Context, structure, and informativeness judgments: An extensive empirical investigation

Ronaldo Vigo^{1,2} · Charles A. Doan^{2,3} · Basawaraj^{1,2} · Derek E. Zeigler^{2,4}

Published online: 8 July 2020
© The Psychonomic Society, Inc. 2020

Abstract

We explored the nature of human informativeness judgments: namely, people's judgments about the quantity of information that object stimuli convey about the category of objects to which they belong. Informativeness judgments play a key role in everyday decision-making situations involving the selection of items from groups that best represent the “group as a whole.” They also provide insight into the nature of prototype formation. We investigated informativeness judgments with an experiment involving 41 category structures – the most comprehensive and rigorous examination thus far. We assess the robustness and generalizability of the results from this experiment by examining the relationship between group-level and individual-level performance. In addition, we show that in most cases (and especially in those involving relatively lower dimensionality structures), these judgments are predicted more accurately and explained more satisfactorily by Representational Information Theory (Vigo in *Information Sciences* 181: 4847–4859, 2011 and in *Information* 4(1):1–30, 2012) and its simplest core model than by standard models of prototypicality. Finally, we argue that prototypicality models are special cases of the more general “representational information” framework.

Keywords Information · Representational information · Prototypicality · Informativeness judgments · Invariance · Relational cognition

Electronic supplementary material The online version of this article (<https://doi.org/10.3758/s13421-020-01053-1>) contains supplementary material, which is available to authorized users.

✉ Ronaldo Vigo
vigo@ohio.edu

✉ Charles A. Doan
cad007@marietta.edu

Basawaraj
bb593707@ohio.edu

Derek E. Zeigler
dzeigler3@csc.edu

- ¹ Department of Psychology, Ohio University, Athens, OH, USA
- ² Consortium for the Advancement of Cognitive Science, Ohio University, Athens, OH, USA
- ³ Department of Psychology, Marietta College, Marietta, OH 45750, USA
- ⁴ Department of Psychology and Education, Columbus State Community College, Columbus, OH, USA

Introduction

Information, as a universal construct, plays a fundamental role in contemporary science. In fact, the construct plays a central role in Data Science, Informatics, Bioinformatics, and Cognitive Science, as well as an important role in much older fields like Physics and Chemistry, to name a few. Yet, perhaps because of its multifaceted, fundamental, general, and ubiquitous nature, it has resisted a precise definition that completely conveys its meaning in a domain-independent fashion. Indeed, entire journal issues have been devoted to the enterprise of determining the best way of defining information (e.g., Logan, 2012). Unfortunately, unlike mass, energy, time, and any other of the fundamental physical quantities, there is no universally agreed upon definition and, consequently, no universal measure of information that is capable of characterizing its nature effectively across the wide variety of existing scientific domains. The well-known measure of information introduced by Shannon (1948), in which information is quantitatively characterized as the logarithm of the probability of a random variable, has been a popular approach to measuring information, but it also has proven to be inappropriate and ineffective in capturing the deeper structure of information (Luce, 2003; Vigo, 2012) as well as its more subtle aspects.

Attempting to resolve some of these issues, Vigo (2011) proposed an alternative way of construing and measuring information referred to as “representational information.” In a nutshell, representational information was defined as the discrete rate of change in categorical complexity whenever subsets of a category (where the term “category” refers to a set of objects related in some way) are removed from the category. This latter measure has shown promise in capturing psychological aspects of information (Doan & Vigo, 2016; Vigo & Basawaraj, 2013). Particularly, it has been applied to successfully predict people’s judgments with respect to the amount of information that an object carries about the set to which it belongs.

Cognitively, such judgments play a fundamental role in the way that humans interact with their environment to acquire knowledge efficiently. For example, people often encounter situations where the need arises to assess the most representative object among a set of objects, or, more directly, to assess how informative a subset of the elements of a given category reflect the information conveyed by the entire category. Henceforth, we refer to such subsets of category elements as “concept cues” and to the category under consideration as an “online concept” due to the constant availability of its associated instances *a priori*. For example, consider a political setting where potential voters may want to review a set of candidates in terms of the one that best represents the views of all the other candidates (perhaps as embodied by the party’s platform). On the other hand, consider a scientific setting where a biologist may find it useful to ascertain which genus may be most representative of a certain family of organisms. In these types of situations, representational information theory (RIT; Vigo 2011, 2012) provides a natural way of characterizing subjective information in a manner that captures the structural essence of information. Representational information has also proven useful in understanding and predicting how people sort objects in constrained and free-sorting tasks and in understanding how people build natural categories (Doan & Vigo, 2016). In fact, the evidence gathered thus far suggests that the ability to assess the amount of information that is conveyed by concept cues is grounded on core representational mechanisms and, as such, may explain the way categories are encoded as concepts in the first place.

Because of its fundamental nature, the ability to assess the degree of informativeness of concept cues can also explain why the idea of prototypicality has played an important intuitive role in concept learning research. More specifically, in the 1970s, concept learning researchers gathered evidence in support of the idea that people represent categorical information as prototypes, where a prototype is roughly defined as a summary single-object representation of the objects of a category. Of course, prototypes can be generated in different ways. Estes (1986), Reed (1972), and Rosch and Mervis (1975) proposed specific ways of characterizing prototypicality and prototypes. For instance, Rosch and Mervis proposed the family

resemblance principle to characterize how people perceive a set of items that are internally interrelated. The principle dictates that each object in a category must share at least one feature with one or more other objects in the category. From a category exhibiting family resemblance, people then generalize to a prototype representation by “assembling” the winning features across all objects in the category to form the prototypical object. In the next section we describe the notions of prototypicality introduced by Estes (1986) and Reed (1972).

Given that at least some forms of conceptual behavior can be accounted for by the existence of prototypes and that prototypes are a type of conceptual representation, a key question emerges regarding the relationship between representational information and prototypicality (Reed, 1972; Rosch & Mervis, 1975; Rosch, 1978): are informativeness judgments as construed by representational information theory consistent with choosing the most prototypical object from a category of objects as prescribed by Rosch and colleagues (Rosch & Mervis, 1975)? To answer this question, recall that an observer can form the central family tendency or prototype by simply taking the features that most members share and draw a composite using these features. This means that the winning features of the family are passed to the prototype, thereby providing a common link to the majority of elements in the category. Thus, once the prototype is generated, degrees of prototypicality could be the basis of informativeness. On the other hand, representative information generates such degrees of informativeness by a very different mechanism: one by which the degree of diagnosticity of each dimension of the objects of the category (and its corresponding values or features) must first be ascertained before computing the contextual contribution of each object to the complexity (or degree of learning difficulty) of the category of objects. Via this final computation, the degree of representativeness of each object, which includes their degree of prototypicality in a broader sense, can be computed. Thus, judgments about the degree of informativeness depend on more than an estimate of the most prototypical item(s) in a set as determined by the two prescriptions examined in this paper. In other words, notions of prototypicality do not serve as good measures of the amount of information carried by concept cues because: (1) they only carry a portion of the essential information that is available in the category and (2) they can only account for single-element concept cues.

As evidence, consider an investigation by Chin-Parker and Ross (2004) that examined how the nature of a learning task would emphasize prototypical versus diagnostic features. Particularly, subjects were required to perform one of two types of tasks: a classification task for which subjects had to predict category membership, and an inference task for which subjects had to predict missing features from category knowledge. This was followed by other tasks to test the role of diagnosticity and prototypicality on these studies. It was observed that classification learners relied primarily on the diagnostic features when

making decisions about category membership, but inference learners were sensitive to both the diagnosticity of the features (albeit significantly less so than the classification learners) and prototypicality (Chin-Parker & Ross, 2004). These results seem to corroborate the representational information approach because, as mentioned, the approach involves both an assessment of the degree of diagnosticity of the dimensions/features of the category and a more general assessment of what it means to be “most representative.” Thus, representational information can account for both types of learning – a point that we will argue for the remainder of this paper.

Now that we have established the need for assessing the most representative object among a set of objects or, more psychologically, the need to assess how informative concept cues are about their corresponding category, in the next section we explain in greater detail representational information. However, the mathematical details are discussed in [Appendix 3](#).

Representational information and notions of prototypicality

Several well-known models that use a measure of prototypicality at their core have been developed to account for categorization behavior in humans (Estes, 1986, 1994; Nosofsky & Zaki, 2002; Reed, 1972; Rosch & Mervis, 1975; Smith, Murray, & Minda 1997; Smith & Minda, 1998). Because our interest is on informativeness judgments, and not on categorization performance, we will focus on the methods used to measure degrees of prototypicality in these models. The first of these methods, which we refer to as the “feature-frequency” prototypicality measure (FFPM), was suggested by Rosch and associates (Rosch, 1978; Rosch & Mervis, 1975). The FFPM determines the degree of prototypicality of each object by taking the sum of the number of objects in the set with which it shares the same features. For example, with respect to a set of objects defined by the dimensions of color, shape, and size, the degree of prototypicality of any object in the set is given by summing three quantities: (1) the number of objects in the set that have the same color as the object in question, (2) the number of objects in the set that have its same shape, and (3) the number of objects in the set that have its same size. Using this approach, degrees of prototypicality (normalized) corresponding to the category structures used in the current experiment are shown in the “FFPM” column of [Tables 4 and 5](#) in [Appendix 1](#).

The second method for determining degrees of prototypicality comes from the multiplicative prototype model (MPM; Estes, 1986, 1994; Nosofsky & Zaki, 2002). The core idea underlying the MPM is that humans classify objects from a category by assessing the similarity of their corresponding exemplars to the category’s prototype. Thus, each similarity comparison between an object’s exemplar (i.e., mental representation of the object) and the prototype of the set of objects

determines the perceived degree of prototypicality of the object. Shepard’s similarity measure (Shepard, 1974, 1987) is used in the model to compute these degrees of prototypicality. The similarity between the particular object and the prototype is computed by taking the negative exponent of the sum of the distances (differences) between their features. The MPM comes in two flavors: the original MPM (Estes, 1986; denoted “MPMP” in [Tables 4 and 5](#) in [Appendix 1](#)) uses the average value across the dimensions of the objects in the category to form its prototype and the second version uses the more common feature-frequency notion of prototype explained above (Nosofsky & Zaki 2002; Rosch, 1978; Rosch & Mervis, 1975). Nosofsky and Zaki (2002) argued that the latter is the most effective among the well-known prototype models in accounting for classification data. Notwithstanding, because the feature-frequency version of the MPM computes the prototype of a category in exactly the same way as the FFPM, we compare the more commonly adopted FFPM to the MPMP (the “original” MPM) in terms of accounting for concept-cue judgments in the current experiment (discussed shortly).

Although there have been multiple studies on prototypicality judgments as the basis for categorization behavior (for a brief summary, see Nosofsky et al., 1994), no study has directly investigated the “degree of informativeness” conveyed by concept cues. By “degree of informativeness,” we mean a direct comparison on some subjective information scale as to how much information a stimulus-object conveys to the observer. One reason for this deficit may be attributed to the fact that the construct of information as proposed by Shannon (1948) does not capture our intuitions about what information means as a cognitive construct. Indeed, the numerous attempts to link this widely used method of measuring information to various psychological phenomena have been effectively thwarted by Devlin (1991), who points at the inadequacy of the Shannon and Weaver information measure in accounting for conceptual information. Luce (2003), on the other hand, recognized that the measure does not capture the structural nature of information. Both criticisms revolve around the seemingly daunting task of capturing the structural contextual effects that emerge from the perceived relationships between the instances of a category during the process of generalization or concept formation in humans. Many experiments in perception science (particularly by Gestalt researchers) and fewer in other areas of cognitive science have documented the undeniable existence of these contextual effects on response behavior (for examples, see Medin, Goldstone & Gentner, 1993; Reicher, 1969). For example, in one study by Medin et al. (1993), subjects were instructed to list common and distinctive properties among two objects from a set of three objects. They found that the properties ascribed to any one object depended on to which of the other two objects it was paired. They concluded that similarity judgments are heavily influenced by context.

Representational information theory or RIT (Vigo, 2011, 2012) offers a possible solution to this problem because it is based on a theory of concept learning known as “Generalized Invariance Structure Theory” or GIST (Vigo, 2013, 2016). GIST posits that observers are invariance pattern detectors. As such, they tune into particular types of invariant relationships between the object-stimuli of categories. These relationships take the fundamental form of simple abstract symmetries in the categorical stimulus and serve to extract the essential information of a category. Two functionally equivalent mechanisms for invariance detection have been proposed in the theory. Under the first mechanism (Vigo, 2013, 2016), invariance detection is the process of assessing the “partial similarity” between object stimuli in a category of objects by binding or suppressing, one at a time, each of their dimensions. This process is referred to in GIST as dimensional binding. The process not only involves rapid attention shifting from one dimension to another, but also involves sustained and deliberate “unattending” or “suppression” with respect to each dimension. Ultimately, the described process results in the detection of categorical invariants. These invariants are the category members that determine which features (and more generally, dimensions) of the categorical stimulus should be used in rule formation (see Vigo, 2009, 2011, 2013, 2016, for a detailed explanation). The proportion of categorical invariants of a categorical stimulus are then represented as an “ideotype” in psychological space (i.e., ideotypes contain the structural information necessary to form concepts). In other words, in order to form explicit language-oriented categorization rules, humans must first implicitly detect these invariance patterns because they make it possible for an observer to ascertain which features associated with a particular categorical stimulus best predict category membership (i.e., which dimensions are most diagnostic). If the diagnostic dimensions of the categorical stimulus are fully detected by the observer and amount to one or two, then the categorical stimulus is simple to learn, otherwise it is not. Under the second and equivalent mechanism (formally described in Appendix 3), observers engage in the process of perturbing or psychologically transforming the given category in such way that the dimensional values (i.e., features) of its objects are changed to opposite values in a range of possible values. Once this is done, observers compare the original category to the perturbed category and determine the common elements from the standpoint of the relevant dimensions. These common elements are the invariants of the category with respect to each of the dimensions. After, as in the first mechanism, observers form the ideotype representation described above. In short, in GIST, when the observer is a good detector of categorical invariants, he or she is a good concept learner and classifier; in turn, these patterns of invariance reduce stimulus complexity by revealing the key information necessary for correctly classifying the members of a categorical stimulus. Clearly, the process of

invariance pattern detection requires the analytic ability to dissect objects in terms of their dimensions in the first place, and not just the ability to shift attention rapidly while deliberately ignoring a particular dimension. A summary of GIST is given in Appendix 3. Here we present the basic tenets of Generalized Representational Information Theory as based on GIST:

- 1) Humans communicate via concepts or, in other words, mental representations of categories of objects (where a category is simply a set of objects that are related in some way)
- 2) therefore, concepts are the mediators of information
- 3) but concepts are relationships between qualitative objects in the environment that are defined dimensionally
- 4) the degree of structural homogeneity of a category (i.e., to what extent its objects are relationally indistinguishable or, equivalently, to what extent the objects are interconnected as determined by the invariance properties of the category) as well as its cardinality (i.e., size) determine the degree of learnability or complexity of its associated concept, and
- 5) information is the rate of change of that complexity.

The first three principles are frequently adopted by researchers in the field of human concept learning (Bourne, 1966; Estes, 1994; Garner, 1974; Vigo, 2009, 2013), while principles four and five form the basis of the theories proposed by Vigo (2009, 2011, 2012, 2013). For example, with respect to the fourth principle, Vigo (2013) developed a theory and model of degree of concept learning difficulty where the perceived degree of structural complexity of a set of objects (i.e., its learnability) is directly proportional to the cardinality or size of the set and inversely proportional to the exponent of its degree of invariance (where the degree of invariance of a category is proposed as a measure of its degree of structural homogeneity). The model, referred to as the “Generalized Invariance Structure Theory Model” (GISTM; Vigo, 2013), accurately accounts for the variance in historical data (e.g., Bourne, 1970; Shepard, Hovland & Jenkins, 1961) and from a recent large-scale experiment involving 84 category structures consisting of two, three, and four dimensions (Vigo, 2013) – a number well beyond the six tested by Shepard and associates.

Lastly, the fifth principle is the basis of Vigo’s representational information theory (Vigo, 2011). The basic idea underlying the principle is that what makes objects informative is the way that their inclusion/exclusion impacts the learnability or complexity of sets. The extent of their impact, whether positive (increase in complexity) or negative (decrease in complexity), determines the amount and quality of information they convey. More specifically, the amount of information conveyed by a subset R of a set of objects S about S is the rate of change in the perceived degree of structural complexity

in S whenever R is removed from S . In other words, the extent to which an object(s) (in extension) or concept cue (object in intension) is informative about its set of origin is determined by the impact that the removal of the object(s) has on the relative difference between the structural complexity of the original set and that of the new set (i.e., the original set minus the object(s)).

The higher this rate of change, the more information is conveyed. Thus, the notion of “the degree of surprise of an event” that is often associated with Shannon information is replaced by the notion of “the rate of change in the learnability of a concept” or, equivalently, the rate of change in the perceived degree of structural complexity of the category from which the concept is learned. Moreover, the quality of the information is represented by the direction of the change: negative rates of change indicate a decrease in perceived complexity, while positive rates indicate an increase in perceived complexity. The sign of the information acts as a moderating factor on the raw amount of information by indicating the difference between information perceived as not useful information (a decrease in the rate of change in complexity) versus useful information (an increase in the rate of change in complexity).

Accordingly, to compute the amount and quality of representational information conveyed by any particular single object about the stimulus set S that contains it, we take the difference between the structural complexity of S' (S without the object in question) and the structural complexity of S ; then, we divide the result by the structural complexity of S to get the percentage rate of change of the structural complexity in S with respect to the particular object. This rate of change, according to GRIT, is the amount of information conveyed by the object about S . The structural complexity of both S' and S is computed by the generalized invariance structure theory model (GISTM; Vigo, 2013, 2016), a model intended to account for the degree of learning difficulty of any dimensionally defined stimulus set (i.e., set of objects) in terms of an invariance-based notion of conceptual complexity.

In addition, note that the object removal process described above also has a corresponding cognitive mechanism based on selective attention and stimulus suppression. Specifically, GRIT posits that observers have the ability of psychologically “subtracting” an object from a category (as is the case with a single element concept cue) temporarily by selectively attending to all the members of the category except the object in question while at the same time suppressing the object in memory. It is this ability that makes it possible for subjects to determine the impact that the absence of any object has on the complexity of the category as a function of its coherence. Empirical investigations by Vigo, Zeigler, and Halsey (2013), Vigo (2016), and Rehder and Hoffman (2005) suggest that the object suppression mechanism in the theory is plausible. Specifically, these researchers have shown that gaze patterns in category learning tasks are often very complex and

systematic in character. For example, subjects often engage in many very rapid attention-shifts per second from feature to feature and from object to object, often ignoring specific features and objects in the category in ways that optimize classification decisions. We assume that such a suppression mechanism (entitled “dimensional binding” in GIST and GRIT) is implicit and subjects are not likely to be able to report verbally what has transpired.

Thus, the results of the coherence and complexity judgments of the category with and without the object are then computed and stored in working memory where the computation of the most informative object takes place per GRIT’s prescription. A sketch of the mathematical details of GRIT and the GISTM are given in Appendix 3 as well as a detailed example of how to compute information values. Representational information values (normalized) for each of the objects in each of the category structures tested are shown in the “GRIT-NP” column of the Tables 4 and 5 in Appendix 1.

Finally, the idea of linking the change in the complexity of a category structure (when some of its elements are removed) to information content addresses the problem of determining the role that context plays on the amount of information humans attribute to concept cues. Hence, we will argue that the most informative concept cues will be those whose role in the online category (i.e., in the set presented in its entirety to the subject) is to decrease its perceived complexity. In other words, the more informative a concept cue, the more its removal will increase the complexity of the original category (or equivalently, the less interrelated its members will be), and vice versa. Then, the difference between the new category and the original will be positive and the percent rate of change of the complexity will be positive. In support of this hypothesis are recent empirical studies by Doan and Vigo (2016) who demonstrated observers adding and removing multidimensional objects to and from categories that minimized the concept learnability as measured by GRIT.

The current experiment shows an additional and novel way of making this key connection empirically. We asked the participants to judge which single object out of a set of two, three, or four objects conveyed the most information about the other objects in the set (or about the set of objects as a whole). The sets were structurally different. The “most informative” object request was purposely used instead of the “most typical” object request because the purpose of our study was to elicit responses that are consistent with human intuitions as to the nature of the concept of information. Moreover, we felt that a “typical object” request would encourage subjects to indulge in the process of counting features. However, our aim was to elicit Gestalt judgments with respect to the relationships detected in each set of objects.

Before proceeding to a detailed description of the experiment, it should be noted that, in related work, Vigo and Basawaraj (2013) conducted similar studies on the nature of

informativeness judgments, but these were very limited in scope and rigor when compared to the current studies. For example, these researchers only examined six category structures compared to 41 in the current study. The methodology and aims were also very different. In their study, no attention was given to the relationship between group level and individual level performance – a key relationship in establishing the robustness of the results and their generalizability. Furthermore, the tested category instances varied from subject to subject unlike in our current study. Other differences, such as the nature of the statistical analysis, make the current research far more comprehensive and methodologically more rigorous to this previous early attempt at answering key questions about the nature of subjective information.

Method

Participants

We recruited a total of 105 participants from Introductory Psychology courses at Ohio University. Each participant received partial course credit for his or her participation. Participants were removed from statistical analysis if they were deemed uncooperative ($N = 9$). Specifically, these uncooperative participants were clicking through the experiment to finish early (evidenced by significantly faster response times) and were found to be texting during the experiment. We were able to identify 16 non-native English speakers among the remaining 96 participants and we report a dual-statistical analysis for the data with ($N = 96$) and without ($N = 80$) these participants given the likelihood that they may have not fully understood the instructions to the task.

Stimuli

A set of three HP XW4600 workstations, each equipped with a Dell 1708FP 15-in. flat panel LCD monitor (5-ms response time), were used to display categorical stimuli consisting of between two and four flasks (flat bottles; Vigo & Basawaraj, 2013) defined over either three or four binary-valued stimulus dimensions. The three-dimensional flask stimuli were defined over the dimensions of *shape* (circular or triangular), *color* (white or black), and *size* (large or small). In addition to being defined over these same binary-valued stimulus dimensions, the four-dimensional flask stimuli were also defined by a dimension of *neck width* (narrow or wide).

We use the $D_n[p]$ notation (Feldman, 2000, 2003; Vigo, 2013) to distinguish the dimensionality (D), number of dimensional values (n), and number of objects (p) characteristic of each structure type within the associated structure family. In the current experiment, each categorical stimulus conformed to one of the 41 unique structure configurations associated

with the $3_2[2]$, $3_2[3]$, $3_2[4]$, $4_2[2]$, $4_2[3]$, and $4_2[4]$ structure families (see Feldman, 2003, for visual depictions of these 41 structure types). By “unique structure configuration” we simply mean that the specific relationships between the dimensional values of the objects belonging to one of the 41 structure types are different from the relationships between the dimensional values of the objects belonging to any of the other 40 structure types. For example, each of the six structure types of the $3_2[4]$ structure family consist of four objects defined over three binary-valued stimulus dimensions. This family of structure types has been extensively studied with respect to the difficulty in correctly classifying members of each unique structure type (Feldman, 2000, 2006; Kruschke, 1992; Kurtz et al., 2012; Nosofsky, 1984, 1986; Nosofsky et al., 1994; Rehder & Hoffman, 2005; Shepard, Hovland, & Jenkins, 1961; Smith, Minda, & Washburn, 2004; Vigo, 2013, 2016; Vigo, Evans, & Owens, 2014).

Although the $3_2[4]$ structure family has been widely studied with respect to human and non-human classification behavior, successful theories and computational/mathematical models of concept learning behavior more importantly account for patterns of categorization behavior within and across families of category structure types (Feldman, 2000; Vigo, 2013, 2016). Currently, we take a similar approach by positing that theories of informativeness should also be able to account for patterns in the data for multiple families of structure types, but with regards to the amount of information conveyance of individual concept cues within and across these same families of structure types.

Finally, if we systematically remap the dimensional values constituting the set of objects, we can produce instances for each structure type that are distinct but structurally equivalent. For example, one instance of the $3_2[4]$ – I structure type may consist of the four small flasks. However, if we remap the dimensions such that white replaces small, then we obtain another instance of a $3_2[4]$ – I structure type consisting of the four white flasks. Exhausting this procedure for the remaining structure types of the $3_2[4]$ family reveals that there are a total of 6, 6, 24, 8, 24, and 2 instances associated with these six structure types, respectively. And although we have only demonstrated this procedure for the $3_2[4]$ structure family, the exhaustive procedure can also be applied to structure types belonging to the remaining five structure families (35 structure types) tested in the current study.

Procedure

Participants were assigned to one of two databases of category structures with their associated randomly selected instances. One group of participants made decisions on instances from a random selection of 20 structure types (9 instances $\times$ 20 types = 180 decisions), whereas another group of participants made decisions on instances from the remaining 21 structure types

Table 1 Simple linear regression analyses per structure family

Structure family	Rotated data?	Data (w/non-native)			Data (w/o non-native)		
		GRIT-NP	FFPM	MPMP	GRIT-NP	FFPM	MPMP
3 ₂ [3]	No	.40 (.124)	.00 (.16)	.01 (.159)	.46 (.143)	.01 (.194)	.02 (.193)
	Yes	.41 (.121)	.00 (.157)	.01 (.156)	.47 (.140)	.01 (.192)	.02 (.191)
3 ₂ [4]	No	.52 (.056)	.40 (.062)	.04 (.079)	.61 (.064)	.55 (.068)	.06 (.099)
	Yes	.83 (.026)	.64 (.038)	.14 (.059)	.84 (.035)	.76 (.043)	.16 (.08)
4 ₂ [3]	No	.37 (.113)	.27 (.121)	.40 (.11)	.37 (.141)	.29 (.15)	.43 (.135)
	Yes	.39 (.107)	.32 (.113)	.46 (.101)	.39 (.135)	.35 (.141)	.49 (.124)
4 ₂ [4]	No	.15 (.052)	.20 (.05)	.13 (.053)	.34 (.061)	.44 (.056)	.23 (.065)
	Yes	.30 (.034)	.38 (.032)	.17 (.037)	.48 (.045)	.59 (.04)	.23 (.055)
All 3D	No	.63 (.08)	.49 (.094)	.12 (.123)	.67 (.085)	.46 (.109)	.12 (.139)
	Yes	.71 (.067)	.55 (.083)	.15 (.114)	.74 (.072)	.51 (.099)	.15 (.13)
All 4D	No	.54 (.07)	.58 (.067)	.15 (.098)	.56 (.079)	.57 (.079)	.21 (.107)
	Yes	.63 (.059)	.69 (.054)	.21 (.086)	.64 (.068)	.64 (.068)	.25 (.098)

Note. R^2 values provided with accompanying *SEE* values in parentheses. Bolded values represent the “best” fit for that row when considering both R^2 and *SEE* values. Fits are not provided for the two-object structure families, as both the GRIT-NP and FFPM make identical predictions for all objects belonging to these structure types

(9 instances $\times$ 21 types = 189 decisions). Upon giving consent, we instructed participants “to select the one flask from each set that you believe conveys or carries the most information about the flasks in that set.” Once any questions were answered, participants were then guided to a room where they completed a four-trial training session and the 180-trial or 189-trial main task on the computer. Each participant made their selection for each presented instance by clicking on a button, labeled “1,” “2,” “3,” or “4,” situated below the presentation of each structure instance. Note that button “1” corresponds to the leftmost object for each of the structure instances, regardless of number of objects presented. However, for the two-object structure instances (e.g., 3₂[2], 4₂[2]), button “2” corresponds to the rightmost object. Conversely, for the three-object structure instances (e.g., 3₂[3], 4₂[3]), button “2” corresponds to the middle object, and button “3” corresponds to the rightmost object. Finally, for the four-object structure instances (e.g., 3₂[4], 4₂[4]), buttons “2” and “3” correspond to the middle two objects, and button “4” corresponds to the rightmost object. Although each subject within a group made decisions on the same structure type instances, these instances and the serial positioning of the objects per instance were presented in a random order per participant.

Results

As mentioned previously, we conducted a dual statistical analysis for the data that included the non-native English-speaking participants ($N_{\text{Group1}} = 44$; $N_{\text{Group2}} = 52$) and for the data that did not include these participants ($N_{\text{Group1}} = 38$; $N_{\text{Group2}} = 42$).

To determine the role structural context has on human information judgments, we calculated the proportion of times each flask was chosen as conveying the most information about the category on a per structure type basis. Because there were (1) multiple instances tested per structure type and (2) the spatial arrangement (e.g., serial position) of objects per instance were randomized per trial, we first had to align the objects per instance for each participant such that we could aggregate proportions across structurally equivalent objects. That is, we ensured logical configurations of objects within particular instances of each of the 41 structure types matched before aggregating their associated data (see [Appendix 2](#) for an accompanying description and example of the aligning process). Next, we conducted statistical analyses on: (1) the aggregate data per structure type, (2) the same data after 40 rotations were performed, and (3) the individual level per structure type. The goal of the rotations is to reduce the effect particular confounding variables, such as differentially salient stimulus dimensions, that may be introduced into the judgment process. Importantly, the rotations do not favor increased predictive capability of any of the tested models; they were merely performed to reduce confounding influences affecting how observers perceive the underlying structure of each categorical stimulus before making their decisions. That being said, it is reasonable to conclude that models purporting to account for the most informative concept cues across different contexts should fare better in accounting for such behavior on rotated sets of data. Both the non-rotated and rotated proportions for each object belonging to the 41 structure types are provided in the column “Concept function” of [Tables 4](#) and [5](#) in [Appendix 1](#). Similarly, model predictions for each structure

Fig. 1 Simple linear regression analyses across the six structure families (41 structure types) for the non-rotated data with Non-native English-speaking participants included. Each row represents fits to one of the three

type (normalized per structure type) are presented in their respective columns of Tables 4 and 5 in Appendix 1.

Regression analysis per structure family We performed simple linear regression analyses to determine how well each model accounts for the non-rotated and rotated proportions across particular structure families and across the 41 structure types. Table 1 provides these regression fits per structure family

tested models, displaying both the fitted data (Column 1) and the associated distribution of residuals (Column 2)

(e.g., $3_2[4]$) for each of the three models assessed in the current study and for both sets of data. We provide Table 1 to help delineate among the better performing model with respect to particular structure families.

Regarding particular structure families, the pattern of results observed for the $3_2[4]$ structure family is particularly worth discussing further as this family has been widely studied to ascertain the nature of categorization behavior in

humans and non-human animals (Nosofsky et al., 1994; Rehder & Hoffman, 2005; Shepard et al., 1961; Vigo, 2013; Vigo, Zeigler, & Halsey, 2013; Vigo, Evans, & Owens, 2014). Typically, degrees of learning difficulty for each of these six structure types is operationalized as the average proportion of errors in sequentially classifying the members belonging to the category. With regard to the six distinct structure types of the $3_2[4]$ family, the highly replicated learning difficulty ordering in terms of average proportion of errors is $I < II < [III, IV, V] < VI$ (Nosofsky et al., 1994; Rehder & Hoffman, 2005; Shepard et al., 1961; Vigo, 2013, 2016; Vigo, Evans, & Owens, 2014; Zeigler & Vigo, 2018; but see Kurtz et al., 2012). The question the current research addresses is how individuals perceive the informativeness of each of the objects (e.g., concept cues) comprising each of these distinct structure types.

With respect to informativeness, Vigo and Basawaraj (2013) conducted the first study inquiring which objects

humans find most informative within each of the six structure types belonging to the $3_2[4]$ family. With respect to types I, II, and VI, they found that subjects selected each of the four objects equally often – indicating that each object conveys the same amount of information about the category as each other object in the category. However, they observed differences in proportions for types III, IV, and V. Specifically, two objects and one object were selected more often and judged as more informative for type III and type IV, respectively. The pattern of results for type V revealed an ordering whereby one object was chosen most often, followed by two objects being chosen equally as often, followed by the last object being chosen the least. Inspecting the data in Table 4 in Appendix 1 reveals the same pattern of equal-object selection for the four objects belonging to types I, II, and VI. Furthermore, it is worth noting that the same two objects and one object for types III and IV were chosen as most informative by participants in the current study. The results for type V are not as

Fig. 2 Shown are stacked bar plots of the data (two leftmost columns in each plot) and of the normalized model predictions for the three-dimensional structure types tested in the current study. Within each plot, each color segment of a particular bar represents one of the objects for that structure type

Fig. 3 Shown are stacked bar plots of the data (two leftmost columns in each plot) and of the normalized model predictions for a subset of the four-dimensional structure types tested in the current study. Within each plot,

each color segment of a particular bar represents one of the objects for that structure type

unambiguous, but the ordinal trend across the four objects matches that previously found by Vigo and Basawaraj (2013). Thus, we believe the current experiment replicates the previous findings for the $3_2[4]$ structure family. Now, which of the three models best accounts for the informative data across these six structure types? Inspecting Table 1, we find that the GRIT-NP provides the best fit, $b = 0.83$, $SE = .077$, $p < .001$, $R^2 = .84$, $RSE = .035$. The FFPM does provide a relatively good fit and, ordinally, only differs from the GRIT-NP in prediction for type V, $b = 1.77$, $SE = .215$, $p < .001$, $R^2 = .76$, $RSE = .043$. The MPMP does not seem to fit the data well – only accounting for the pattern of selections for types IV and V, $b = 0.17$, $SE = .082$, $p = .053$, $R^2 = .16$, $RSE = .08$.

Regression analysis across structure families Regarding the regression analyses across all 41 structure types, Fig. 1 shows similar fits between the GRIT-NP and FFPM when the non-native English-speaking participants are included (Column 1).

GRIT accounts for 58% of the variance, while the FFPM accounts for 55% of the variance. However, similar fits without these participants show a widening of this difference – with the GRIT-NP accounting for more of the variation (60%) compared to the FFPM (53%) or MPMP (18%). Also, supporting this interpretation in favor of GRIT is the more favorable distribution of residuals compared to the other two models (see Column 2 of Fig. 1), as revealed by the non-significant Shapiro-Wilk normality test, $W = .985$, $p = .13$ (i.e., no significant departure from normality was found). Clearly, this result lends additional support to the validity and quality of the GRIT-NP predictions when compared to the significant normality tests for the FFPM, $W = .959$, $p < .001$, and MPMP, $W = .913$, $p < .001$.

Importantly, additional analyses reveal that each model does better accounting for the concept-cue data after rotations are applied (GRIT-NP: $b = 0.83$, $SE = .049$, $p < .001$, $R^2 = .68$, $RSE = .069$; FFPM: $b = 1.05$, $SE = .073$, $p < .001$, $R^2 = .59$, $RSE = .077$; MPMP: $b = 0.22$, $SE = .035$, $p < .001$, $R^2 = .22$,

Fig. 4 Shown are stacked bar plots of the data (two leftmost columns in each plot) and of the normalized model predictions for a subset of the four-dimensional structure types tested in the current study. Within each plot,

each color segment of a particular bar represents one of the objects for that structure type

$RSE = .107$. This increase in predictive accuracy is important considering that the goal of rotations is to reduce confounding effects individual stimulus dimensions may introduce into the judgment process independent of the structural effects existing between structure types and structure families. Although including the data from the non-native English-speaking participants negatively affects the performance of each model (i.e., lower R^2 values) with respect to the rotated data, the GRIT-NP ($R^2 = .66$, $RSE = .061$) still slightly outperforms the FFPM ($R^2 = .64$, $RSE = .063$).

Cross-validation We performed an extensive cross-validation by bootstrapping 1,000 samples of 60 objects and their associated proportions of being “most informative” (randomly selected across the 41 structure types) for each of the three models. Consistent with the results thus far, the GRIT-NP provided the most stable fits to all four sets of data

(Unrotated/Rotated $\times$ with/without Non-native), with the most stable fit (averaged statistics, $\bar{R}^2$ and $\overline{RSE}$, across each set of 1,000 samples) for the rotated set of data without the non-native English-speaking participants, $\bar{R}^2 = .68$, $SD = .092$; $\overline{RSE} = .068$, $SD = .01$. The most stable fit for the FFPM involves the rotated set of data with the non-native English-speaking participants, $\bar{R}^2 = .64$, $SD = .101$; $\overline{RSE} = .061$, $SD = .01$. Lastly, the most stable fit for the MPMP involves the same data set as that for GRIT, $\bar{R}^2 = .22$, $SD = .087$; $\overline{RSE} = .104$, $SD = .009$. In addition to these particular analyses, we also calculated the average Shapiro-Wilk statistic ($\bar{W}$) and p -value ($\bar{p}$) for each of the three models across the four datasets (Unrotated/Rotated $\times$ with/without Non-native; $N = 4,000$). These values from fits using the GRIT-NP ($\bar{W} = .915$, $\bar{p} = .15$) were higher than either the FFPM ($\bar{W} = .894$, $\bar{p} = .06$) or the MPMP ($\bar{W} = .899$, $\bar{p} = .006$). Importantly, this

pattern of values across the three models matches the pattern observed with the model fits provided in Fig. 1, indicating further stability for the fits across all the tested models.

Finally, Figs. 2, 3, and 4 display stacked bar plots for each of the three-dimensional ($N = 12$) and four-dimensional ($N = 29$) structure types tested in the current study for the data set including the non-native English-speaking participants. Within each bar of a plot, the bigger the occupied area by a certain color the more informative that particular object was judged to be (data) or predicted to be (model). The leftmost column in each plot displays the proportion of times each object (e.g., flask) was selected as conveying or carrying the most information about the other flasks in the set (or about the set as a whole). The adjacent column to the right in each plot displays the data after performing the 40 rotations per structure type. The three remaining columns depict the normalized model predictions of the GRIT-NP, FFPM, and MPMP, respectively. Altogether, the regression, correlational, and bootstrapping analyses suggest the GRIT-NP outperforms the other two models in accounting for the most informative object data within and across structure families and dimensionality.

Individual-level results

In addition to conducting regression and bootstrapping analyses on the aggregate data, we also conducted a comprehensive individual-level analysis across all 41 structure types with the 40 random rotations applied per type per subject. This analysis involved utilizing a simple method for determining the similarity between pairs of stacked bar plots (Vigo, 2019) using Shepard's similarity measure (Shepard, 1987). More specifically, we used this method to compare the normalized stacked bar plots of each model to those from each participant on a per structure type basis. The results of this comparison for each of the three models are provided in Table 2 and represent the average similarity between each model's predictions to each individual's data ($N = 1,968$ stacked bar plot comparisons).

Table 2 Results of stacked bar plot comparisons, in terms of average similarity, across all individuals and 41 structure types ($N = 1,968$ comparisons)

	Normalized predictions		
	GRIT-NP	FFPM	MPMP
Average	.701	.567	.531
SD	.16	.165	.128

Note. These similarity values are based on the Euclidean metric (i.e., $r = 2$) and no rescaling (i.e., $c = 1$). The similarity values can range between 0 and 1, with values closer to 1 indicating higher levels of similarity between the stacked bar plots (see Vigo, 2019, for details)

These comparisons prove to be particularly useful as they reveal differences between the GRIT-NP, FFPM, and MPMP that may be hard to ascertain via a visual comparison alone. Regarding the results, it appears that the GRIT-NP provides a better account of the individual-level data in terms of average similarity ($M = .701$, $SD = .16$) than the other two models. Importantly, these results are consistent with the aforementioned regression and bootstrapping analyses on the aggregate data.

Discussion

A central hypothesis of the current study is that the structural contexts in which objects appear differentially affect how individuals perceive how much information they convey about their category of origin. We empirically tested this hypothesis by conducting an extensive experiment on 41 unique category structures. To account for the results, we compared model performance between two “prototypicality” measures, both of which have been used as core components in formal models aiming to predict categorization behavior in human and non-human animals (Estes, 1986; Nosofsky & Zaki, 2002; Smith et al., 2004), and the representational information measure from Generalized Representational Information Theory (GRIT; Vigo, 2011, 2012, 2013). Generally, the prototypicality measures assume that single-element concept cues that are more similar to the category prototype will be perceived as conveying more information about the category than those that are less similar to the category prototype. Empirically, this should result in the former concept cues being selected more often in a forced choice task. On the other hand, in GRIT, representational information is defined as the percent rate of change in the conceptual complexity of the category. In simple terms, if removing a concept cue from a category results in a harder category to learn conceptually, then that concept cue conveys more information about the category than a similar concept cue that results in an easier category to learn once it is removed from the category. This makes sense because the degree to which removing the object is disruptive (in terms of percent rate of change) to the structural coherence of the category indicates how the object is interconnected to other objects in the category. This interconnectivity signifies that the object shares key structural properties with the other objects in the category.

But how well does each model predict the selection behavior of participants in the current experiment? That is, can each model accurately account for the perceived degree of informativeness of each concept cue about its category of origin? Before we describe and interpret model fits, it is helpful to assess how related the three models are to each other with respect to predictions at the family and dimensional levels of analysis. As shown in Table 3, there is a moderate degree of

Table 3 Correlation matrix comparing the normalized model predictions

Structure family	Normalized predictions		
		FFPM	MPMP
3 ₂ [3]	GRIT-NP	.77	.81
	FFPM		.94
3 ₂ [4]	GRIT-NP	.85	.52
	FFPM		.45
4 ₂ [3]	GRIT-NP	.64	.60
	FFPM		.93
4 ₂ [4]	GRIT-NP	.71	.51
	FFPM		.48
All 3D	GRIT-NP	.80	.60
	FFPM		.59
All 4D	GRIT-NP	.72	.50
	FFPM		.54
3D and 4D	GRIT-NP	.83	.50
	FFPM		.49

overlap among the three models in their predictions. One legitimate comparative strategy is to assess model performance for particular category structures with differing predictions. Alternatively, one may report the family and dimensional level fits, as these fits are likely more stable in that they aggregate results across different category structures. We report the latter analysis, while acknowledging a fruitful future research endeavor may be to focus on the structures that meaningfully differentiate among the tested models.

First, each of the three models differentially accounts for the forced-choice behavior when considering structure families in isolation, with the GRIT-NP accounting best for the selection of concept cues for the 3₂[3] and 3₂[4] families, the FFPM accounting best for the 4₂[4] family, and the MPMP accounting best for the 4₂[3] family. In explaining this pattern of results, it is important to consider the increased processing demands for the human observer with this corresponding increase in dimensionality and cardinality of the categories. When the number of dimensions and/or the number of objects are beyond the scope of the computational resources needed for more accurate informativeness judgments, it may be the case that observers resort to computationally simpler strategies such as counting features. This particular simpler strategy, and possibly others, aligns more closely to the representational and processing assumptions of the two prototype models assessed in the current study. However, it is also important to note that a greater number of dimensions and objects do not necessarily imply greater processing difficulty. For example, there are categories consisting of many dimensions and many objects that are relatively simple to process (take a category consisting of seven dimensions and

ten objects, each with the same value for one of its seven dimensions). This category is structurally similar to a type I structure in the 3₂[4] family, consisting of one fully diagnostic dimension, while its structural complexity is relatively low. Thus, a better criterion for determining computational demands would be its degree of internal coherence. In other words, the degree of invariance (or coherence) of the category may play a major role in determining which strategy observers would resort to in assessing informativeness. For example, it may be that categories that are in the extreme values of a U function/curve (i.e., high in coherence and very low in coherence as measured by their degree of invariance) may be more susceptible to feature counting. This tradeoff between computational expense and accuracy is an intriguing prospect indeed and should be rigorously and specifically tested in follow-up work. Relevantly, recent empirical work on one-shot categorization behavior provides evidence that such a tradeoff may be affected by the variability in experienced features and/or the sample size of experienced exemplars (Morgenstern, Schmidt, & Fleming, 2019).

In addition to this analysis across families of structure types, we can compare how well each model accounts for the selection of concept cues across the classic “SHJ” category types that consist of four objects each defined over 3-binary stimulus dimensions (3₂[4]). As detailed in the *Results* section above, the results of the current experiment generally replicate the results of Vigo and Basawaraj (2013), thus providing evidence for the stability of the findings for these six structure types. Only the GRIT-NP and FFPM account for the results on structure types I, II, III, and VI. Similarly, each of these three models makes identical predictions at the ordinal level for structure type IV. Importantly, the GRIT-NP and FFPM may be differentiated with respect to the predictions for type V, where the chosen concept cues seem to align more closely with the predictions of the GRIT-NP and MPMP. Given this evidence and considering model performance across all six structure types (see Table 1), it seems the GRIT-NP accounts for the concept cue results for this structure family more accurately than either of the other two models.

Upon aggregating proportions across structure families, we can determine how well each model accounts for concept-cue behavior dimensionally and across levels of dimensionality. Regarding dimensionality, it appears the GRIT-NP accounts best for three-dimensional concept cue behavior and both the GRIT-NP and FFPM account best for four-dimensional concept cue behavior. Across levels of dimensionality, the plots in Fig. 1 show that the GRIT-NP provides the best fit when aggregating performance across all 6 structure families (i.e., 41 structure types). This aggregate analysis includes how well each model accounts for behavior across the two-object structure families, a major shortcoming of the MPMP, as evidenced by comparing the stacked bar plots for the 3₂[2] and 4₂[2] families in Figs. 2 and 3, respectively.

Lastly, effective models of classification behavior should be able to also account for the most informative concept-cue data of the current experiment and also for concept-cue data for more complex categories than those currently tested, such as those with members defined over continuous dimensions. This ability would support the generality, comprehensiveness, and parsimony of the model.¹ Among the models tested, the GRIT-NP – and the GISTM, which is at the core of the GRIT-NP measure (Vigo, 2013, 2016) – meets this important standard with respect to binary-valued stimulus dimensions and is a promising candidate with respect to more complex dimensional configurations. In fact, both GRIT and GIST (the theory from which GRIT is derived) are capable of making predictions for stimuli defined over continuous dimensions. Moreover, the major difference between GRIT and its precursor (representational information theory or RIT) is this ability. Thus, one may easily compute representational information values for concept cues of categories consisting of continuous dimensions with the GRIT framework.

In contrast, Nosofsky and Zaki (2002) report results showing that a categorization model implementing the FFPM prototype method in the MPM outperforms a similar model implementing the MPM with the average prototype calculation currently considered. However, the FFPM does not perform as well as the GRIT-NP in accounting for the present

informativeness judgment results. Additionally, it is not clear presently how one would compute model predictions with either the FFPM or MPMP for continuous stimulus dimensions. For example, the FFPM procedure of counting features is computationally clear with respect to discrete features but not for continuous features for which there may be different degrees across the objects in the category. We suggest this lack of generalizability severely limits the applicability of these prototype models, a particularly notable limitation when one considers the wide range of categories that can be studied and assessed with continuous dimensions. Taken together, we contend that the results of the current extensive experiment align more parsimoniously with the generalized representational information theory framework while introducing impediments for measures based on the discussed notions of prototypicality.

Open Practices Statement The data for the reported experiment are available as Electronic Supplementary Material, and the experiment was not preregistered.

Acknowledgements The authors would like to thank Jay Wimsatt, Mikayla Doan, and the anonymous reviewers for their helpful comments.

Appendix 1

Table 4 Proportions and normalized model predictions (per structure type) for each concept cue of the tested three-dimensional category instances

$D_n[p]$	Structure type	Concept function (rotated proportions) (non-rotated proportions)	GRIT-NP	FFPM	MPMP
$3_2[2]$	1	$xyz + xyz'$ (.511, .489) (.543, .457)	{.50, .50}	{.50, .50}	{.74, .26}
	2	$xyz + xy'z'$ (.521, .479) (.545, .455)	{.50, .50}	{.50, .50}	{.89, .11}
	3	$xyz + x'y'z'$ (.497, .503) (.477, .523)	{.50, .50}	{.50, .50}	{.97, .03}
$3_2[3]$	1	$xyz + xyz' + xy'z$ (.614, .196, .189)	{.70, .15, .15}	{.40, .30, .30}	{.58, .21, .21}

¹ Additionally, it is important to note that all three models tested did not contain free parameters and that the variant of the GRIT core model tested is the simplest of the theory because it does not utilize the ability of the more complex and more powerful variants to specify levels of discriminability between the object-stimuli of the categories using what is referred to in GRIT and GIST as the tau-threshold bound-parameter (for details, see Vigo, 2012, 2016).

Table 4 (continued)

$D_n[p]$	Structure type	<i>Concept function</i> (rotated proportions) (non-rotated proportions)	GRIT-NP	FFPM	MPMP
$3_2[4]$	2	(.614, .146, .239) $xyz + xyz' + x'y'z$ (.138, .636, .226)	{.48, .48, .03}	{.50, .33, .17}	{.68, .24, .08}
	3	(.138, .636, .226) $xyz + xy'z' + x'yz'$ (.339, .337, .324) (.303, .351, .346)	{.33, .33, .33}	{.33, .33, .33}	{.33, .33, .33}
	1	$xyz + xyz' + xy'z + xy'z'$ (.252, .236, .246, .265) (.19, .157, .314, .339)	{.25, .25, .25, .25}	{.25, .25, .25, .25}	{.54, .20, .20, .07}
	2	$xyz + xyz' + x'y'z + x'y'z'$ (.254, .266, .237, .244) (.278, .297, .242, .183)	{.25, .25, .25, .25}	{.25, .25, .25, .25}	{.57, .21, .21, .02}
	3	$xyz + xyz' + xy'z + x'yz'$ (.346, .374, .141, .138) (.374, .347, .156, .124)	{.34, .34, .16, .16}	{.31, .31, .19, .19}	{.54, .20, .20, .07}
	4	$xyz + xyz' + xy'z + x'yz$ (.531, .141, .17, .159) (.531, .072, .136, .261)	{.54, .15, .15, .15}	{.33, .22, .22, .22}	{.48, .17, .17, .17}
	5	$xyz + xyz' + xy'z + x'y'z'$ (.294, .255, .288, .162) (.294, .239, .305, .162)	{.44, .24, .24, .09}	{.29, .29, .29, .14}	{.56, .20, .20, .03}
	6	$xyz + xy'z' + x'yz' + x'y'z$ (.238, .238, .257, .267) (.165, .277, .248, .31)	{.25, .25, .25, .25}	{.25, .25, .25, .25}	{.74, .09, .09, .09}

Table 5 Proportions and normalized model predictions (per structure type) for each concept cue of the tested 4-dimensional category instances

$D_n[p]$	Structure type	<i>Concept function</i> (rotated proportions) (non-rotated proportions)	GRIT-NP	FFPM	MPMP
$4_2[2]$	1	$wxyz + wxyz'$ (.503, .497) (.441, .559)	{.50, .50}	{.50, .50}	{.74, .26}
	2	$wxyz + wxy'z'$ (.507, .493) (.416, .584)	{.50, .50}	{.50, .50}	{.89, .11}
	3	$wxyz + wx'y'z'$ (.499, .501) (.488, .512)	{.50, .50}	{.50, .50}	{.97, .03}
	4	$wxyz + w'x'y'z'$ (.507, .493) (.479, .521)	{.50, .50}	{.50, .50}	{.98, .02}
$4_2[3]$	1	$wxyz + wxyz' + wxy'z$ (.645, .182, .174) (.645, .154, .202)	{.63, .18, .18}	{.37, .32, .32}	{.58, .21, .21}
	2	$wxyz + wxyz' + wx'y'z$ (.628, .229, .144) (.628, .229, .144)	{.47, .47, .05}	{.40, .33, .27}	{.68, .24, .08}
	3	$wxyz + wxyz' + w'x'y'z$ (.249, .528, .223) (.249, .528, .223)	{.47, .47, .05}	{.45, .36, .18}	{.72, .26, .02}
	4	$wxyz + wxy'z' + wx'yz'$ (.326, .321, .353) (.26, .36, .38)	{.33, .33, .33}	{.33, .33, .33}	{.33, .33, .33}
	5	$wxyz + wxy'z' + w'x'yz$	{.33, .33, .33}	{.45, .27, .27}	{.81, .10, .10}

Table 5 (continued)

$D_n[p]$	Structure type	Concept function (rotated proportions) (non-rotated proportions)	GRIT-NP	FFPM	MPMP
$4_2[4]$	6	(.641, .186, .173) (.641, .141, .218) $wxyz + wxy'z' + w'x'yz'$ (.282, .345, .373) (.282, .276, .441)	{.33, .33, .33}	{.36, .36, .27}	{.43, .43, .14}
	1	$wxyz + wxyz' + wxy'z + wxy'z'$ (.248, .252, .253, .247) (.241, .338, .214, .206)	{.25, .25, .25, .25}	{.25, .25, .25, .25}	{.54, .20, .20, .07}
	2	$wxyz + wxyz' + wxy'z + wx'yz$ (.429, .167, .20, .203) (.429, .168, .238, .165)	{.38, .21, .21, .21}	{.29, .24, .24, .24}	{.48, .17, .17, .17}
	3	$wxyz + wxyz' + wxy'z + wx'yz'$ (.31, .301, .201, .188) (.313, .298, .26, .13)	{.32, .32, .18, .18}	{.28, .28, .22, .22}	{.54, .20, .20, .07}
	4	$wxyz + wxyz' + wxy'z + wx'y'z'$ (.283, .285, .279, .152) (.283, .291, .273, .152)	{.39, .24, .24, .13}	{.27, .27, .27, .20}	{.57, .21, .21, .02}
	5	$wxyz + wxyz' + wxy'z + w'x'yz$ (.472, .178, .174, .176) (.472, .188, .164, .176)	{.39, .24, .24, .13}	{.32, .25, .25, .18}	{.54, .20, .20, .37}
	6	$wxyz + wxyz' + wxy'z + w'x'yz'$ (.288, .317, .173, .221) (.288, .317, .173, .221)	{.39, .24, .24, .13}	{.31, .31, .23, .15}	{.57, .21, .21, .02}
	7	$wxyz + wxyz' + wxy'z + w'x'y'z'$ (.383, .248, .257, .111) (.383, .251, .254, .111)	{.39, .24, .24, .13}	{.29, .29, .29, .13}	{.57, .21, .21, .01}
	8	$wxyz + wxyz' + wx'y'z + wx'y'z'$ (.254, .247, .248, .251) (.242, .301, .237, .22)	{.25, .25, .25, .25}	{.25, .25, .25, .25}	{.66, .24, .08, .02}
	9	$wxyz + wxyz' + wx'y'z + w'xy'z$ (.389, .195, .213, .203) (.389, .195, .245, .171)	{.32, .32, .18, .18}	{.31, .23, .23, .23}	{.63, .23, .08, .08}
	10	$wxyz + wxyz' + wx'y'z + w'xy'z'$ (.28, .281, .216, .223) (.313, .248, .251, .188)	{.32, .32, .18, .18}	{.29, .29, .21, .21}	{.66, .24, .08, .02}
	11	$wxyz + wxyz' + wx'y'z + w'x'y'z$ (.407, .187, .303, .104) (.407, .187, .303, .104)	{.25, .25, .25, .25}	{.29, .21, .29, .21}	{.66, .24, .08, .02}
	12	$wxyz + wxyz' + wx'y'z + w'x'y'z'$ (.323, .27, .255, .151) (.323, .27, .255, .151)	{.32, .32, .18, .18}	{.27, .27, .27, .18}	{.66, .24, .08, .02}
	13	$wxyz + wxyz' + w'x'y'z + w'x'y'z'$ (.258, .252, .246, .244) (.235, .216, .259, .291)	{.25, .25, .25, .25}	{.25, .25, .25, .25}	{.71, .25, .02, .02}
	14	$wxyz + wxy'z' + wx'yz' + wx'y'z$ (.243, .248, .26, .25) (.295, .131, .333, .241)	{.25, .25, .25, .25}	{.25, .25, .25, .25}	{.74, .09, .09, .09}
	15	$wxyz + wxy'z' + wx'yz' + w'xyz'$ (.235, .245, .26, .259) (.147, .28, .36, .213)	{.25, .25, .25, .25}	{.25, .25, .25, .25}	{.25, .25, .25, .25}
	16	$wxyz + wxy'z' + wx'yz' + w'xy'z$ (.248, .256, .25, .246) (.298, .255, .198, .249)	{.25, .25, .25, .25}	{.29, .29, .21, .21}	{.74, .09, .09, .09}
	17	$wxyz + wxy'z' + wx'yz' + w'x'y'z$ (.277, .275, .261, .186) (.358, .219, .237, .186)	{.25, .25, .25, .25}	{.27, .27, .27, .18}	{.79, .09, .09, .02}
	18	$wxyz + wxy'z' + w'x'yz + w'x'y'z'$ (.225, .261, .247, .267) (.234, .178, .359, .228)	{.25, .25, .25, .25}	{.25, .25, .25, .25}	{.79, .09, .09, .02}
	19	$wxyz + wxy'z' + w'x'yz' + w'x'y'z$	{.25, .25, .25, .25}	{.25, .25, .25, .25}	{.85, .10, .03, .03}

Table 5 (continued)

$D_n[p]$	Structure type	Concept function (rotated proportions) (non-rotated proportions)	GRIT-NP	FFPM	MPMP
		(.247, .247, .251, .255) (.201, .271, .266, .263)			

Appendix 2

Aligning and rotating instances

As a secondary method of analysis, we performed 40 random logical rotations, each of which preserved the alignment of objects occupying identical logical or structural roles, on each structure type to obtain a robust proportion estimate for individual objects within each structure type. An example of the alignment process is provided in Fig. 5.

Unaligned Instances of $3_2[4] - IV$					Aligned Instances of $3_2[4] - IV$				
Boolean cube	O_1	O_2	O_3	O_4	O_1	O_2	O_3	O_4	Boolean cube Axes

Fig. 5 Displayed are three instances of structure type IV from the $3_2[4]$ family of category structures. Column 1 displays the object positions for each instance on a 3-dimensional Boolean cube. Column 2 displays a visual example of each instance with a particular randomized serial position of objects (i.e., flasks). Although these instances are structurally equivalent, they are not aligned properly. To see this more clearly, note how O_2 of the first instance (object ‘000’) acts as a “fulcrum” on the Boolean cube and plays a unique structural role. However, O_1 (object

‘110’) and O_3 (object ‘011’) act as “fulcrums” for the second and third instance, respectively. Each of these three objects occupies the same structural role with respect to this structure type and their observed choice proportions should be appropriately aggregated. Column 3 displays the exact same instances after they have been properly aligned. In other words, O_2 now occupies the same structural role across the three category instances. This can be clearly seen upon inspecting the Boolean cubes (column 4) after the axes have been rotated (column 5)

The rotations ensure that objects occupying identical logical or structural roles are aggregated before statistical analyses are conducted and they mitigate confounding influences, such as differentially-salient stimulus dimensions. Table 6 provides an example of two rotations performed on the six instances of structure type $3_2[4] - II$. Note that there are a finite number of possible rotations for each instance, which is demonstrated in the table with the particular rotation being repeated for instances 1, 3, and 5. Despite this redundancy, there were other rotations (between rotation #2:#39) where the objects did vary in their locations for these instances (as is demonstrated for instances 2, 4, and 6).

Table 6 Example rotations for instances of structure type $3_2[4] - II$

Instance #: $\{O_1, \dots, O_4\}$	Rotation #1	...	Rotation #40
1: $\{010, 110, 101, 001\}$	$\{010, 110, 101, 001\}$	...	$\{010, 110, 101, 001\}$
2: $\{000, 001, 111, 110\}$	$\{000, 001, 111, 110\}$	...	$\{001, 000, 110, 111\}$
3: $\{010, 011, 101, 100\}$	$\{011, 010, 100, 101\}$	...	$\{011, 010, 100, 101\}$
4: $\{000, 100, 111, 011\}$	$\{100, 000, 011, 111\}$	...	$\{000, 100, 111, 011\}$
5: $\{000, 010, 111, 101\}$	$\{010, 000, 101, 111\}$	...	$\{010, 000, 101, 111\}$
6: $\{001, 011, 110, 100\}$	$\{100, 110, 011, 001\}$	...	$\{011, 001, 100, 110\}$

Figure 6 provides a visualization of the rotation process for Instance 6 in Table 6. With respect to this structure type, note how pairs $\{O_1, O_2\}$ and $\{O_3, O_4\}$ occupy identical logical roles in each instance. Rotations involve systematically changing the positions of these pairs for this particular structure type. For example, $\{O_3, O_4\}$ occupy positions $\{O_2, O_1\}$ for Rotation #1 of Instance 6, but are switched to $\{O_4, O_3\}$ for Rotation #40 of Instance 6.

Fig. 6 Boolean cube visualization of the rotation process for Instance 6 of structure type $3_2[4] - II$. Please note that the path provided for each

rotation (in brackets) represents only one possible path for that particular rotation

In Table 7, we replace the object notation (e.g., 011) in Table 6 with the choice proportions observed for this structure type in the current experiment. To obtain the robust proportion estimate across the four objects of this type, we simply average the proportions associated with each object across the 40 rotations. This estimate for structure type $3_2[4] - II$ is shown in the “Average for 40 Rotations” column in Table 7. The average proportions for each of the 41 structure types can be found in column ‘Concept function’ of Tables 4 and 5 in Appendix 1.

Table 7. Demonstration of data aggregation across rotations for structure type $3_2[4] - II$

Data for Exp. 1	Rotation #1	...	Rotation #40
1: $\{.27, .47, .16, .10\}$	$\{.27, .47, .16, .10\}$	...	$\{.27, .47, .16, .10\}$
2: $\{.28, .18, .16, .39\}$	$\{.28, .18, .16, .39\}$	...	$\{.18, .28, .39, .16\}$
3: $\{.34, .10, .14, .42\}$	$\{.10, .34, .42, .14\}$	...	$\{.10, .34, .42, .14\}$
4: $\{.42, .33, .19, .07\}$	$\{.33, .42, .07, .19\}$	...	$\{.42, .33, .19, .07\}$
5: $\{.30, .39, .16, .16\}$	$\{.39, .30, .16, .16\}$	...	$\{.39, .30, .16, .16\}$
6: $\{.08, .22, .43, .27\}$	$\{.27, .43, .22, .08\}$	...	$\{.22, .08, .27, .43\}$
Averages	$\{.27, .36, .20, .18\}$	...	$\{.26, .30, .27, .18\}$
	Average for 40 Rotations		
	$\{.25, .27, .24, .24\}$		

Appendix 3

3A. Representational Information Theory (adapted from Vigo (2016))

We begin by introducing one of the core models of Generalized Representational Information Theory (GRIT; 2012, 2016). The model is referred to as the GISTM and is the basis for Generalized Representational Information Theory. Note that Vigo (2011) introduced Representational Information Theory (RIT), which is a precursor to GRIT. The main difference between RIT and GRIT is that GRIT is not limited by categories defined only over binary dimensions. A second difference is that GRIT uses the law of invariance where the degree of invariance is squared. In our discussion below, we only discuss category structures defined over binary dimensions because in our experiment in this article, we investigate only such category structures. Thus, we focus on examples involving these types of structures. For comprehensive details, we recommend Chapter 9 of the book entitled *Mathematical Principles of Human Conceptual Behavior* (Vigo, 2016).

In what follows, a well-defined category is a set containing objects that are defined in terms of a preset finite number of dimensions and their values. In Representational Information Theory (Vigo, 2011, 2012, 2016), the amount and quality of subjective information h_s conveyed by a subset $R=\{a_1, \dots, a_n\}$ of a well-defined category F is defined as the percentage change (expressed as a proportion) in the structural complexity of F whenever R is subtracted from F . This is expressed by the following equation if we let $G = F - R$ (i.e., G is obtained by subtracting from the original category F the subset or “concept cue” R of F).²

$$h_s(R|F) = \frac{\psi(G) - \psi(F)}{\psi(F)} = \frac{|G| \cdot e^{-\frac{D_0}{D}\Phi^2(G)} - |F| \cdot e^{-\frac{D_0}{D}\Phi^2(F)}}{|F| \cdot e^{-\frac{D_0}{D}\Phi^2(F)}} \quad (1)$$

In the equation above, $R \subseteq F$ or $R \in \wp(F)$, $G = F - R$ (i.e., $F = G \cup R$), $|G|$ and $|F|$ are the number of elements in G and F , respectively, and Φ is the degree of categorical invariance of a set of well-defined objects (see Vigo, 2013, 2016, for a detailed discussion of structural complexity and categorical invariance). Equation 2 below is a more detailed and general version of Eq. 1 (Vigo, 2016, p.173).

$$\begin{aligned} h_s(\{a_1, \dots, a_n\} | G \cup \{a_1, \dots, a_n\}) &= \frac{\psi(G) - \psi(G \cup \{a_1, \dots, a_n\})}{\psi(G \cup \{a_1, \dots, a_n\})} \\ &= \frac{|G| \cdot e^{-\frac{D_0}{D}\Phi^2(G)} - |G \cup \{a_1, \dots, a_n\}| \cdot e^{-\frac{D_0}{D}\Phi^2(G \cup \{a_1, \dots, a_n\})}}{|G \cup \{a_1, \dots, a_n\}| \cdot e^{-\frac{D_0}{D}\Phi^2(G \cup \{a_1, \dots, a_n\})}} \end{aligned} \quad (2)$$

² Note that there is a typographical error in Vigo (2016), page 173, where the second occurrence of “ $|F|$ ” in the numerator of equation 9.7 and the “ $|F|$ ” in the denominator of equation 9.7 should both appear instead as “ $|F'|$ ”.

$$\psi(F) = |F| e^{-\frac{D_0}{D}\Phi^2(F)} \quad (3)$$

As may be seen in Eqs. 1, 2, and 3, the degree of complexity ψ of a well-defined category of objects is defined as the product between the number of objects in the set $|F|$ and the negative exponent of its degree of categorical invariance Φ (Vigo, 2013, 2016), where $\frac{D_0}{D}$ is a structure discriminability index for structures, in which D is the number of dimensions over which the category is defined and $D_0 = 2$ (the minimum non-trivial number of dimensions that define a category). The degree of categorical invariance Φ of F is defined with the following metric (a more general notion of degree of invariance $\widehat{\Phi}$ for which Φ is a special case is discussed in Vigo (2013, 2016); it is the more general notion that is at the core of GRIT, but the simpler version shown here suffices for our analysis involving only binary dimensions):

$$\begin{aligned} \Phi(F) &= \left[\sum_{i=1}^D \left[0 - \frac{|F \cap T_i(F)|}{|F|} \right]^2 \right]^{1/2} \\ &= \left[\sum_{i=1}^D \left[\frac{|F \cap T_i(F)|}{|F|} \right]^2 \right]^{1/2} \end{aligned} \quad (4)$$

To explain the above equation, let's encode the features of the objects in the category F using the digits “1” and “0” so that each of its objects may be represented by a vector of zeros and ones. Now, suppose that F consists of the following three objects: a small black triangle, a small black circle and a large white circle. For example, the vector (1, 1, 1) stands for the first object when $x=1=triangular$, $y=1=small$, and $z=1=black$. Thus, the entire categorical stimulus can be represented by $F = \{(1, 1, 1), (0, 1, 1), (0, 0, 0)\}$. If we perturb (i.e., transform) this categorical stimulus F with respect to the shape dimension (dimension 1) by assigning the opposite shape value to each of the object-stimuli in the set, we will get the perturbed categorical stimulus $T_1(F) = \{(0, 1, 1), (1, 1, 1), (1, 0, 0)\}$ which indicates a transformation along the first dimension. More generally, if F is a categorical stimulus defined over D binary dimensions (where $D \geq 1$), then, for any dimension i ($1 \leq i \leq D$), the transformation T_i on F is defined as follows: $T_i(F) = \{(x_1, \dots, x_i', \dots, x_D) | (x_1, \dots, x_i, \dots, x_D) \in F\}$ where $x_i' = 1$ if $x_i = 0$ and $x_i' = 0$ if $x_i = 1$.

Now, if we compare the original categorical stimulus to the perturbed set, we see that they have two object-stimuli in common. The object-stimuli that “survive” the transformation (from the standpoint of membership in the categorical stimulus) are called “invariants.” Thus, two out of three objects are invariants or remain the same. Intuitively, the ratio of invariants to number of objects in the categorical stimulus may be construed as a measure of the partial structural homogeneity or structural coherence of the categorical stimulus with respect to the dimension of shape and can be written more formally as

Fig. 7 Logical manifold transformations across the dimensions of shape, color, and size for a $3_2[3]$ type structure instance

$|F \cap T_i(F)|/|F|$. In this expression, $|F|$ stands for the number of object-stimuli in the categorical stimulus F and $|F \cap T_i(F)|$ for the number of object-stimuli that F and its perturbed counterpart $T_i(F)$ share.

The first pane of Fig. 7 above illustrates this qualitative transformative process. Doing this for each of the dimensions generates the partial invariance scores of the categorical stimulus F which are then arranged as a vector referred to as the *logical manifold* (*structural manifold* in the general theory) of the Boolean categorical stimulus F . The described procedure defines the Λ operator, which operates on Boolean categorical stimuli to generate their respective logical manifolds as shown by Eq. 5 below.

$$\Lambda(F) = \left(\frac{|F \cap T_1(F)|}{|F|}, \frac{|F \cap T_2(F)|}{|F|}, \dots, \frac{|F \cap T_D(F)|}{|F|} \right) \quad (5)$$

Each component of the resulting logical manifold vector represents a degree of partial invariance for the categorical stimulus. But we also wish to characterize the global or total degree of gestalt invariance of the categorical stimulus and that is what Eq. 3 (the invariance law or GISTM) above accomplishes. We do this by taking the Euclidean distance of each logical manifold from the zero logical manifold whose components are all zeros (i.e., $\mathbf{0} = (0, \dots, 0)$). The zero logical manifold represents the condition of absolute structural categorical disorder where object stimuli in the category are not interrelated from the standpoint of categorical invariance. Degrees of similarity to this baseline condition represent degrees of organization determined by the Λ operator. Thus, Eq. 3, the law of invariance (or GISTM), follows directly from Shepard's similarity measure and it describes the overall degree of categorical invariance Φ (or $\hat{\Phi}$, in the more general theory) of a categorical stimulus F defined over D dimensions.

3B. How to compute representational information

Next, we give an example of how to compute the information values for each of the objects in a stimulus set (i.e., categorical stimulus) using the GISTM. The reader will find more details and examples in Vigo (2012, 2016). We begin by computing the degree of categorical invariance of the stimulus set. First, we can represent this stimulus set by the following concept function:

$$F = x'yz + x'y'z' + xy'z + xy'z' \quad (1)$$

Or, alternatively, we can encode the dimensional values of the stimulus set using the digits “1” and “0” such that $x = 1 = \text{white}$, $y = 1 = \text{round flask}$, and $z = 1 = \text{large}$ and get $F = \{011, 000, 101, 100\}$. This set can be transformed along the color dimension by assigning the opposite color value to all its objects to obtain the *perturbed* set $\{111, 100, 001, 000\}$. Comparing the original set to the perturbed set, we observe that they have two objects in common with respect to the dimension of color. That is, two out of four objects remain the same. This ratio is a measure of the partial invariance of the category with respect to the dimension of color. Conducting a similar analysis with respect to the dimensions of shape and size, an arranging the values as an ordered set or vector, we get the logical or structural manifold of the concept function F above (in other words, we apply the lambda operator defined in appendix C1 on the concept function F as shown in the following equation).

$$\Lambda(F) = \left(\frac{2}{4}, \frac{0}{4}, \frac{2}{4} \right) \quad (2)$$

Relative degrees of global invariance can then be measured by taking the Euclidean distance of each structural or logical

manifold from the zero logical manifold whose components are all zeros (i.e., $(0, \dots, 0)$). Thus, the overall degree of invariance Φ of the stimulus set is given by:

$$\Phi(F) = \sqrt{\left(\frac{2}{4}\right)^2 + \left(\frac{0}{4}\right)^2 + \left(\frac{2}{4}\right)^2} \approx 0.71 \quad (3)$$

Using the degree of categorical invariance, Vigo (2013, 2016) computes the perceived degree of structural complexity ψ of any well-defined category of objects as being directly proportional to its cardinality (i.e., size) and indirectly proportional to the exponent of its degree of invariance as shown below:

$$\psi(\{011, 000, 101, 100\}) = pe^{-\frac{D_0}{D}\Phi^2(F)} = 4e^{-.67(.5)} \approx 2.86(4)$$

Now that we have obtained the degree of perceived structural complexity of the stimulus set using the structural complexity measure, we are now in a position to calculate representational information as shown below. In general, a set of objects is informative about a category whenever the removal of its elements from the category increases or decreases the structural complexity of the category as a whole. That is, the amount of representational information (RI) conveyed by a representation R of a well-defined category F is the rate of change of the structural complexity of F . Simply stated, the representational information carried by an object or objects from a well-defined category F is the percentage increase or decrease (i.e., rate of change or growth rate) in structural complexity that the category experiences whenever the object or objects are removed. Note that this is usually expressed as a proportion that may be rendered as a percentage by multiplying by 100.

In Representational Information Theory (Vigo, 2011) and Generalized Representational Information Theory (Vigo, 2012, 2016), the amount and quality of subjective information h_s conveyed by a subset R of a well-defined category F is defined as the percentage change in the structural complexity of F whenever R is subtracted from F . This is expressed by the following equation:

$$h_s(R|F) = \frac{\psi(G) - \psi(F)}{\psi(F)} = \frac{|G|e^{-\frac{D_0}{D}\Phi^2(G)} - (|F|e^{-\frac{D_0}{D}\Phi^2(F)})}{(|F|e^{-\frac{D_0}{D}\Phi^2(F)})} \quad (5)$$

In the equation above, $R \subseteq F$ or $R \in \wp(F)$, $G = F - R$, $|G|$ and $|F|$ are the number of elements in G and F respectively, and Φ is the degree of categorical invariance of a set of well-defined objects (see Vigo, 2016 for a detailed discussion of structural complexity and categorical invariance). As may be seen in the equation above, the degree of complexity ψ of a well-defined category of objects is defined as the product

between the number of objects in the set $|F|$ and the negative exponent of its degree of categorical invariance Φ (Vigo, 2013, 2016), where $\frac{D_0}{D}$ is a discriminability index for structures in which D is the number of dimensions over which the category is defined and $D_0=2$ (the minimum non-trivial number of dimensions that define a category).

Using this we can calculate the information conveyed by the first element, that is the subset $R = \{011\}$, of the set $F = \{011, 000, 101, 100\}$. The set G is now defined as:

$$G = F - R = \{011, 000, 101, 100\} - \{011\} = \{000, 101, 100\} \quad (6)$$

The logical manifold of the set G is obtained as described earlier. We transform the set along each of its three dimensions and compare the perturbed set with the original to observe the number of objects that are common between the two sets. The logical manifold is:

$$\Lambda(G) = \left(\frac{2}{3}, \frac{0}{3}, \frac{2}{3}\right) \quad (7)$$

Using the logical manifold, we now obtain the degree of categorical invariance as

$$\Phi(G) = \sqrt{\left(\frac{2}{3}\right)^2 + \left(\frac{0}{3}\right)^2 + \left(\frac{2}{3}\right)^2} \approx 0.94 \quad (8)$$

and subsequently obtain the degree of perceived structural complexity as:

$$\psi(G) = pe^{-\frac{D_0}{D}\Phi^2(G)} = 3e^{-.67(.88)} \approx 1.66 \quad (9)$$

Substituting the values of $\psi(F)$ and $\psi(G)$ above gives us the information conveyed by the object 011 in $R = \{011\}$ about its set of origin $F = \{011, 000, 101, 100\}$: that is,

$$h_s(R|F) = \frac{\psi(G) - \psi(F)}{\psi(F)} = \frac{1.66 - 2.86}{2.86} = -0.4196 \approx -0.42 \quad (10)$$

Repeating the above analysis, we can obtain the amount of information conveyed by all possible single element representations of F . The table below shows these approximate information values.

R	F	G = F - R	$\psi(G)$	$h_s(R F)$
{011}	{011, 000, 101, 100}	{000, 101, 100}	1.66	-0.42
{000}	{011, 000, 101, 100}	{011, 101, 100}	2.22	-0.22
{101}	{011, 000, 101, 100}	{011, 000, 100}	2.22	-0.22
{100}	{011, 000, 101, 100}	{011, 000, 101}	3.00	0.05

The cognitive basis for this removal is selective attention. GRIT posits that observers have the ability of psychologically “subtracting” an object from a category (as is the case with a single element concept cue) temporarily by selectively attending to all the members of the category except the object in question (while at the same time suppressing the object in memory). The results of the coherence and complexity judgments of the category with and without the object are then computed and stored in working memory where the computation of the most informative object takes place per GRIT’s prescription described above.

References

- Bourne, L.E. (1966). *Human conceptual behavior*. Boston, MA: Allyn and Bacon.
- Bourne, L.E. (1970). Knowing and using concepts. *Psychological Review*, 77(6), 546–556.
- Chin-Parker, S., & Ross, B.H. (2004). Diagnosticity and prototypicality in category learning: A comparison of inference learning and classification learning. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 30(1), 216–226.
- Devlin, K. (1991). *Logic and information*. New York, NY: Cambridge University Press.
- Doan, C.A., & Vigo, R. (2016). Constructing and deconstructing concepts: On the nature of category modification and unsupervised sorting behavior. *Experimental Psychology*, 63(5), 249–262.
- Estes, W.K. (1986). Array models for category learning. *Cognitive Psychology*, 18, 500–549.
- Estes, W.K. (1994). *Classification and cognition* (Vol. 22). New York, NY: Oxford University Press.
- Feldman, J. (2000). Minimization of Boolean complexity in human concept learning. *Nature*, 407, 630–633.
- Feldman, J. (2003). A catalog of Boolean concepts. *Journal of Mathematical Psychology*, 47, 75–89.
- Feldman, J. (2006). An algebra of human concept learning. *Journal of Mathematical Psychology*, 50, 339–368.
- Garner, W.R. (1974). *The processing of information and structure*. New York, NY: Wiley.
- Kruschke, J. (1992). ALCOVE: An exemplar-based connectionist model of category learning. *Psychological Review*, 99(1), 22–44.
- Kurtz, K.J., Levering, K.R., Stanton, R.D., Romero, J., & Morris, S.N. (2012). Human learning of elemental category structures: Revising the classic result of Shepard, Hovland, and Jenkins (1961). *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 39(2), 552–572.
- Logan, R.K. (Ed.). (2012). Information: Its different modes and its relation to meaning [Special issue]. *Information*, 3(1–4) & 4(1).
- Luce, R.D. (2003). What ever happened to information theory in psychology? *Review of General Psychology*, 7(2), 183–188.
- Medin, D., Goldstone, R.L., & Gentner, D. (1993). Respects for similarity. *Psychological Review*, 100(2), 254–278.
- Morgenstern, Y., Schmidt, F., & Fleming, R.W. (2019). One-shot categorization of novel object classes in humans. *Vision Research*, 165, 98–108.
- 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. (1986). Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology: General*, 115(1), 39–57.
- Nosofsky, R.M., & Zaki, S.R. (2002). Exemplar and prototype models revisited: Response strategies, selective attention, and stimulus generalization. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 28, 924–940.
- Nosofsky, R.M., Gluck, M.A., Palmeri, T.J., McKinley, S.C., & Glauthier, P.G. (1994). Comparing models of rule-based classification learning: A replication and extension of Shepard, Hovland, and Jenkins (1961). *Memory and Cognition*, 22, 352–369.
- Reed, S.K. (1972). Pattern recognition and categorization. *Cognitive Psychology*, 3(3), 382–407.
- Rehder, B., & Hoffman, A.B. (2005). Eye tracking and selective attention in category learning. *Cognitive Psychology*, 51, 1–41.
- Reicher, G.M. (1969). Perceptual recognition as a function of meaningfulness of stimulus material. *Journal of Experimental Psychology*, 81(2), 274–280.
- Rosch, E. (1978). Principles of categorization. In E. Rosch & B. B. Lloyd (Eds.), *Cognition and categorization*. Hillsdale, NJ: Erlbaum. Reprinted in: Margolis, E. and Laurence, S. (Eds.) (1999). *Concepts: Core readings*. Cambridge, MA: MIT Press.
- Rosch, E. & Mervis, C.B. (1975). Family resemblances: Studies in the internal structures of categories. *Cognitive Psychology*, 7(4), 573–605.
- Shannon, C.E. (1948). A mathematical theory of communication. *Bell System Technical Journal*, 27, 379–423.
- Shepard, R.N. (1974). Representation of structure in similarity data: Problems and prospects. *Psychometrika*, 39(4), 373–421.
- Shepard, R.N. (1987). Toward a universal law of generalization for psychological science. *Science*, 237(4820), 1317–1323.
- Shepard, R.N., Hovland, C.L., & Jenkins, H.M. (1961). Learning and memorization of classifications. *Psychological Monographs: General and Applied*, 75(13), 1–42.
- Smith, J.D., & Minda, J.P. (1998). Prototypes in the mist: The early epochs of category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 24(6), 1411–1436.
- Smith, J.D., Murray, M.J., & Minda, J.P. (1997). Straight talk about linear separability. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 23(3), 659–680.
- Smith, J.D., Minda, J.P., & Washburn, D.A. (2004). Category learning in rhesus monkeys: A study of the Shepard, Hovland, and Jenkins (1961) tasks. *Journal of Experimental Psychology: General*, 133(3), 398–414.
- Vigo, R. (2009). Categorical invariance and structural complexity in human concept learning. *Journal of Mathematical Psychology*, 53(4), 203–221.
- Vigo, R. (2011). Representational information: A new general notion and measure of information. *Information Sciences*, 181(21), 4847–4859.
- Vigo, R. (2012). Complexity over uncertainty in generalized representational information theory (GRIT): A structure-sensitive general theory of information. *Information*, 4(1), 1–30.
- Vigo, R. (2013). The gist of concepts. *Cognition*, 129(1), 138–162.
- Vigo, R. (2016). *Mathematical principles of human conceptual behavior: The structural nature of conceptual representation and processing* (Paperback Ed.). New York, NY: Routledge.

- Vigo, R. (2019). A simple descriptive method and standard for comparing pairs of stacked bar graphs. *arXiv: 1910.11524*.
- Vigo, R., & Basawaraj, B. (2013). Will the most informative object stand? Determining the impact of structural context on informativeness judgments. *Journal of Cognitive Psychology*, 25(3), 248–266.
- Vigo, R., Zeigler, D.E., & Halsey, P.A. (2013). Gaze and informativeness during category learning: Evidence for an inverse relation. *Visual Cognition*, 21(4), 446–476.
- Vigo, R., Evans, S.W. & Owens, J.S. (2014). Categorization behavior in adults, adolescents and ADHD-adolescents: A comparative investigation. *The Quarterly Journal of Experimental Psychology*, 68(6), 1058–1072.
- Zeigler, D.E., & Vigo, R. (2018). Classification errors and response times over multiple distributed sessions as a function of category structure. *Memory & Cognition*, 46(7), 1–17.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.