

Introduction

Categorisation of objects is usually 'effortless' - but if an object is unfamiliar it can be very difficult to find the right category or concept.

Concepts, categories and words

Concept - a general idea formed in the mind - 'general' = concepts apply to every one of a class of things - a category.

e.g. Concept of 'cat' is a general idea about all cats.

Concepts are internal to the mind, but the categories that concepts are about are external.

Concepts can map to more than one category - e.g. the concept of 'chest' could relate to a category of body parts or items of a category of furniture.

Several concepts can map to the same word - e.g. Cat (concepts of domestic cat, feline, big cat ...) Link between concepts and words is complex and not simple ... but for the purposes of this chapter assumption is that words pick out concepts in a straightforward way.

Categorisation

Bruner et al - categorising = make different things equivalent; group objects/events/people around us into classes; respond in terms of class membership, not uniqueness.

Implies concepts at work when similarities in behaviour shown towards different objects and when difference in behaviour are shown to different objects.

e.g. pat two different dogs = similar behaviour = same category; pat a dog and not a plant = different behaviour = different category.

Chapter 5 - Concepts

[Discursive approach to categorisation - e.g. **Potter and Wetherell** - subtleties in how categories shown based on social context and categorisation not tied to language - e.g. **Sappington and Goldman** - in horses: neither aspect considered in this course - cognitive psychology interested in what can be generalised from categorisation processes in humans.]

One technique used to investigate categorisation:

Sorting tasks - e.g. **Ross and Murphy** - categorisation of foods (on cards.) Sometimes eggs go into the same category as other breakfast foods; other times with dairy products. Groups indicate the underlying concept.

Eggs in different categories consistent with **Barsalou** - categorisation reflects a person's goals - e.g. if you were asked to group things together that you'd rescue if a fire happened, they wouldn't be things you'd normally put in the same category.

The wider story of concepts

Concepts are usually implicit in our behaviour, but at times are explicitly discussed. E.g. **Eco** - documents the discovery of the platypus and its eventual categorisation as a monotreme. Demonstrates the complexity of the process of categorisation.

E.g. 2 - legal and moral concepts - in the UK, laws are applied differently to adults and children - therefore, as it is difficult to identify the boundary between adulthood and childhood, and arbitrary categorisation has to be made by parliament.

E.g. 3 - clinicians - **Macintyre** - diagnosis of ME based on one major criterion plus 4 out of 8 minor criteria being present - diagnosis is a form of categorisation.

Concepts and cognition

w.r.t. recognition [chapter 4] - 'semantic classification' is what concepts are about. So classifying an animal as a 'cat' can be viewed as a further type of recognition.

Concepts are the building blocks of semantic memory - e.g. the persistent knowledge that 'cats are animals' Semantic memory expresses relationships between concepts [chapter 8].

Some concepts are lexical concepts - a concept for which there is a single word - e.g. concept of 'cat' represents what we believe the word 'cat' to mean.

Concepts have a role in reasoning - e.g. can infer that a cat has a heart, as we know a cat is a mammal and mammals have hearts. Concepts simplify the task of remembering information.

Explaining Categorisation

Similarity 1: Classical view of concepts

Aristotle, Sutcliffe (1993): things belong to categories as they have common properties.

(i) If something is a category member, then it must have properties common to the category's members.

(ii) if something has the properties common to a category's members, then it must be a member of that category.

e.g. bachelors.

Support for this view from empirical investigations - e.g. **Bruner et al** - showed people categorise according to if instances possessed necessary and sufficient conditions of a category.

However, this view can be criticised in a number of ways.

(i) Typicality

Rosch - by asking participants rating of typicality (goodness of exemplar) of instances in a category, demonstrated that not all members of a category are equally good examples (as would be expected by the classical view.) e.g. some dogs are more doggy than others. Rating scale (1-7) used.

Example ratings from **Rosch's** work:

<u>Fruit</u>		<u>Vehicle</u>	
Apple	1.3	Car	1.0
Plum	2.3	Scooter	2.5
Pineapple	2.3	Boat	2.7
Strawberry	2.3	Tricycle	3.5
Fig	4.7	Skis	5.7
Olive	6.2	Horse	5.9

Rosch => ratings are indicative of the internal structure of categories.

Barsalou - typicality does not correlate with familiarity of an instance and is only weakly correlated to frequency. Penguins are more atypical birds than robins and would remain so even if we saw lots of penguins.

Rips et al and **Rosch** - examined the relationship between typicality and time it takes participants to verify a sentence making a category judgement - category or sentence verification. Highly typical sentences (a robin is a bird) are verified more quickly than atypical sentences (a penguin is a bird.)

More support for categories having an internal structure - **Rosch and Mervis** - from the property or attribute listing method. Results show less typical instances shared properties with fewer category member - e.g. Robins have properties (flight, eats worms, builds nests) shared with many other birds; Penguins (swims) - shared with few others.

This research therefore strongly challenges the explanatory power of the classical view by providing evidence that categories have a rich internal structure. However, as the classical view does not say anything about internal structure, if it exists, it doesn't disprove the classical view unless internal structure reflects category membership!

If typicality implies graded membership of a category, then the classical view is wrong. If it does not, it may be compatible.

(ii) Borderline cases

If category membership is 'all or nothing', then there should be no borderline cases.

Intuitively, there are categories (such as colours) that do have borderline cases - e.g. at what point does red become orange?

McCloskey and Glucksberg - evidence supporting this intuition for a wide range of categories. Method based on asking for categorisation judgements.

Yes/No responses to questions required (is a robin a bird?). Some instances were also rated for typicality. The researchers then considered the level of agreement shown in their judgements, between participants and their own judgements when re-tested.

Good agreement found on highly typical and atypical instances; but disagreements across time and individuals for intermediate instances - e.g. bookends as an item of furniture.

However, disagreement may just reveal a lack of knowledge or perspective-dependence aspects (e.g. tomato usually used as a vegetable but it is a fruit) rather than that the classical view is wrong (there should be no borderline cases.) However, not obvious that **McCloskey and Glucksberg's** examples involved either - so appears to undermine the classical view.

(iii) Intransitivity of categorisation

Transitivity - if A is in B, and B is in C, then A is in C.

Hampton - people's categorisation judgements are not consistent with transitivity - e.g. car seats are a kind of chair; chairs are a kind of furniture; but car seats are not a kind of furniture.

(iv) Lack of definitions

Wittgenstein - 'games' example - classical view would expect there must be something in common between card games, Olympic games, board games, ball games etc. - but his position is that when you look for common properties there are none.

His argument is that most categories are indefinable.

However, does not prove that natural categories cannot be defined, so someone may be able to provide a definition eventually. However:

Putman, Kripke - 'all cats are robots controlled from Mars' thought experiment - if true, a cat would no longer be a mammal (an assumed property), so we would be wrong about a key defining property of a cat is that they are mammalian. But, it would still be a cat even after such a discovery. Therefore, we can be wrong about virtually any defining property of a category.

Similarity II - prototype theory of concepts

Idea: concepts are organised around the "central tendency" of a category - the **prototype**. May be a real instance, but usually a 'best' member, formed by aggregation of characteristics. **Rosch** - belief that this happens in nature - certain attributes cluster together, prototypes describe these natural clusters. E.g. prototype for 'bird' may have properties such as feathers, wings, beak, ability to fly and cluster together the way that feathers, lips, gills and swings through trees do not.

Category membership is therefore determined by how similar it is to the prototype.

Some similarity to the classical view - except that an instance doesn't need to match on all properties that are distinctive of a category.

Similar to the way clinicians diagnose ME for example - different properties can carry different weights - and so prototype theories readily explain the typicality effects discovered by **Rosch et al.**

1. Instances differing in typicality are assumed to differ in terms of weighting of the value on which they match the concept.

2. 'A robin is a bird' is more easily verified than 'a penguin is a bird' as most attributes that match for a robin have more highly weighted properties than those for a penguin - for which more matches will be needed before the criteria for category membership is reached.

3. Typicality will correlate with how widely the members of the category share attributes.

However, there are still challenges to this approach:

(i) The meaning of typicality effects

Armstrong et al - found typicality effects occur for concepts that are definitional - e.g. 'odd number', 'even number', 'plane geometry figure': a square is 'more typical' than a triangle and than ellipse. His conclusion was not that such concepts are organised around a prototype, but that the existence of typicality effects cannot be taken as conclusive evidence that membership is determined by similarity to a prototype.

Proposed a dual-process model to explain the results - concepts possess a 'core', used to judge category membership and a set of procedures we use to (rapidly) identify instances of a category.

Classical view = 'core'; prototype theory = 'identification procedures' ... but therefore inherits some of the problems faced by each approach!

(ii) Context-sensitivity of typicality effects

Problem - typicality effects appear to change with context - if prototype theories are a good explanation for categorisation, then stability would be expected in the prototype.

Roth and Shoben - typicality effects changed by linguistic context. E.g 'B. enjoyed riding/milking the animal; the ___ quite liked it too' - typicality of the animal words used to complete the sentence depended on the verb - e.g. horse/cow. [see also chapter 2 - priming effects]

Medin and Shoben - spoon example. Small spoons rated more typical than large spoons; metal spoons more typical than wooden ones. However, large wooden spoons rated more typical than small wooden spoons.

The contribution of 'large' and 'small' to typicality depends on if the spoon is made of metal or wood.

Instability is a significant difficulty for prototype theory - as (i) it is at odds with **Rosch's** argument that prototypes = clusters of correlated properties reflecting the natural world; (ii) not clear what mechanism is at work if weighting of values or diagnostic power of attributes are changeable; (iii) **Medin and Shoben** - contributions to typicality of different properties are mutually dependent - yet prototype theories assume attributes and their values are independent of each other.

(iii) Complex concepts

A common assumption is that concepts express what we understand a word to mean - e.g. concept of 'red' is expressed by the word red and a concept of 'car' by the word car.

If complex concepts are combinations of lexical concepts, then how is 'red car' explained?

If concepts are structured around prototypes, they could combine by combining their prototypes - which might work for 'red car' but not for 'pet fish' or 'stone lion' or 'killer firework'. **Fodor** - complex concepts create real problems for most theories of concepts.

Common-sense theories - the theory based view

Classical and prototype theories are similar as they explain concepts in terms of simple feature sets. Problems with these theories have led to questions about (a) the importance (or otherwise) of similarity to categorisation and (b) the role played in categorisation of larger knowledge structures - the 'theory'-theory of concepts.

Similarity based accounts have had some success - e.g. **Hampton**, using **McCloskey and Gluckberg's** data for borderline cases found typicality was a good predictor explaining 46-96% of the variance in categorisation probability. However, other factors were good (but not as good as familiarity) - e.g. lack of familiarity; the extent to which something was 'technically speaking' a category member (a dolphin is a mammal but is more like a fish). This result therefore suggests categorisation draws on deeper knowledge.

(i) Problems with similarity

Goodman - it is 'a pretender, an impostor, a quack'.

In prototype theory, similarity is defined as 'sharing properties with' rather than 'is similar to' - so therefore it is not really similarity that is driving categorisation at all.

Another problem is there is no obvious limit to the number of properties any two objects may share. **Murphy and Medin** - lawnmowers and plums are similar as they weigh less than 10,000kg, did not exist 10m years ago, both cannot hear, both can be dropped

... so if similarity (as shared properties) is to explain categorisation, then a definition of what counts as a relevant property is required. **Murphy and Medin** therefore suggest similarity is shorthand for something else that makes categories coherent.

(ii) Role of common-sense theories

Murphy and Medin - concepts are explanation based (not similarity based.) e.g. category of 'intoxicated' unlikely to include the property 'jumps into the pool fully clothed at a party' - but instead, categorising someone as 'intoxicated' plays a role in explaining their behaviour.

Categorising a robin as a bird helps to explain (albeit partially) why it has feathers and a beak. Without categorising a robin as a bird we would need an explanation as to why it had feathers and a beak.

Similarity-based approaches fail to easily explain non-independence of attributes. **Murphy and Medin** argue relationships between attributes point to concepts being embedded in larger knowledge structures. As they believe many categorisation judgements become automated, the role of our underlying theories is hidden, so leading to a mistaken conclusion that categorisation is driven by similarity. Where we encounter novel instances (e.g. robot cats) then our underlying theories are used.

Support for this approach: **Rips** - objects belonging to distinct categories used (e.g. pizza, US quarter.) Largest size estimate for US quarter smaller than for the smallest size estimate for a pizza. Participants then asked to consider a third, intermediate size object - about which only the size was disclosed.

If asked which category the 3rd object was likely to belong to, participants gave a different response to being asked which the object was likely to be more similar to. The judgements dissociated - more likely to be a pizza; but more similar to a quarter.

Other dissociations between categorisation and similarity have been demonstrated - e.g. Kroska and Goldstone - showed participant a set of 3 phrases that described a putative emotion. E.g:

'Threat of harm or death' [Fear]
'Being accepted, belonging' [Joy]
'Experiencing pleasurable stimuli or sensations' [Joy]

Participants categorised scenario as 'fear'; but judged it to be more similar to an instance of 'joy'.

Category membership judgements - influenced by properties considered central to a category; similarity judgements influenced by characteristic properties.

Keil - support for knowledge of deeper causal principles at work in categorisation comes from looking at children's attempts to categorise.

Discovery and transformation procedures used.

Discovery - children told of a novel hybrid looking/behaving like a zebra; but had horse insides, and was the offspring of two horses.

4y.o. - animal is a zebra (influence is superficial chars)
7y.o. - animal is a horse (influence is biological relevant properties.)

Similar results found using a transformation procedure - raccoon undergoes a series of transformations to make it look and behave like a skunk. 4y.o. - animal is a skunk; 7y.o. animal is a raccoon.

Age-related change in categorisation found - 'the characteristic-to-defining shift'.

Murphy argues probably not the case that older and younger children have qualitatively distinct ways of categorising; more likely younger children do not know enough about biological categories to use them.

(iii) Difficulties with the 'theory'-theory

Some of the research doesn't necessarily support the 'theory'-theory - rather, it just suggests that similarity isn't the explanation for categorisation.

Not clear what is meant by 'theory' - difficult to formalise (in a way similarity is not -e.g. tables showing the weightings of attributes in prototypes.)

Murphy - background knowledge that influences concepts is too simplistic to be a proper scientific 'theory' - but is simply 'knowledge'.

Hard to see how combining theories can explain complex concepts.

Simply replaces difficulties with the notion of similarity with the mysterious notion of a 'theory'.

Psychological essentialism

Medin, Medin and Ortony - "people believe and act as though category members have certain essential properties in common."

Essential properties may constrain superficial properties -e.g. genetic makeup of an animal constrains its appearance, behaviour ... (dangerous dogs?)

Essential property - if an object did not have them, it would not be that object. Like defining properties in the classical view, but most people will not know what they all are for a category - merely believe that it has some.

Psychological essentialism proposes concepts contain a 'placeholder' for an essence. It could be empty, if the person's knowledge is not sufficient to know what they are. Not everyone's placeholder for the same category is empty - e.g. a chemist may know the essential properties of gold.

This approach is consistent with most of the evidence used to support the 'theory'-theory. Specific support comes from how categories develop.

E.g. **Gelman and Wellman** - 4/5y.o. believe insides of objects more important than outsides in determining category membership - a dog is still a dog if you remove its outside, but not its inside. Children are therefore essentialist - something internal explains category membership.

P.E. challenged by **Malt** - if H₂O is the 'essence of water', then categorisation of liquids as water should be influenced by how much water something contains.

Not so - pond water = water (78.8% H₂O); but tears != water (88.6% H₂O). However, categorisations influenced by the source of the water, location and its function. Presence or absence of H₂O therefore not the only factor in categorising a liquid as water.

Braisby et al - 'robot cat' thought experiment can be used to investigate P.E.

Statements expressing essentialist intuitions:

(a) 'Tibby is a cat, though we were wrong about her being a mammal'

and the contrary intuition

(b) 'Tibby is not a cat, though she is a robot controlled from Mars'

50% participants thought (a) true and (b) false; many gave contradictory judgements (a) and (b) both true or false.

Braisby and Franks - these results show P.E. not supported, but that implied concepts change their content according to context and perspective.

Malt - people are psychologically essentialist for natural categories, but not for others. Experts judged

to be in a better position to categorise a 'halfway object', such as a robin-sparrow or trout-bass (75% would ask the expert) but for objects like ship-boat 63% said you could call it whatever you wanted to.

Braisby - evidence from **Malt's** experiment not conclusive - examined the extent to which categorisation judgements are modified for genetically modified biological categories.

Only half of participants changed their initial category judgements to match those of an 'expert' for genetically modified instances. 25% would modify their judgement to fit in with 'shoppers' - so **Braisby** argues only 25% modify their judgements because it is expert opinion - so this doesn't provide evidence for P.E.

Much evidence said to support P.E. only relates indirectly to beliefs in essential properties. E.g. **Gelman and Wellman** result - a further inference is required to relate the insides (of a dog) to essences.

Similarly, **Stevens** - notion of essence/essential properties is not required to explain empirical data from **Gelman and Wellman** expt - not a parsimonious explanation.

P.E. also doesn't help explain complex concepts, like other theories.

Where next?

None of the theories seem good at explaining concepts. So what sense can be made from them and the evidence generated?

(i) Is all categorisation the same?

The different explanations perhaps imply categorisation is not a single process. For example, to explain the general robustness of categories against the evidence that people categorise differently in different circumstances, it could be argued there may be a determinate number of different kinds of

categorisation, which could be framed by the different theories of concepts.

e.g. Classical view useful where definitions are required; prototype view = rapid categorisation required; theory-based views = categorisation that require considered judgements -e.g. to explain something; P.E. = categorisation when we want consistency with a scientific/expert view of the world.

Smith and Sloman - some support by replicating the dissociation between similarity and categorisation judgements found by **Rips**. They got the same dissociation only if participants required to think out loud. Argue result indicates 2 categorisation modes: (1) Similarity based; (2) Rule-based. We use one or the other depending on how a categorisation task is presented to us.

Similarity mode - categorisation conforms to similarity based accounts.

Rule mode - categorisation conforms to theory or explanation based.

Suggests categorisation is not a single process.

Also, we label things (e.g. a stone lion as a lion) that imply category membership, even though we don't believe they are members. **Malt et al** - same true for containers - a 'bottle' is not a bottle just because it looks like a prototypical bottle, but because things have historically been labelled so - e.g. shampoo bottle.

(ii) Are all concepts the same?

Well-defined categories (e.g. even number) are amenable to definition => classical view is a good explanation. (but doesn't explain typicality effects.)

Fuzzy categories - (e.g. red), with borderline cases, are amenable to prototype theories. 'Chair' may be fuzzy in the same way.

Theory-based/essentialist approaches best suited to categories for which 'common-sense' theories are common - e.g. sparrow, gold - and match well with categories for which scientific theories have been developed.

Essentialism may also be a good explanation of people's social categories (e.g. introvert) - a deeper, explanatory knowledge is required of the causal principles of category membership.

(iii) Are all categorisers the same?

Medin et al - participants from three different occupational groups; different experience and knowledge of trees. Asked to categorise 48 different trees.

Taxonomists - reproduced a scientific way of sorting the trees

Maintenance workers - gave more emphasis to superficial characteristics - e.g. if a tree was broad leaved or not. Included a 'weed' tree - ones that are particularly difficult to maintain.

Landscapers - factors like landscape utility, size, beauty.

Lynch et al - typicality ratings of trees varied between experts and novices. Ratings of experts influenced by similarity to ideals - best examples were of extreme height rather than average height; novices influenced largely by familiarity.

These studies suggest people do not necessarily categorise things in the same way; extent of knowledge influences categorisation and so this is reflected in their concepts.

Tyler and Moss (offprint)

Fundamental human ability is our formation of domains and categories.

One important distinction we make is between inanimate and living objects. Happens very early in life.

Neuroimaging studies show selective activation in cortical regions associated with different domains and in brain-damaged patients with deficits in a single domain (e.g. living/non-living things) or category (e.g. animals, fruits, tools) - category specific semantic deficits.

Data from brain damaged patients therefore suggests that different parts of the brain are used to hold information about different domains/categories - BUT - most patients with 'living things' deficits have antero-medial temporal lobe damage. HOWEVER, in normal patients, this region is not activated in neuroimaging studies.

Demonstrates further that a unified theory of categorisation is some distance from being found, if such a unifying theory exists at all.

Link to connectionism:

Prototype theory - support from **Rumelhart and McClelland** pattern associator models - e.g. see many dogs, can 'recognise' a prototypical dog that has not been seen.

Conclusion

Category knowledge multi-layered, covering:

Knowledge of causal properties;
Knowledge relevant to explaining membership;
Knowledge of properties of instances;
Knowledge of function;
Knowledge of superficial properties

We can call on different types of category knowledge at different times for different purposes.

Competing theories offer good explanations of different sets of phenomena; but this does not rule out their being a single, unifying theory that we have not yet discovered.

Improving understanding of categorisation and how it operates essential to us being able to develop a good understanding of the mind.