable or even superior data fits must not only first be found, it must also be independently motivated. Given that theories can be stretched beyond all recognition through the addition of suitable post hoc auxiliary assumptions, a crucial factor in evaluating competing accounts must not only be whether an account can be made compatible with a particular pattern of data, but also whether it in any way predicted it.

In Experiment 1, there is nothing in featural theories of similarity that would naturally give rise to the predictions made on the basis of transformations in this experiment. The sequences of filled circles lend themselves naturally to a featural decomposition on a one by one basis due to the fact that the "object" is readily parsed into a set of individual circles. Many of the relevant "features" of the geometric shape stimuli in Experiment 2 become cognitively salient only through transformational contrast between the two comparison objects (for example, the feature "orientation" highlighted by the transformation "rotation"). Consequently, transformations are explanatorily prior. The use of simple formations of Lego bricks in Experiment 3 demonstrates the central representational weakness of featural accounts - their inability to deal with structured representations and thus adequately represent relational information. The challenge presented by Experiment 3 is to identify even a remotely suitable featural description, given the inherent relational nature of materials and transformations. The limitations of featural accounts exposed by this series of experiments is equally shared by spatial models of similarity, whether they are based on multi-dimensional scaling or standard connectionist networks. It must again be stressed that from the perspective of RD theory, featural and spatial accounts of similarity are not

wrong, they are simply too restricted to cope with the flexibility of transformations available to the cognitive system.

Interestingly, issues related to this research have been raised in philosophy. Goldman (1986) suggests that the lawfulness of human similarity judgments might be furthered by an inherent preference ranking for transformations, which comes in to play where multiple transformational sequences could link the same stimulus pair. This question links closely with a central issue for future research, that of the relative "cost" or "weight" of individual transformations. Single transformations need not be equal in cost or 'effort'. Such inequalities arise automatically in the theory of Kolmogorov complexity, the general mathematical framework on which RD theory draws. Here, deletions, for example, tend to be less costly than insertions, because deletions only require a specification sufficient to identify the component for deletion, whereas insertions require a complete specification of the additional component. What weightings of this kind, if any, are intrinsic to the cognitive system is an issue we are currently seeking to determine through the investigation of perceived similarity for different single transformations. Information as to relative transformational costs will be crucial for more detailed cognitive modeling and thus constitute a major issue for future research. Another potentially interesting area is to apply the approach to different domains; particularly those that appear to require structured representations where RD can be utilized in a straightforward way. For example, two postures of a hand, in terms of a specification of joint angles can be compared simply. Given the transformations likely to be salient in cognitive processing involving motor control, we might expect that 'similar' hand positions would correspond to positions that can readily be transformed into each other.

We have presented a new account of similarity, Representational Distortion, according to which the judged similarity between a pair of items is determined by the complexity of the transformation between the mental representations of those items. We have tested the central tenet of the account in three experiments, finding that transformational complexity is, indeed, inversely related to similarity. These results present a challenge for other accounts of similarity, based on feature comparison or spatial distance; and they indicate that the view that similarity can be explained in terms of transformation merits further theoretical and empirical investigation.

Acknowledgements

This work was funded by the Leverhulme Trust, and NC is supported by EU Grant RTN-HPRN-CT-1999-00065.

References

Chater, N. (1999). The search for simplicity: A fundamental cognitive principle? *Quarterly Journal of Experimental Psychology*, 52A, 273-302.

Chater, N. & Hahn, U. (1997) Representational Distortion, Similarity, and the Universal Law of Generalization. In, *Proceedings of the Interdisciplinary Workshop on Similarity and Categorization*, SimCat97, University of Edinburgh. Dept. of Artificial Intelligence, University of Edinburgh.

Franks, J.J. & Bransford, J.D. (1975) Abstraction of Visual Patterns. *Journal of Experimental Psychology: Human Perception and Performance*, 4, 65-74.

Goldman, A.I. (1986) *Epistemology and Cognition*. Cambridge:MA.Harvard University Press.

Hahn, U. & Chater, N. (1997). Concepts and similarity. In K. Lamberts and D. Shanks (Eds.). *Knowledge, concepts and categories*, (pp. 43-92). Hove, England: Psychology Press.

Hahn, U. & Chater, N. (1998a). Similarity and rules. Distinct? Exhaustive? Empirically distinguishable? *Cognition*, 65, 197-230.

Hahn, U. and Chater, N. (1998b) Understanding Similarity: a Joint Project for Psychology, Case-Based Reasoning and Law. *Artificial Intelligence Review*, 12, 393-427.

Imai, S. (1977) Pattern Similarity and Cognitive Transformations. *Acta Psychologica*, 41, 433-447.

Li, M. & Vitanyi, P. (1997). *An introduction to Kolmogorov complexity and its applications* (2nd edition). New York: Springer-Verlag.

Nosofsky, R. M. (1986). Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology: General*, 115, 39-57.

Posner, M. I. & Keele, S. W. (1968). On the genesis of abstract ideas. *Journal of Experimental Psychology*, 77, 353-363.

Shepard, R. N. (1987) Toward a universal law of generalization for psychological science. *Science*, 237, 1317-1323.

Tversky, A. (1977) Features of similarity. *Psychological Review*, 84, 327-352.