



City Research Online

City, University of London Institutional Repository

Citation: Hampton, J. A. (2020). Investigating differences in people's: Concept representations. In: Marques, T. & Wikforss, A. (Eds.), *Shifting Concepts: The Philosophy and Psychology of Conceptual Variability*. (pp. 67-82). Oxford, UK: Oxford University Press. ISBN 9780198803331 doi: 10.1093/oso/9780198803331.003.0005

This is the accepted version of the paper.

This version of the publication may differ from the final published version.

Permanent repository link: <https://openaccess.city.ac.uk/id/eprint/27626/>

Link to published version: <https://doi.org/10.1093/oso/9780198803331.003.0005>

Copyright: City Research Online aims to make research outputs of City, University of London available to a wider audience. Copyright and Moral Rights remain with the author(s) and/or copyright holders. URLs from City Research Online may be freely distributed and linked to.

Reuse: Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

City Research Online:

<http://openaccess.city.ac.uk/>

publications@city.ac.uk

Hampton, J.A. (2020). Investigating Individual Differences in Concepts. In T.Marques & Å.Wikforss (Eds.), *Shifting Concepts. The Philosophy and Psychology of Conceptual Variability* (pp. 67-82). Oxford: Oxford University Press.

James A. Hampton

Author's Proof.

City, University of London

hampton@city.ac.uk

Abstract: Semantic memory tasks can focus on intensions (features and properties) or extensions (reference and categorization). There is evidence for stable individual differences in each of these tasks, but these differences do not correspond across tasks. Two further studies are reported showing that under certain conditions the correspondence can be demonstrated. Reasons for the lack of connection are discussed in terms of different systems for conceptual understanding involving similarity versus theory-based conceptualisation.

Keywords: Intension, Extension, Individual differences, Concepts, Categorization

Language serves many purposes in human society, and one of its prime functions is as a tool for communication. As you read this sentence, you are taking the words and their grammatical marking and deriving an understanding of what the author is trying to convey. A key element in this system of communication is the actual words that are used, and it can only work successfully if people have a common understanding of what those words mean. The words of a given language have a conventionally agreed meaning, or set of meanings, which, for languages such as English, can be found in a dictionary. Dictionaries allow us to determine when someone is using a word correctly or in error. They act as a brake on language evolution, slowing down the natural process by which word meanings change over generations, and at the same time they improve the likelihood that we will understand one another.

Given the crucial importance of a shared understanding of word meanings for underpinning successful communication, it is interesting to ask just how strong is that shared understanding. People appear generally to understand each other most of the time, but are there actually stable individual differences in how people grasp any given concept and the meaning of its associated term? The fact of successful communication within a language community has been argued by some philosophers to be evidence for the identity of concepts held by different agents. The so-called *Publicity Constraint* (e.g. Rey, 1983) proposes that if two individuals were not using the same identical concept, then it would no longer be possible to resolve any argument of fact between them that depended on that concept. They would necessarily be talking at cross-purposes if they had different meanings for a key word in the debate. While different individuals may have different beliefs or ideas *about* a particular kind or class of thing, the *concept* about which their ideas are in conflict itself has to be the same concept.

The extent to which this Publicity Constraint works in practice is the subject of much debate. My focus here will be to review some empirical research into the question of shared understandings. As such, I will be leaving aside the question of whether it is possible to investigate people's concepts

directly, or whether in fact empirical enquiry into people's concepts is actually tapping into something rather different – for example their *conceptions* of those concepts (Rey, 1983). In this chapter I will review a set of studies that have looked directly at how much agreement and stability there is in people's reports of their concepts (or conceptions of concepts). I will then describe some research recently published with Alessia Passanisi and others that tries (and fails) to map the similarity of people's concept intensions (the properties they consider important for a concept) onto their concept extensions (the degree to which they consider that different exemplars are typical of the concept). I will also describe some subsequent attempts to follow up this disturbing null result.

Individual differences in concept tasks

It will be helpful first to outline the kind of conceptual differences that will be the focus of this chapter. We will not, for example, be considering differences in word meanings between US and UK English (**pants** as underwear or **pants** as trousers), nor the well-known problems of finding exact translations between languages (English **chair** versus French **chaise** and **fauteuil** for example). The differences we will be considering are more subtle and emerge from methodology aimed at eliciting people's personal meanings for words in various more or less direct ways. They are found within samples of speakers within the same language community, and relate to an aspect of word meaning that has been variously described as a *stereotype* (Putnam, 1975) or *prototype* (Rosch, 1975; Lakoff, 1987).

The notion of a prototype for a concept has been proposed to explain some very widespread phenomena to be found in studies of people's understanding of concepts (Hampton, 2006). Briefly, there are two primary aspects to any concept that can be explored psychologically. The *intension* of a concept is the set of associated features or properties that people consider to be a part of the concept, and that determine category membership. For example, for the concept BIRD it would include features such as having feathers, hatching from eggs and having two legs. The *extension* of a concept by contrast is the set of entities (objects, creatures, situations) in the world to which the

concept term can be applied – it is the category of things that “fall under” the concept. The two aspects, intension and extension, should in principle be closely related. It is in virtue of possessing the intensional properties of a concept, that an individual entity will be included in the extension of that concept. Any feathered creature that hatches from eggs and has two legs and a beak will be a bird, and any creature lacking any of these features will be something else.

It should be noted that my use of the terms Intension and Extension here differs from the classic meaning of the terms introduced by Frege (1892), and as commonly used in philosophy and semantics. I will be using the terms to refer to behavioural measures of what a person may represent in their mind about a concept. Traditionally, Intension is what determines Extension, so there is no possibility that the two do not precisely correspond. When we consider the psychology of an individual possessor of a concept however, it is possible that the beliefs they hold about an Intension (how they think they define a concept in terms of its properties) may not in fact map well onto the beliefs that they hold about the reference or Extension of the concept (what they think the concept term refers to in the world). An intriguing example of this mismatch was a study by Levitis, Lidicker and Freund (2009) in which they showed that members of scientific societies concerned with the study of “behaviour” frequently had self-contradictory views about what features defined the term for them, and what examples in the world they would count as falling under the term. My use of Intension and Extension should be understood then as a short-hand for “information that people are able to provide about their understanding of a concept’s defining features and their beliefs about its application or reference in the world.”

Prototype theory (Hampton, 2006, Rosch & Mervis, 1975) takes this difficulty in defining terms to be characteristic of most of our everyday concepts. Following Wittgenstein (1953), Rosch and Mervis (1975) argued that the notion of all and only category members possessing a common set of defining features simply does not work for a large number (indeed a majority) of semantic concepts in everyday language. Hampton (2006) summarises the evidence. On the intensional side, when people

are asked to generate features of a concept they make no differentiation between features common to the class (birds having feathers) and those which are only found in typical instances (for example birds flying). In fact determining a good set of defining features for any given concept has been a major intellectual challenge going back to Plato's Socratic dialogues.

Furthermore, the case of birds is exceptional. Most other categories have no clear set of defining features which taken together can determine the correct extension of the concept. Wittgenstein (1953) famously pointed out that if we consider the category of GAMES, there is no common element running through them all, but rather a set of *family resemblances*, or attributes which are common to the majority of cases, though not necessarily to all. Research in the 1970s confirmed this intuition by showing how it is often impossible to generate clear-cut common feature definitions from people's descriptions of their conceptual categories (Hampton, 1979; McNamara & Sternberg, 1983).

Turning to extensions, there are similar problems. On the one hand, there are often borderline cases to membership in the extension. The concepts are vague in much the same way that adjectives such as TALL or BALD are vague. There seems to be no higher authority to which one can appeal to decide if an AVOCADO should be called a FRUIT, or if TRAMPOLINING should be counted as a SPORT. People disagree with each other, and are prone to change their opinion from one week to the next (McCloskey & Glucksberg, 1978), or depending on the context in which the question is asked (Hampton, Dubois & Yeh, 2006). Together with this vagueness at the category boundary, Rosch (1975) also discovered the phenomenon of *Typicality*. The members of a category differ from each other in terms of how typical they are of the concept. Robins and sparrows are more typical as birds than penguins and ostriches, even though all four are *bona fide* members of the category. Typicality has been found to affect a wide range of tasks such as perceptual processing, ease of recall, and inductive reasoning (Murphy, 2002). Hampton (1988; 1995) showed that for many categories the degree to which an item belongs in a category (its graded membership) and its typicality in the

category should be treated as two measures based on the same underlying scale of closeness or similarity to the concept prototype.

Taken together, the evidence for prototypes is overwhelming. The individual differences in concepts that will be the focus of this chapter relate to the prototypes possessed by different speakers. To access someone's prototype for a concept (be it BIRD, FRUIT or SCIENCE), there are several tasks that have been commonly used. In *production* tasks, people are asked to list the properties that they consider important or defining for a concept (the intension) and people are asked to list the exemplars that they consider to fall within the named category (the extension). Because of the rather idiosyncratic data that production tasks can generate, lists of features and exemplars are usually then submitted to *rating* tasks in which people rate the degree to which each property is an important part of the concept, and the degree to which each exemplar is typical of the category. In addition to typicality judgments, one can also obtain categorization judgments where lists of items including clear members, borderline cases and clear non-members are categorized as in or out of the category, and the probability of a positive categorization can be measured.

For this range of different tasks probing people's semantic memories, it is possible in each case to determine whether people have their own personal take on a concept that differs from other people's. The procedure requires that people perform that same task on two separate occasions, usually a week or more apart. The average consistency with which someone gives the same responses on each occasion can then be compared with the average amount of agreement between one person and another. If a person's responses on the second occasion are closer to their responses on the first occasion than they are to those of some randomly chosen other individual, then there is *prima facie* evidence that people have what we might term stable "*idio-prototypes*" in their personal idiolect, which differ from those of other people.

Evidence for idio-prototypes

Interest in exploring idio-prototypes began in the 1980s with a series of studies by Barsalou (1987) who investigated the stability of typicality judgments, and Bellezza (1984a; 1984b) who tested the stability of responses in a range of different semantic tasks. To take an example, Barsalou provided students with lists of category members to be judged in terms of their Typicality in a category. For example, they may have ranked different activities such as baseball, swimming, athletics and archery for their typicality as examples of Sport. Participants worked through several category lists, and then returned after a certain period of time to do the task again. Correlating their judgments over time, Barsalou found that the correlations declined over the first 24 hours, but then levelled off at around 0.8. On the other hand, he reported that the average correlation between any two individuals doing the same task was much lower, between 0.3 to 0.6. Similar results were found by Bellezza for production tasks where students generated either category exemplars or category definitions on two occasions a week apart. Within-participant consistency (correlation) was around 0.7 for category exemplars, and 0.5 for definitions, while between-participant agreement or consensus was just 0.4 and 0.2 respectively.

The finding that people's semantic judgments are stable and replicable and that the level of agreement with others is much lower than the level of consistency for an individual over time provides good evidence that people have idiosyncratic meanings for these concepts. It is possible that at least some of the effect may reflect more general individual differences, such as the way that people choose to interpret the instructions, or differences between individuals in the strategy that they may use to perform the tasks. However the size and consistency of the effect across different tasks suggests that there is more to the result than this.

Using similarity to assess individual differences in concepts

Theories of semantic memory and concepts in psychology share a common fundamental assumption. Intensions and extensions should be strongly related. Extensional categories of reptiles and amphibians should be distinguished by their different intensions. Similarly, the intensional

differences between reptiles and amphibians should be discoverable (perhaps with some difficulty) by examining their extensions. Scientific categorization typically evolves through a cyclical process of refining intensions/definitions of classes and extensions/reference of terms until a satisfactory classification has been achieved. The two sides of a concept, intension and extension, are thus tightly bound. The question then arises as to whether the individual differences that we see in prototype production and judgments tasks will match across the two types of semantic information. We know from Barsalou, Bellezza and others that there are quite strong and stable individual differences in extension (typicality and membership judgments) and in intension (production of definitions and ratings of feature importance). So it should therefore be possible to show that individual patterns of similarity and difference between individuals in how they judge extensions should be mirrored in patterns of similarity and difference in how they judge intensions.

Consider a simple example, from work by Verheyen and Storms (2013). In their study they were interested in whether concepts such as Sport might show individual differences in conceptual contents. After all, there are some sports involving individual skilled activities which take place indoors, such as darts, billiards, chess or snooker. High levels of skill are involved, but little in the way of physical exertion. On the other hand, there are outdoor sports which are tests of strength or endurance more than skill. Marathon running, triathlon, and long distance swimming would be examples. So it is possible that when activities are borderline to the category (like chess or hill walking) some people may favour one type of sport, and other people may favour the other. This was in fact what Verheyen and Storms found. Fitting a quantitative model to categorization probabilities, they discovered that dividing their participant group into subgroups produced a significant improvement in fit to the data, with each subgroup providing a different set of categorization probabilities to the items at the borderline. They also showed that for some categories one could speculate about the underlying intensional basis for the group split, although their evidence for this was weak. They did not have data on individual views about intensions, and so could only provide post hoc accounts of the different categorization profiles.

Hampton and Passanisi (2016) followed up an earlier set of studies that I conducted in the 1980s, and looked for more direct evidence that individual or group differences in extension would mirror differences in intension. The method used was quite simple. A group of students were asked to make extensional judgments for six semantic categories – in this case judgments of typicality of a list of category members (exemplars). The same group also made intensional judgments about the same categories, namely to judge the importance of a list of features for defining the category.

Correlations within each set of judgments were then used to produce two similarity matrices for the group of students, showing for each individual a measure of how similar their responses were to each of the other individuals in the group. One similarity matrix was produced for the extensional (typicality) judgments, and another for the intensional (feature importance) judgments, and we then looked to see if they corresponded. If a pair of individuals had similar ideas about which Sports were most typical, did they also share similar ideas about what features of Sports were most important for deciding membership.

The strange result that we obtained was that there appeared to be zero relationship between these two sets of similarities. A pair of participants who had similar ideas about typicality showed no greater similarity than average for feature importance ratings, and vice versa. There appears to be a disconnect between the ways in which people vary in their extensional beliefs and the way in which they vary in their intensional beliefs.

Is this null result to be believed? What if the variation between individuals reflects random responding and they all have fundamentally the same concept prototypes? That could then account for the lack of correlation between the similarity matrices. To test this and other challenges, we provided several controls in our study. First, we tested people's judgments again a week or two later and recomputed the similarity matrices. If there is little or no difference between people and the similarities reflect randomness, then the similarity in Typicality (or Feature importance) at Week 1 should not correlate with similarity for the same measure at Week 2. In fact, we found that there

was a significant positive correlation across occasions in the similarity matrices. So the similarity and difference between individuals in each task was stable over time. There are genuine differences in how people make the two judgments, but the differences on the extensional task do not correspond to the differences on the intensional task. A second control was to consider that differences between people may reflect some general strategy for responding to the question. Perhaps different groups interpret each task in different ways, or they use the rating scale in different ways, and so the similarities reflect general effects, and not the contents of people's prototypes. To answer this possibility, we calculated the correlation between the similarity matrices for Typicality (or Feature Importance) for different pairs of categories. General strategic differences should show up as positive correlations between the similarity matrices for different categories. We found that there was indeed a small positive correlation across categories for each task, but this was much lower than the correlation for the similarity matrices for the same category over time, confirming that the stable individual differences related to specific categories. A pair of individuals with similar ideas of rating typicality in Sports were much less likely to be similar when rating typicality in another category such as Science.

Finally, we questioned whether we had the right set of features and exemplars to provide a good test of the hypothesis. If the features were not relevant to determining degrees of typicality in the category that could explain the null result. Using data from the Leuven concept database (De Deyne et al., 2008; Verheyen & Storms, 2013), we showed that the number of features that an item possessed did indeed correlate well with the mean typicality of that item. The correlation averaged .80 based on the 12 features that we had sampled, compared with a correlation of .82 using the full set of 29 to 39 features in the norms. So our features were representative of the full set.

A further study: connecting the two tasks

In an attempt to bring the two aspects of conceptual representation closer together in our data, Alessia Passanisi and I conducted a new study. For this study we had 188 participants divided into 8

groups of 21-25 each on the basis of three between-subjects factors. All participants judged the typicality of exemplars, and the importance of features for two different categories, Sports and Fish. All participants did all four tasks. Across groups we varied the order of doing the four tasks, and the instructions they were given. The design is shown in Table 1. While counterbalancing whether the Fish or Sports category was done first, the two tasks were always done for one category before moving on to the other category. This procedure was a marked change from our earlier method where participants performed the first task on 6 categories before moving on to the other task. In this study we deliberately had the extension and intension tasks for a particular category placed next to each other, in order to increase the influence of one on the other.

In addition the order of the two tasks for each category was counterbalanced. We were interested in seeing whether, when the two tasks are set side-by-side in this way, there would be a greater influence of one upon the other. Having just decided which sports are most typical, we expected there would be a better chance that people would then also judge the features of those typical sports more important. If that turned out to be the case, we could then also tell whether the influence was symmetrical (extension on intension and vice versa) or not. For example, if people primarily represent concepts by their exemplars (Storms, 2004) then extension should affect intension more than the reverse.

Our third factor was the instructions that were given. One set of four groups had “standard” instructions as used in our earlier studies, asking the participants to rank the typicality of 12 exemplars and then to rank the importance of 12 features as two separate tasks. The “connected” instructions gave the instructions for both tasks at the beginning. After instructions for one set of rankings (e.g. exemplar typicality), the booklet then explained that participants would subsequently be asked to do the other (e.g. feature importance). In the version where the Feature Importance task came first, instructions (translated here from Italian) were as follows:

“First, we want you to tell us how important you think each property is for deciding whether something is in the category. An important property is one that you find often in members of the category – most members have it, and that is also not so often found in other kinds of thing. So it is distinctive to that category.”

[An example was given here]

“After this task, we want you to give us a judgment of how representative or typical you think each word is of its category. A typical example of a category is one that has lots of properties in common with the other members of the category, thus it would be a good example to represent what that category is normally like. Judge typicality in terms of the features that you have chosen as most important for the same category.”

Equivalent instructions were given for those doing the tasks in the opposite order, again referring to the first task when giving instructions for the second.

We hoped that this combination of the two instructions into a single page before the start of either task would encourage participants to connect the two tasks. In that case, the pairwise similarity matrices for the two tasks may be expected to show a positive correlation.

The tasks were administered as paper booklets to a large class of students at Kore University in Enna, Sicily. Five participants were excluded on the basis of strong evidence of collusion with other participants. Reliability (Cronbach’s alpha) for each of the ranking scales was high (Mean = .944, range from .83 to .99).

The data were analysed to generate similarity matrices for each of the 8 conditions for each of the four tasks (extension and intension rankings for Sports and for Fish). Correlations between the similarity matrices were then calculated. Following our earlier study (Hampton & Passanisi, 2016) we compared the correlation between similarities for Typicality and Importance within the same category to the same correlation calculated across different categories. The latter correlation shows

the possible effect of general strategic differences between individuals, so the difference between these two correlations will show the extent to which concept-specific individual differences in extension and intension can be mapped. Table 2 shows the results for Sports and Fish and for the two types of instruction.

Table 2: Results of the study. Correlations of the Similarity matrices for Exemplar Typicality and Feature Importance rankings, within and between the two categories.

	Standard Instructions	Connected Instructions
Sports Within Category	.25	.03
Fish Within Category	.29	.24
Between categories (control)	.16	.12

The first thing to note is that the effect of instructions was in the wrong direction. Asking people to think about both tasks at the same time and to coordinate their responses led to reduced correlation between the similarity patterns for the two tasks, particularly for the Sports category. Otherwise there was some evidence here for the Within-category correlations to be higher than the Between-category control. To get an idea of the statistical accuracy of these figures, an estimate of the standard deviation of the correlations was obtained using a randomisation procedure where the participants' data in the two similarity matrices were set out of alignment. The resulting mean correlation was zero, and the standard deviation was .09.

In sum, having participants do the two procedures one after the other for the same category did have some positive effect on the correlation between extensional and intensional similarity. Instructing the participants to deliberately consider one task when performing the other however did not have the expected result. If anything, the two tasks were less related under these conditions. The unexpected differences between the two categories calls for a further study using a larger number of categories.

Implications

The results of these experiments represent something of a paradox. There is little reason to doubt that the conceptual categories that we form are based on intensional properties. It is a commonplace since Aristotle that we judge whether a concept applies to an item by considering shared properties. The whole tradition of scientific discovery has been based on the classification of the world on the basis of properties, which are then selected and redefined to provide categories that optimise our ability to understand and predict the processes involved. The results from Hampton and Passanisi (2016) do not undermine this general principle. In fact we reported that the degree to which exemplars possess the category features was correlated at 0.8 with judgements of typicality. After all what else could determine whether some activity is a sport or a science other than its having enough of the right kind of properties? What the results suggest is that when one person's prototype concept differs from another's, the *dimensions of difference* for extensional information (typicality and graded membership of categories) are unrelated to the dimensions of difference for intensional information (importance of properties for defining the category). Only by placing the tasks side by side did we find any evidence for correspondence in the individual similarities and differences in responses, and that was far from strong.

Table 1: Design of the unpublished study by Hampton & Passanisi

Group	First	Second	Third	Fourth	Instructions
A	Typicality Sports	Importance Sports	Typicality Fish	Importance Fish	Standard
B	Typicality Fish	Importance Fish	Typicality Sports	Importance Sports	Standard
C	Importance Sports	Typicality Sports	Importance Fish	Typicality Fish	Standard
D	Importance Fish	Typicality Fish	Importance Sports	Typicality Sports	Standard
E	Typicality Sports	Importance Sports	Typicality Fish	Importance Fish	Connected
F	Typicality Fish	Importance Fish	Typicality Sports	Importance Sports	Connected
G	Importance Sports	Typicality Sports	Importance Fish	Typicality Fish	Connected
H	Importance Fish	Typicality Fish	Importance Sports	Typicality Sports	Connected

One way to understand this result is in terms of a dual process theory of categorization (Ashby et al. 1998). It has been hypothesized that the brain has two mechanisms for learning to categorize. One involves a similarity-based associative learning in which people learn to associate category labels with particular regions of similarity space (Gärdenfors, 2000). The other involves the induction of an explicit rule for categorizing on the basis of observed properties. When a category can be readily discriminated with a rule (e.g. members of category A tend to be larger in size than Category B) then an explicit rule can be learned. But when two dimensions need to be integrated to determine the category (for example members of category tend to be larger but also rounder) then associative processes are more likely to be used. We can speculate therefore that when judging typicality of category exemplars, such as whether snooker or wrestling is a better example of sport, people are basing their judgment on an associative similarity between these activities and their general prototype for sport. The position of these items within a similarity space relative to the prototypical centre of the category would be used to make the ranking. On the other hand when asked to judge whether exercise or skill is more important as a feature of sport, people are retrieving (or attempting to construct) a rule-based representation of the category. They are thinking of what they have heard or read about the value and meaning of sport within a socio-cultural context. Do people enjoy it, is it healthy, what should count as a sport in general? These two modes of thinking both show systematic and stable individual variation, according to our paper (Hampton & Passanisi, 2016). But they tap into different parts of the semantic memory system. While intensions and extensions have to be strongly coordinated at a social level (that is at the level of group data and common understanding) the individual variation is overlaid onto this coordination in a way that is largely independent for the intensional and extensional representational system.

A more positive result

To conclude, I will report some results from a recent study initiated by Farah Djalal and Tom Heyman with Gert Storms at the Catholic University in Leuven, to which I have subsequently contributed (Djalal, Heyman, Storms & Hampton, 2017). In a conceptual replication of Hampton and Passanisi (2016), the tasks for intension and extension were changed. The extension task used was a binary category membership judgment rather than a typicality judgment. For each of 8 categories, participants were shown a set of 15 pictures and had to click on pictures of exemplars that were category members. Prior to this, they completed a property generation task. Rather than rating or ranking the importance of features, the participants were asked to generate their own properties as if explaining the terms to someone who did not understand them. The properties generated were then used for measuring intensional similarity, while correlation of the binary category membership judgments was used for measuring extensional similarity.

Initial results confirmed the earlier findings of Hampton and Passanisi (2016). The correlation between the extensional and intensional similarities was close to zero (mean of .03 and -.07 for two different ways of computing intensional similarity). However a different analysis did show a connection between the two sides of concept representations. The properties generated to the category names were combined with the exemplars from the extensional task into an Exemplar by Property matrix. Separate groups of four students then filled in each matrix putting a 1 if each exemplar (in rows) had each property (in columns), and a 0 otherwise. From this matrix it was then possible to generate predicted values for a graded membership scale, based on the sum of feature applicability scores for each exemplar. Simply put, the more features that an exemplar possessed, then the better member of the category it should be. (As in Rosch and Mervis, 1975, no attempt was made to provide differential weighting to the features). That was indeed the case. The next step was to consider each individual in the original group, and calculate a predicted degree of membership for the set of exemplars based on two separate sets of properties. One set contained those properties which the individual themselves had generated (*individual properties*) and the other set contained those properties which they did not generate (*residual properties*). A logistic regression analysis was

then used to predict the person's own categorization of exemplars on the basis of the two different sets of properties. This analysis was run for each individual participant, taking their own set of individual properties (and residual properties) to predict their own categorization judgments. It was found (as expected) that both individual and residual properties played a role in the prediction for any given individual. However when the data were shuffled so that individual A's properties were used to predict other individuals' (B, C, D etc) categorizations, the weight in the model for individual properties declined, while the weight for residual properties increased. When the individual model parameters were averaged across all the participants in the sample, this pattern was seen in over 98% of 1000 random shuffles of the data.

This analysis is quite complex, and is also different in many ways from the comparison of similarity matrices that we used before. It involves showing that the properties that a given person generated for a category were more predictive of their own categorization than of other people's categorization of the exemplars in that category. (Recall also that the extensional measure here was a yes/no membership decision about pictures of exemplars).

What has been shown by this new study? While the use of the previous method of analysis – correlating similarity matrices for extensions and intensions showed the same lack of correspondence as before, when an individual's own generated properties were used to predict their own exemplar categorization a better correlation was found than when either the properties or the categorization behaviour were from a different individual. It is possible but probably unlikely that the use of different intensional and extensional measures was responsible for the positive result this time. It is possible that asking people to generate their own properties rather than rank a fixed set will require them to access a deeper level of personal meaning. It is known that what people say and what people do are often at odds (Nisbett & Wilson, 1977). So people may explicitly rank a particular property as important on the basis of some higher level theory or belief that they hold, whereas when it comes to generating their own features they may rely more on their extensional

beliefs about exemplars. However, the generation task alone was not sufficient to show a link between intension and extension. When similarity matrices based on generated properties were compared with similarity for extensions, they showed no more correspondence than before. It was only when those properties generated by an individual were used to predict their own categorization of exemplars that a greater level of correlation between predicted and actual categorization was seen. At this stage the most likely explanation is one of statistical sensitivity or power. It may be that measuring similarity between individuals is not a sensitive enough measure. Reliability of the similarity matrices over time in Hampton & Passanisi varied around .3 to .4, which is quite low, even though clearly greater than zero. Time may tell whether more powerful designs can be found to show a significant effect with this method of analysis

Conclusions

We set out to show that the two aspects of prototype concept representations, intensions and extensions can be linked at the level of individual differences. Repeated attempts to show that similarity between individuals for intensions can be mapped to similarity between individuals for extensions have produced null results. Hampton and Passanisi (2016) reported four experiments with no positive result. Our replication with Sports and Fish reported above showed that when the two tasks are done one immediately after the other, a low level of correlation can be seen, although instructing people to explicitly make the connection had a negative impact (for some reason only for Sports). The new study by Djalal et al. (2017) confirmed the lack of a correlation between similarity matrices, based this time on a property generation and exemplar categorization task. However a different analysis, possibly more sensitive to individual variation, was able to show a significant link between a person's own individually generated properties and their own categorization.

How good then is the coordination of different individuals' conceptual representation? There is clear evidence that people differ from each other, both in judging the typicality and membership of category items, and in judging the importance of, or generating category features. As a next step in

the research it will be interesting to know how these differences relate to other domains, such as social categories, understanding of political positions or judgments about crime and morality. How labile are the individual differences that we have seen? They persist over time, but can they be easily influenced by recent experience? In a year when a male British tennis player became World Number One, how does the typicality of tennis change for British as opposed to American (or Australian) people? From a more philosophical viewpoint, can these differences in prototypes be kept apart from differences (or identity) in beliefs about the essential nature of conceptual categories? Time may tell.

Acknowledgements: The author would like to thank Greg Murphy and the Editors for helpful comments on this chapter.

References

- Ashby, F. G., Alfonso-Reese, L. A., Turken, A. U., & Waldron, E. M. (1998). A neuropsychological theory of multiple systems in category learning. *Psychological Review*, 105(3), pp. 442–481
- Barsalou, L. W. (1987). The instability of graded structure: implications for the nature of concepts. In U. Neisser (Ed.), *Concepts and conceptual development: Ecological and intellectual factors in categorization* (pp. 101-140). Cambridge: Cambridge University Press.
- Bellezza, F. (1984a). Reliability of retrieval from semantic memory: Common categories. *Bulletin of the Psychonomic Society*, 22, pp. 324-326.
- Bellezza, F. (1984b). Reliability of retrieval from semantic memory: Noun meanings. *Bulletin of the Psychonomic Society*, 22, pp. 377-380.
- De Deyne, S., Verheyen, S., Ameel, E., Vanpaemel, W., Dry, M. J., Voorspoels, W. et al. (2008). Exemplar by feature applicability matrices and other Dutch normative data for semantic concepts. *Behavior Research Methods*, 40, pp. 1030-1048.
- Djalal, F. M., Hampton, J. A., Storms, G., & Heyman, T. (in press). Relationship between Extensions and Intensions in Categorization: A Match Made in Heaven?. *Journal of Experimental Psychology: Learning, Memory, and Cognition*.
- Frege, G. (1892). Über Sinn und Bedeutung. Translated in (P. Geach & M. Black, eds.) *Translations from the Philosophical Writings of Gottlob Frege*. Oxford: Blackwell, 1952.
- Gärdenfors, P. (2000). *Conceptual spaces: the geometry of thought*. Cambridge MA: MIT Press.
- Hampton, J.A. (1979). Polymorphous concepts in semantic memory. *Journal of Verbal Learning and Verbal Behavior*, 18, 441-461.

- Hampton, J.A. (1988). Overextension of conjunctive concepts: Evidence for a unitary model of concept typicality and class inclusion. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 14, 12-32.
- Hampton, J.A. (1995). Testing the prototype theory of concepts. *Journal of Memory and Language*, 34, 686-708.
- Hampton, J. A. (2006). Concepts as Prototypes. In Ross, B.H. (Ed), *The Psychology of Learning and Motivation: Advances in Research and Theory*, 46, pp. 79-113.
- Hampton, J.A., Dubois, D., & Yeh, W. (2006). The effects of pragmatic context on classification in natural categories. *Memory & Cognition* 34, pp. 1431-1443
- Hampton, J.A., & Passanisi, A. (2016). When intensions don't map onto extensions: Individual differences in conceptualization. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 42, 505-523.
- Lakoff, G. (1987). *Women, Fire, and Dangerous Things*. Chicago: Chicago University Press.
- Levitis, D.A., Lidicker, W.Z. Jr., & Freund, G. (2009). Behavioural biologists don't agree on what constitutes behaviour. *Animal Behavior*, 78, 103-110. doi:10.1016/j.anbehav.2009.03.018.
- McCloskey, M., & Glucksberg, S. (1978). Natural categories: Well-defined or fuzzy sets? *Memory & Cognition*, 6, pp. 462-472.
- McNamara, T. P. & Sternberg, R. J. (1983). Mental models of word meaning. *Journal of Verbal Learning and Verbal Behavior*, 22, pp. 449-474.
- Murphy, G. L. (2002). *The Big Book of Concepts*. Cambridge, MA: MIT Press.
- Nisbett, R. E., & Wilson, T. D. (1977). Telling more than we can know: Verbal reports on mental processes. *Psychological Review*, 84, 231-259.

Putnam, H. (1975). The Meaning of "Meaning". *Minnesota Studies in the Philosophy of Science*, 7, 131-193.

Rey, G. (1983). Concepts and Stereotypes. *Cognition*, 15, 237-262.

Rosch, E. R. (1975). Cognitive representations of semantic categories. *Journal of Experimental Psychology: General*, 104, 192-233.

Rosch, E. R., & Mervis, C. B. (1975). Family resemblances: studies in the internal structure of categories. *Cognitive Psychology*, 7, 573-605.

Storms, G. (2004). Exemplar models in the study of natural language concepts. *The Psychology of Learning and Motivation*, 45, 1-39.

Storms, G. (2004). Exemplar models in the study of natural language concepts. *Psychology of Learning and Motivation-Advances in Research and Theory*, 42, 1-40.

Verheyen, S., & Storms, G. (2013). A mixture approach to vagueness and ambiguity. *PLoS ONE* 8(5): e63507.

Wittgenstein, L. (1953). *Philosophical Investigations*. New York: MacMillan.