Bayesian Modelling of Perceptual Categorisation in Conceptual Spaces

Nina Poth[0000−0002−0216−694X]

Institute for Philosophy II, Ruhr-Universität Bochum, 44801 Bochum, Germany
nina.poth@rub.de

Keywords: Perceptual categorisation, Bayesian inference, Conceptual Spaces

1 Introduction

Why do we categorise two objects (e.g., a pear and a banana) as instances of the same concept (e.g., the concept FRUIT), despite their perceptual differences? This is the two-step decision problem of perceptual categorisation (PC). Accordingly, the first step is to infer the 'correct' concept (e.g., FRUIT, as opposed to EDIBLE) that subsumes each object. The second step is to decide whether the second object falls under the same or a different concept[1]

I argue that Bayesian modelling [6, 8] combined with the Conceptual Spaces framework of knowledge representation [3] provides an optimal solution to the problem of PC. Bayesian modelling of PC offers an elegant approach to the rationality of PC, but it implicitly relies on an account of psychological similarity representations.

2 Bayesian modelling of PC

On the Bayesian approach, PC is a rational inference task that involves a set of countable, possibly infinite, and mutually exclusive hypotheses, $h_1, \ldots, h_n, \ldots \in \mathcal{H}$. Each hypothesis pairs a given set of stimuli, $x_1, x_2, ..., x_n$, with a concept, C. $\mathcal{H}$ defines a probability distribution, $pr : h \to [0, 1]$. The assignment of probabilities to hypotheses is conditional on the given evidence, $e : x_1, ..., x_n$. To model the agent's degree of belief in h after observing e, we use Bayes' theorem:

$$pr(h|e) = \frac{pr(e|h)pr(h)}{\sum_{h' \in \mathcal{H}} pr(e|h')pr(h')},$$

[1] One reason to characterise PC as a two-step problem is that some agents seem to generalise behaviour on the basis of nonconceptual contents [2]. On this basis, inferring the concept and generalising a response are logically independent steps. The Bayesian approach that I consider rests on the assumption that the first step is involved.

where the posterior probability, $pr(h|e)$, is a function of the product of the likelihood, $pr(e|h)$, and prior, $pr(h)$. This is taken relative to the sum of the products of likelihoods and priors for all alternative hypotheses, $h^{'} \in \mathcal{H}$. The likelihood represents how probable it is to observe e, given h was true. The prior represents the probability of h regardless of e.

A partial answer to the first step (of inferring the correct C that subsumes x) is given in Brössel's (2017) recent account of the rational relations between colour perception and beliefs about colour, in which the correct C is that which covers x in conceptual space.

To answer the second step, we need to determine under what conditions it would be correct for the agent to treat y like x. I evaluate two plausible generalisation policies: Shepard's (1987) *weak sampling* policy and T&G's (2001) *strong sampling* policy. On weak sampling, the extent to which agent S should treat y like x depends fully on whether or not y falls under the same concept as x (i.e., an all-or-none policy). Under strong sampling, it depends on the size principle: agents should assign greater probabilities to hypotheses that pair x and y each with a small concept (see [5, pp. 306-307] for an illustration).

Strong sampling seems preferable to model linguistic communication because it allows for hypothesis selection when multiple hypotheses are compatible with the evidence [4]. For example, it explains why a pear and a tomato are relatively unlikely to fall under the same concept. A hypothesis citing FRUIT will be better confirmed than a hypothesis citing EDIBLE for an encounter of the pear. Likewise, a hypothesis citing VEGETABLE will be more probable than hypotheses citing EDIBLE for an encounter of the tomato. Consequently, it is more likely that the pear and the tomato will be treated differently because their preferred concepts (FRUIT vs. VEGETABLE) are distinct. This is so even if they both fall under EDIBLE (which is all that would count in weak sampling). Furthermore, strong sampling can account for efficient learning, as only a few examples are needed to confirm such hypotheses [8].

However, Bayesian modelling leaves important questions unanswered. Firstly, neither weak nor strong sampling provide an account of concept individuation. So what determines whether the instances belong to the same or to different concepts? Secondly, why are smaller hypotheses better confirmed by the examples than larger hypotheses? Thirdly, which hypotheses are a priori (im)plausible? For example, in Goodman's (1955) classical riddle of induction, all Emeralds observed to be green fall under GREEN but also under GRUE. Still, GREEN is preferable because it is more natural. Just considering the size of hypotheses cannot account for this intuition, since these concepts seem to have the same size. I argue that Conceptual Spaces provides possible answers to these questions.

3 Conceptual Spaces and knowledge representation

Conceptual Spaces structures knowledge representations based on similarity relationships; similarity is a monotonically decreasing function of distance in a geometric space that is endowed with dimensions that represent perceptual qualities such as tastes or colours [3]. An instance is a vector in this space. A concept is a region in this space, and its area represents the concept's intension. This is the average dissimilarity among possible instances of the concept. For example, instances of pears and instances of bananas both fall inside the FRUIT region. However, those pears that fall inside the RED portion of the colour domain will be less similar to bananas than those pears that fall inside the YELLOW portion. Conceptual Spaces relates the evidence to hypotheses in virtue of the relative positions among points and regions. Therefore, what determines whether the instances belong to the same concept is how close they are in psychological similarity space.

On this basis, we can explain why smaller hypotheses are better confirmed by the examples than larger hypotheses. Roughly, smaller concepts allow for fewer variations in the outcomes of a random sampling experiment. If we observe a sample with a small variation of outcomes, this observation is likely to be a consequence of a small concept. Correspondingly, the confirmatory relations between instances and probabilistic hypotheses reduce to the overlap between concepts and the areas covered by their instances. Among those concepts that subsume the available evidence, the possibilities provided by smaller concepts (e.g., FRUIT) will overlap more closely with the actual examples than those provided by larger concepts (e.g., EDIBLE).

Finally, the structure of the space constrains the prior selection of hypotheses. Both [6] and [3] assume that natural concepts are *convex* regions. Alternatively, [7] argues that especially complex concepts are natural to the extent that they are *characteristic* for instances that fall under them. In either case, the topological structure of the space rejects many concepts a priori, making Bayesian modelling of PC more psychologically plausible.

References

1. Brössel, P.: Rational relations between perception and belief: the case of color. Review of Philosophy and Psychology, **8**(4), 721-741 (2017).
2. Deroy, O.: Categorising without concepts. Review of Philosophy and Psychology **10**(3), 465-478 (2019).
3. Gärdenfors, P.: Conceptual spaces: The Geometry of Thought. MIT Press, Cambridge MA (2000).
4. Lewis, M. L., & Frank, M. C.: Still suspicious: the suspicious-coincidence effect revisited. Psychological science, **29**(12), 2039-2047 (2018).
5. Perfors, A., Tenenbaum, J. B., Griffiths, T. L., & Xu, F.: A tutorial introduction to Bayesian models of cognitive development. Cognition, **120**(3), 302-321 (2011).
6. Shepard, R. N.: Toward a universal law of generalization for psychological science. Science **237**(4820), 1317-1323 (1987).
7. Strössner, C.: Criteria of Naturalness in Conceptual Spaces. Manuscript submitted for publication.
8. Tenenbaum, J. B., & Griffiths, T. L.: Generalization, similarity, and Bayesian inference. Behavioral and brain sciences **24**(4), 629 (2001).
9. Goodman, N.: The new riddle of induction. In: Fact, Fiction, and Forecast. Bobs-Merrill (1965).