

Similarity as Tractable Transformation

Moritz Müller (moritz.mueller@math.uni-freiburg.de)

Department of Mathematics
79104 Freiburg, Germany

Iris van Rooij (i.vanrooij@donders.ru.nl)

Radboud University Nijmegen
Donders Institute for Brain, Cognition, and Behaviour
6525 HR Nijmegen, The Netherlands

Todd Wareham (harold@cs.mun.ca)

Department of Computer Science
St. John's, NL A1B 3X5 Canada

Abstract

According to the transformational approach to similarity, two objects are judged to be more similar the simpler the transformation of one of the object representations into the other. This approach draws inspiration from the mathematical theory of Kolmogorov complexity, but otherwise remains an informal theory to this day. In this paper we investigate several different ways in which the informal theory of transformational similarity can be understood, providing a formalization for each possible reading. We then study the computational (in)tractability of each formalization for a variety of parameter settings. Our results have both theoretical and empirical implications for transformational approaches to similarity.

Keywords: similarity; representational distortion; computational complexity; intractability; parameterized complexity

Introduction

Consider the two sequences $\bullet\circ\circ\bullet\circ\circ\bullet$ and $\circ\circ\bullet\circ\circ\bullet$. Even though these two sequences differ in every element they are nevertheless quite similar. Moreover, these two sequences may be judged to be more similar than two sequences that share more elements, such as, for example, $\bullet\circ\circ\bullet\circ\circ\bullet$ and $\bullet\circ\circ\bullet\circ\circ\bullet$. A possible explanation for this is that the first two sequences are related by a single simple transformation (e.g., inversion), whereas no such simple transformation seems to relate the second two sequences. This explanation accords with the Representational Distortion (RD) theory of similarity (Chater & Hahn, 1997; Chater & Vitányi, 2003; Hahn, Chater, & Richardson, 2003).

According to RD, in general, two object representations are judged to be more similar the fewer basic transformations are required to transform one object representation into the other. The basic idea underlying RD already existed in the late 70s (Imai, 1977), but in recent years it has gained in explanatory strength both on empirical (Hahn et al. (2003); but see also Larkey and Markman (2005)) and theoretical grounds (Chater & Vitányi, 2003). As with any computational-level theory, the plausibility of RD depends not only on how well it can predict or describe human similarity judgments, but also on the existence of tractable algorithms for computing similarity under the RD model (Frixione, 2001; Tsotsos, 1990; van Rooij, 2008).

The roots of RD seem to be cause for worry about RD's ability to meet the tractability constraint. The theory of RD is inspired by the mathematical notion of Kolmogorov complexity. The Kolmogorov complexity $K(a)$ of an object representation a is the length of the shortest program that, when run on the empty input, constructs a . The notion can also be used to express the transformational distance from an object representation b to a , by replacing in the former statement 'empty input' by b . The resulting conditional Kolmogorov complexity is denoted $K(a|b)$. It is known that K is an uncomputable function—i.e., there does not exist any algorithm that computes $K(a)$ or $K(a|b)$ for all a and b . This means that conditional Kolmogorov complexity as a measure of similarity, as proposed by Chater and Hahn (1997) and elaborated in Chater and Vitányi (2003),¹ does not even meet the computability constraint, let alone the tractability constraint on computational-level theories.

Chater and Vitányi (2003, p. 347) acknowledge that their measure of similarity is an “‘ideal’ notion in the sense that it ignores the limitations on processing capacity.” To render the transformational approach to similarity more psychologically and computationally realistic, Hahn et al. (2003) proposed the current version of RD. This version holds on to the idea that similarity between object representations a and b is a function of their transformational closeness, but rather than referring to the shortest program transforming a to b , it restricts attention to programs that are sequential applications of operations from a particular set of ‘basic transformations’.

The question now arises if RD is tractably computable. Currently, it is impossible to answer this question, because RD remains so far a verbal theory. In order to determine the (in)tractability of RD we need to make the computational problem of finding shortest transformations mathematically

¹For completeness, we remark that Chater and Vitányi (2003) formulated a symmetric measure of similarity (based on $K(a|b) + K(b|a)$) because they were interested in its metric properties and demonstrating that it can explain Shepard's Universal Law of Generalization. Empirical studies show that similarity judgments need not be symmetrical (Tversky, 1977). As our focus is on the computational (in)tractability of transformational approaches to similarity, we do not force transformational distances to be symmetric.

precise. In this paper we investigate several different ways in which this can be done. We will show that each proposal, if unrestricted, faces the problem of intractability. Our investigation does not stop at this observation, however. We adopt a method for identifying sources of this intractability (van Rooij, Stege, & Kadlec, 2005; van Rooij, Evans, Müller, Gedge, & Wareham, 2008) and having identified such sources we make recommendations to RD theorists on how their theories may be restricted so as to ensure tractability.

Computational-level Models

According to Hahn et al. (2003) “RD theory seeks to characterize the computational level problem involved in determining similarity.” A first possible, informal characterization of the computational problem involved is the following:

REPRESENTATIONAL DISTORTION (version 1) [RD1]

Input: Two representations a and b and a set of basic transformations T .

Output: The length of a shortest sequence of basic transformations from T transforming a to b .

Note that in this problem formulation, the set of basic transformations T is not a constant but can in principle vary independently from a and b . This seems in line with the thinking of Hahn et al., as they hypothesized different T for the different types of stimuli that they used in their experiments. A question raised by the possibility of varying T is how the cognizer knows which T to adopt in a particular situation. One possibility could be that the modality of a and b determines T (e.g., a picture can be rotated and a sound can be change in pitch, but not vice versa). This cannot accommodate, however, the variations in T allowed by Hahn et al., as all their stimuli were in the visual modality.

It is also theoretically plausible that T depends not on the nature of a and b alone, but also the context in which a and b are being compared.² To illustrate consider the example sequences in Table 1. In this table, we have $a_3 = c_3$ and $b_3 = d_3$. Ignoring the context in which the comparisons are made, the comparison of a_3 with b_3 may appear equivalent to the comparison of c_3 with d_3 . But note that in the context in which the pair a_3, b_3 appears it is more natural to see the two sequences as related via the transformation of ‘mirroring’, simply because all a_i, b_i pairs are related by the transformation ‘mirroring’. On the other hand, in the context in which the pair c_3, d_3 appears it is more natural to see the two sequences as related via the transformation of ‘inversion’, simply because all c_i, d_i pairs are related by the transformation ‘inversion’.

The preceding shows that the context in which a and b are compared determines in part which basic transformations are in the set of transformations used for their comparison. This

²This view seems consistent with Imai’s: “let us define four cognitive transformations which we assume basic ... within the context of the set of configurations used in our experiment ... [but] when the set of configurations are less restricted, the basic set of transformations must be expanded” (Imai, 1977, pp. 434-435).

Table 1: Illustration of the effect of context on the type of transformations used to related pairs of objects.

index	a	b	c	d
1	●●○○○○	○○○○●●	●●○○○○	○○●●●●
2	●●○○○○	○○○○●●	●●○○○○	○○●●●●
3	●●●○○○	○○○●●●	●●●○○○	○○○●●●

idea is expressed by the following reformulation of the computational problem underlying RD.

REPRESENTATIONAL DISTORTION (version 2) [RD2]

Input: Two representations a and b , a set of basic transformations \mathcal{T} , and a ‘context’ C .

Output: A number that equals the length of a shortest sequence of basic transformations from $T_C \subseteq \mathcal{T}$ transforming a to b , where T_C is a set of transformations that is ‘most relevant’ for ‘context’ C .³

Generally both ‘context’ and ‘relevance’ remain elusive concepts in cognitive science, but in the specific context of RD-based similarity judgments we may nevertheless try to make more precise what they could mean.

We start with the notion of ‘context’. The illustration in Table 1 shows that one form of context for a comparison of a and b is other comparisons that are made at the same time as or briefly before the comparison between a and b . This special type of context yields the following special case of the foregoing problem:

REPRESENTATIONAL DISTORTION (version 3) [RD3]

Input: Two representations a and b , a set of basic transformations \mathcal{T} , and a set of pairs of object representations X with $(a, b) \in X$.

Output: A number that equals the length of a shortest sequence of basic transformations from $T_X \subseteq \mathcal{T}$ transforming a to b , where T_X is a set of transformations in \mathcal{T} that is most relevant for context X .

Two possible interpretations seem open for RD2 and RD3:

- (i) The set T_C may be (non-inferentially) given to the cognizer; or
- (ii) the cognizer needs to compute T_C per context.

According to option (i), it is assumed that no computation is required to know what subset $T_C \subseteq \mathcal{T}$ is considered relevant for the context at hand. In the remainder of this paper, let RD2 and RD3 be so interpreted.

According to option (ii), part of the problem of computing the similarity between a and b is the computation of this T_C . The latter is consistent with the view expressed by Larkey and Markman (2005, p. 1071), when they write “transformational

³Here, \mathcal{T} may be thought of as the master set of all possible transformations associated with a specific ‘sensory modality’, e.g., all basic visual transformations.

accounts require additional processing to determine the transformations that distort one representation into the other.”⁴

Interpretation under option (ii) requires us to make precise what makes a set of transformations $T_C \subseteq \mathcal{T}$ ‘relevant’ for some context, which in turn requires an elaboration of what ‘context’ is in RD2. For the special type of context X defined in RD3, we propose the following hypothesis: Cognizers consider a set T_C to be *relevant* for judging the similarity between a pair of representations $(a, b) \in X$ if many pairs in X can be transformed with a short sequence of transformations from T_C , while T_C is as small as possible with this property. Here, “many” and “short” can be interpreted as being larger or smaller than a given threshold. Following these stipulations, the task of finding a set of relevant transformations can be formalized by the following computational problem.

RELEVANCE

Input: A set of basic transformations \mathcal{T} , a set of pairs of object representations X , and integers s and w .

Output: A set of transformations $T_X \subseteq \mathcal{T}$ of minimum size such that for at least s pairs $(a, b) \in X$ there exists a sequence of transformations in T_X having length at most w that when applied to a yields b .

Our operationalisation of ‘relevance’ can be seen as adopting the same principle of parsimony (or simplicity) that has been argued to underly RD (Chater & Vitányi, 2003), i.e., assume no more basic transformations than necessary to relate as many as possible representation pairs in the current set in ways as simple as possible. This operationalisation yields the fourth and final version of RD considered in this paper:

REPRESENTATIONAL DISTORTION (version 4) [RD4]

Input: Two representations a and b , a set of basic transformations \mathcal{T} , a set of pairs of object representations X with $(a, b) \in X$, and integers s and w .

Output: A number that equals the length of a shortest sequence of basic transformations from $T_X \subseteq \mathcal{T}$ transforming a to b where T_X is a solution of $\text{RELEVANCE}(\mathcal{T}, X, s, w)$.

The notions of ‘object representation’ and ‘transformations’ have so far remained informal. Below we present formalizations for each.

Formalizing Representation and Transformation

Following Chater and Vitányi (2003) we assume that an object representation is a finite binary string. Nothing seems lost with this assumption, as any finite object can be represented by a finite binary string. Chater and Vitányi (2003, p. 354) argued for this generality as follows:

⁴We think that option (ii) is theoretically to be preferred, because the number of possible contexts simply seems too vast to explicitly store T_C for every possible context C . Be that as it may, we present (in)tractability results for both interpretations, leaving it up to the RD modeler which of these to adopt.

“[each object we are interested in] can be described by using, for example, English. That means we can describe every object by a finite string in some fixed finite alphabet. By encoding the different letters of that alphabet in bits (0’s and 1’s) we reduce every description or representation of the object to a finite binary string. A similar argument presumably holds for the physical manner by which an object is represented in an agent’s cognitive system.”

As for the formalization of the notion of transformation we choose to be equally general, defining a transformation as a Boolean circuit (see Fig. 1 and 2). Again, this entails no loss of generality because any function from binary strings (of some fixed length n) to binary strings (of some fixed length m) can be computed by a Boolean circuit. Transformations, like ‘mirroring’ and ‘inversion’, acting on strings as a whole are viewed as families of circuits, one circuit for each length doing the transformation.

A Boolean Circuit




Figure 1: Illustration of a Boolean circuit. In our formalization of RD, the binary string on the input nodes represents an encoding of object a and the binary string on the output nodes represents an encoding of object b .

Now that we have formalized the notions of ‘object representation’ and ‘transformations’, RD1 and RD4 have become well-defined problems whose (in)tractability can be subjected to formal mathematical analysis. Although RD2 and RD3 remain informal, as they leave e.g. the notion of ‘relevance’ undefined, formal (in)tractability results for RD1 can be translated to informal (in)tractability results for RD2 and RD3.⁵ For technical reasons, formal results cannot be derived for informal problems such as RD2 and RD3, but for purposes of assessing the plausibility of cognitive models the distinction between formal and informal (in)tractability results can be safely ignored, as we will do in the remainder of this paper.

Representational Distortion is Intractable

Using the formalizations above, it can be shown that computing similarity under the considered RD models can only be done by algorithms that use a superpolynomial (e.g., ex-

⁵To do so we will assume that for all possible T in RD1 there exists a context C such that $T_C = T$ in RD2, and there exists a context X such that $T_X = T$ in RD3. Then any algorithm solving RD2 or RD3 also solves RD1.

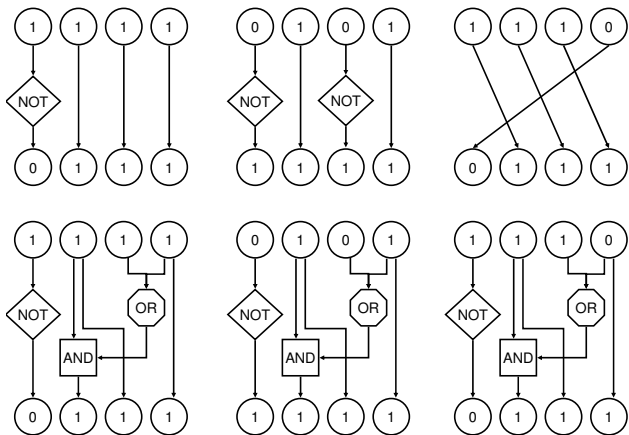


Figure 2: Illustration of RELEVANCE under the Boolean circuit formalization. The three pairs of objects $(1111, 0111)$, $(0101, 1111)$, and $(1110, 0111)$ can be each be transformed by one of the three circuits shown at the top. No one of these circuits can transform these three pairs of objects and pair $(0011, 1001)$ in Figure 1, but the circuit at the bottom is capable of transforming them all. If all these circuits were psychologically possible, i.e. in \mathcal{T} , then in our formalization of RELEVANCE the circuit at the bottom would be considered relevant for the set of pairs $X = \{(1111, 0111), (0101, 1111), (1110, 0111), (0011, 1001)\}$, whereas the other circuits would not be.

ponential) amount of time—i.e., the time cannot be upper-bounded by a function n^c , where n is a measure of input size and c is some constant.⁶ Super-polynomial time algorithms are generally considered computationally intractable because they take unrealistically long for all but small inputs. To illustrate, consider an exponential-time algorithm running in a time proportional to 2^n . Such an algorithm would need to make on the order of 1,000,000,000 computational steps for an input of size $n = 30$. For larger inputs, say $n = 60$, the number of computational steps gets close to the number of seconds that has past since the birth of the universe.

The upshot of the intractability of the considered RD models is that they are all psychologically implausible⁷ unless the right restrictions are posed on the hypothesized domain of inputs. To find such restrictions we will attempt to identify sources of intractability in RD models.

Identifying Sources of Intractability

We adopt the method for identifying sources of intractability described in van Rooij et al. (2008) (see also van Rooij and Wareham (2008)). The method works as follows.

⁶For proofs see the online *Supplementary Materials*, available at <http://www.nici.ru.nl/~irisvr/supplement09.pdf>. These results assume that $P \neq NP$, a mathematical conjecture that is unproven but has strong empirical support. The interested reader is referred to Garey and Johnson (1979) for more details.

⁷For a full treatment of common objections (based on, e.g., heuristics, approximation, parallelism) to the tractability constraint that we assume here, see van Rooij (2008).

First, one identifies a set of problem *parameters* $K = \{k_1, k_2, \dots, k_m\}$ in the problem P under study (for us, the different versions of RD discussed in the Introduction each constitute such a problem P). Then one tests if it is possible to solve P in a time that is exponential (or worse) only in K but polynomial in the size of the input.⁸ If this is the case, then P is said to be *fixed-parameter (fp-) tractable* for parameter set K , and otherwise it is said to be *fp-intractable* for K .

Observe that if a parameter set K is found for which P is fp-tractable then the problem P can be solved quite efficiently, even for large inputs, provided that the members of K are relatively small. In this sense the “unbounded” nature of K can be seen as a reason for the intractability of P . Therefore we call K a *source of intractability* of P .

RD models have several natural parameters, each of which may be a source of the intractability inherent in the general problems postulated by these models. Table 2 gives an overview of the parameters that we consider here.

Table 2: Overview of parameters that may be sources of intractability for Representational Distortion models.

Name	Definition
k	length of the shortest sequence of transformations transforming a to b
t	size of the set T
ℓ_1	maximum of the lengths of a and b
ℓ_2	maximum of the lengths of a , b , and all intermediate representations created in transforming a into b
t_c	size of the set T_C
m	size of the set \mathcal{T}
w	the maximum dissimilarity of pairs in X that can be related by transformations from T_C

Results and Discussion

We present a list of fp-tractability and fp-intractability results for sets of parameters selected from Table 1.⁹ For the proofs we refer to the *Supplementary Materials* published online (see footnote 6). We start by considering RD1:

RD1 is:

1. fp-intractable for parameter set $\{t, \ell_1\}$.
2. fp-intractable for parameter set $\{k, \ell_1\}$.
3. fp-tractable for parameter set $\{t, k\}$.
4. fp-tractable for parameter set $\{\ell_2\}$.

⁸More formally, this would be a time on the order of $f(k_1, k_2, \dots, k_m)n^c$, where f is an arbitrary computable function, n is a measure of the overall input size, and c is a constant

⁹We will work under the assumption $FPT \neq W[1]$. Like $P \neq NP$, this mathematical conjecture is unproven but has strong empirical support. The interested reader is referred to Downey and Fellows (1999) and Flum and Grohe (2006) for more details.

By Result 3, the assumption that t (the number of transformations) is relatively small and the assumption that k (the length of the shortest transformation sequence) is relatively small are together sufficient to render RD1 tractable.¹⁰ By Results 1 and 2, neither of these assumptions can be dispensed with. Results 1 and 2 also show that small values of ℓ_1 (the maximum length of the two given representations) cannot itself make RD1 tractable; this, in combination with Result 4 (which shows that small values of the closely-related parameter ℓ_2 can make RD1 tractable), highlights the importance of intermediate (and not just given) object-representations to the complexity of RD1.

How plausible is it that ℓ_2 , k and/or t are relatively small? It seems psychologically implausible to assume a relatively small ℓ_1 , because humans often judge the similarity of quite complex objects (e.g., two buildings, two faces, two movies); as $\ell_1 \leq \ell_2$, this limits the utility of Result 4 for rendering RD1 tractable in practice. It seems plausible, however, that k is severely limited in size, as humans unlikely have unbounded sensitivity for degrees of (dis)similarity of highly dissimilar objects and it is reasonable to assume a relatively small threshold on k above which no further transformation is attempted and the two objects are simply judged as maximally dissimilar (Hahn et al., 2003, p. 26). If T would be interpreted as ‘all psychologically possible transformations, independent of context’ then its size, t , would unlikely be small. Yet, a small bound on t may be psychologically plausible if we consider the set of transformations, T , to be selected per context. In that case, however, the RD modeler could better adopt one of the models RD2, RD3 or RD4, in which this extra selection step is explicated.

Recall that we consider RD2 and RD3 under option (i), i.e., that the context specific set of transformations, T_C , is given to the cognizer at no computational cost. Under quite mild assumptions on possible formalizations of RD2 and RD3 under option (i) we were able to prove the following two results.

RD2 and RD3 are both:

5. fp-tractable for parameter set $\{t_c, k\}$.
6. fp-tractable for parameter set $\{\ell_2\}$.

Result 5 shows that RD models which assume no computational cost in selecting the set of transformations relevant for a particular context can be rendered tractable by making the psychologically plausible assumptions that both t_c and k are relatively small. An important question raised by this observation is whether it is psychologically plausible to assume that T_C is non-inferentially given (cf. footnote 4).

We investigated whether it was really necessary to make this assumption by analyzing the fp-(in)tractability of RD4.

¹⁰To be more precise, using the best algorithm known to date, RD1 can be computed in a time t^k times a polynomial in the input size. This means that the similarity of a to b can be efficiently computed on inputs in which the set of possible transformations T is not too large (say, $t \leq 10$) and the object representations a and b are not too dissimilar (say, $k \leq 5$). If a more efficient algorithm is found, feasible values of t and k will increase.

Recall that RD4 is the model based instead on option (ii), in which computing the relevant set T_C is a subcomputation involved in computing the similarity between objects a and b . In RD4 this subcomputation is modeled by the computational problem RELEVANCE. We have the following results:

RD4 is:

7. fp-intractable for parameter set $\{m, \ell_1\}$.
8. fp-intractable for parameter set $\{w, t_c, \ell_1\}$.
9. fp-tractable for parameter set $\{m, w\}$.
10. fp-tractable for parameter set $\{\ell_2\}$.

Result 9 shows that the model RD4 is tractable if the modeler is willing to assume that both m and w are relatively small. In the light of Results 7 and 8 one cannot dispense with either of these assumptions. Yet, once again, we must ask how plausible these assumptions are.

As for assuming a small bound on w (the maximum dissimilarity of pairs in the context set X), we are unsure if this is plausible or not. One may imagine that in real-world settings only those objects that can be related to each other by simple transformations will naturally come to be seen as belonging to the same set X . This could be the case, for example, if different X s were to correspond to natural categories of objects (Pothos & Chater, 2002; Rosch & Mervis, 1975). However, in experimental settings, it may well be possible to create artificial sets X with large w . An empirical prediction that we can derive from RD4 is that humans would take particularly long to determine the set of transformations that is relevant for a set X with large associated w , and consequently also long to compute the similarity between objects in such a set.

Granting even that in the real-world w may reasonably be assumed to be small, we are nevertheless left with the problem that m (the size of the master set \mathcal{T}) can be quite large. We see no psychologically plausible way to ensure that \mathcal{T} can be kept small, as there seems to be an in principle unbounded number of ways in which cognizers may learn to see relationships between pairs of objects (French & Anselme, 1999).

Conclusion

We have presented computational-level models for Representational Distortion theory. As it turns out, all RD models considered are computationally intractable in general. We investigated which assumptions may render the models tractable and found several options open to RD modelers.

To start, the two-part assumption that similarity judgments only concern rather simple objects (i.e. objects with a short representation) *and* all intermediate object-representations encountered in the transformation process are rather simple as well (let us call this Assumption 1) suffices to render RD models tractable. It seems reasonable in general to assume that transformations do not dramatically blow up the sizes of manipulated object-representations, but as human cognizers seem to be able to judge the similarity of quite complex objects in the real world, Assumption 1 as a whole lacks psychologically plausibility.

A more promising option for RD modelers is to assume that humans compute exact transformational distances only for objects that are at least somewhat similar (Assumption 2). This assumption is consistent with the view expressed by Hahn et al. (2003, p. 26), who write “for pairs of items for which no transformational relationship is discernible, RD predicts simply that these items are maximally dissimilar.” Third, there is the option to assume that the set of transformations that is used for a given context of comparisons is quite small (Assumption 3). This assumption holds for the situations investigated by Imai (1977) and Hahn et al. (2003), where the similarity of objects could be rated by only a handful of basic transformations. Fourth, the RD modeler could assume that the set of transformations to be used in a particular context is non-inferentially given to the cognizer (Assumption 4). This assumption seems inconsistent with the view of critics of RD, such as Larkey and Markman (2005, p. 1071), but to our knowledge RD theorists themselves have not explicitly rejected this assumption. Our results show that Assumptions 2–4 suffice to render RD models tractable.

For the RD theorist feeling uncomfortable with Assumption 4, we note that the assumption could be dropped, but in that case it seems that one is forced to assume that the number of transformations in total (i.e., that could ever be used in any context) is not too large (Assumption 5). We see evidence for doubting this in the apparently unbounded ability of people to detect relationships between seemingly unrelated objects (French & Anselme, 1999). Also Hahn et al. (2003, p. 28) note that their results “point in the direction of the candidate set of cognitively salient transformations being at least rather large.” That being said, as long as Assumption 5 cannot be rejected with confidence, it remains a possibility to use the assumption in rendering RD models tractable.

In closing, we remark that our analyses are of course not exhaustive, in the sense that we only considered a small set of candidate sources of intractability (those listed in Table 2). There may exist other sources of intractability in RD models that we have not considered. With our analyses we hope to have illustrated how systematic intractability analyses can aid RD modelers, and cognitive modelers in general, in building computationally plausible models. As argued by Hahn et al. (2003, p. 28), transformational models of similarity ideally take into account:

“(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, and (III) any constraints on the ability of the cognitive system to discover simple transformations between mental representations.”

Notably, our analyses have contributed significantly towards (III), even without making any detailed assumptions about (I) and (II). If RD modelers find a way to introduce some motivated constraints on the nature of relevant representations and the nature of relevant transformations, many new options for rendering tractable RD models may become available.

References

- Chater, N., & Hahn, U. (1997). Representational distortion, similarity, and the universal law of generalization. In *Proceedings of the Interdisciplinary Workshop on Similarity and Categorization* (p. 31-36). Edinburgh: Department of Artificial Intelligence, University of Edinburgh.
- Chater, N., & Vitányi, P. (2003). The generalized universal law of generalization. *Journal of Mathematical Psychology*, *47*, 346-349.
- Downey, R. G., & Fellows, M. R. (1999). *Parameterized Complexity*. Berlin: Springer.
- Flum, J., & Grohe, M. (2006). *Parameterized Complexity Theory*. Berlin: Springer.
- French, R. M., & Anselme, P. (1999). Interactively converging on context-sensitive representations: A solution to the frame problem. *Revue Internationale de Philosophie*, *3*, 365-385.
- Frixione, M. (2001). Tractable competence. *Minds and Machines*, *11*, 379-397.
- Garey, M. R., & Johnson, D. S. (1979). *Computers and intractability: A guide to the theory of NP-completeness*. San Francisco, CA: W.H. Freeman.
- Hahn, U., Chater, N., & Richardson, L. B. (2003). Similarity as transformation. *Cognition*, *87*, 1-32.
- Imai, S. (1977). Pattern similarity and cognitive transformations. *Acta Psychologica*, *41*, 433-447.
- Larkey, L. B., & Markman, A. B. (2005). Processes of similarity judgment. *Cognitive Science*, *29*, 1061-1076.
- Pothos, E. M., & Chater, N. (2002). A simplicity principle in unsupervised human categorization. *Cognitive Science*, *26*, 303-343.
- Rosch, E., & Mervis, C. B. (1975). Family resemblances: Studies in the internal structure of categories. *Cognitive Psychology*, *7*, 573-605.
- Tsotsos, J. K. (1990). Analyzing vision at the complexity level. *Behavioral and Brain Sciences*, *13*, 423-469.
- Tversky, A. (1977). Features of similarity. *Psychological Review*, *84*, 327-352.
- van Rooij, I. (2008). The tractable cognition thesis. *Cognitive Science*, *32*, 939-984.
- van Rooij, I., Evans, P., Müller, M., Gedge, J., & Wareham, T. (2008). Identifying sources of intractability in cognitive models: An illustration using analogical structure mapping. In B. C. Love, K. McRae, & V. M. Sloutsky (Eds.), *Proceedings of the 30th Annual Meeting of the Cognitive Science Society* (p. 915-920). Austin, TX: Cognitive Science Society.
- van Rooij, I., Stege, U., & Kadlec, H. (2005). Sources of complexity in subset choice. *Journal of Mathematical Psychology*, *49*(2), 160-187.
- van Rooij, I., & Wareham, T. (2008). Parameterized complexity in cognitive modeling: Foundations, applications, and opportunities. *Computer Journal*, *51*, 385-404.