Criteria of naturalness in conceptual spaces

Corina Strößner

Ruhr University Bochum, Universitätsstrasse 150, 44780 Bochum, Germany
corina.stroessner@rub.de
https://www.corinastroessner.com/

One of the central aims of conceptual spaces theory [5] is to find criteria of what makes a concept natural, that is, easy to process and cognitively useful. The most prominent proposal is the convexity thesis. It has been proposed by Gärdenfors as early as 1990, when he states “a property, that is, a region of a conceptual space, is natural only if the region is convex” [4, p. 88], which means that any point between two points of the region belongs to it as well. The criterion has been repeated in many of his other works, often in different variations. Most supporting arguments for convexity requirement have been established in domains, that is, sets of closely related dimensions. A classical example of such a domain is the colour space with the dimensions HUE, SATURATION, and BRIGHTNESS. Unfortunately, domain-specific concepts like colours, smells, shapes are not the primary examples of natural concepts. In particular, they are not able to capture natural correlations, as most noun concepts (e.g., APPLE) do. The main aim of the paper is to develop criteria of naturalness for multi-domain concepts with respect of the representations of features that are associated to the concept and the probabilistic correlation patterns which are captured by it. I will discuss three criteria: the existence of characteristic features, the coverage of inhabited areas of conceptual spaces and the shape of the probability distribution in conceptual regions.

Characteristic features

Natural multi-domain concepts are recognizable by *characteristic features* in several distinct conceptual spaces. These are highly associated to the concept and allow to categorize instances. For example, there is a specific smell of coffee which allows to recognize it. In this sense, coffee has the characteristic feature of coffee smell. The sentence ‘Coffee typically smells like coffee’ sounds hopelessly circular but a representation of the smell in a conceptual space is not a word but represents a perception that is strongly associated to a concept but does not depend on it. I do not demand that characteristic features are necessarily convex. Convexity is without doubt important if one aims to partition a domain into concepts, i.e., the taste space into taste concepts. However representations of features like THE TASTE OF A STRAWBERRY are not like property concepts (SWEET or BITTER), and primarily restricted by the multi-domain concept and the correlations it captures. In order to be cognitively perceivable as one feature, it is still reasonable that feature representation are topological restricted in some way but the more important demand is that the region corresponds as closely as possible to the concept.

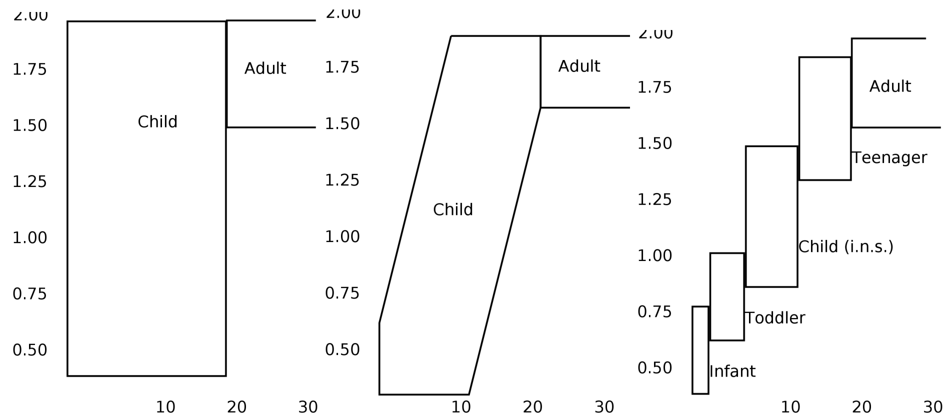

(a) (b) (c)

Fig. 1. Fig. 1a shows a convex representation. Fig. 1b represents the concept of child with a correlation of age and height. In 1c, fine-grained convex concepts capture the correlation.

Inhabitedness and correlations

The demand that a concept has several characteristic features implies that there is a strong correlation between these features. Apple shapes are correlated with apples tastes. This brings me to the second aspect of multi-domain concepts: the capturing of correlations. When discussing how concepts capture correlation, a combination of probabilistic arguments and conceptual spaces is needed. Bechberger and Kühnberger [1] argued that a representation of correlation conflicts with the convexity criterium. For example, the fact that height and age of children are correlated is best captured by a representation of CHILD that is not convex, as seen in Fig. 1b compared to Fig. 1a.[1] Their representation avoids the inclusion of uninhabited regions, such as a 1.50 m tall one-year-old, and thus captures the intuition that a natural concept should only cover inhabited regions: If S and P are regions in a conceptual space CS and $S \subset P$ and the probability mass captured by S and P is the same (i.e., $Pr(S) = Pr(P)$), then S is preferred over P. As the example of [1] illustrates, there is a tension between narrowness and convexity if S is non-convex and P is convex. According to the convexity criterion, P is preferable but the principle of narrowness demands that non-inhabited parts of P should not belong to the conceptual representation. This conflict is resolved if a correlation is well captured *by* a concept. For example, INFANT captures a correlation between very young age and many physical and mental attributes such as size. As seen in Fig. 1c, a finer distinction within minors reconciles narrowness and convexity.

[1] Both representations would be convex in Euclidean space. However, SIZE and AGE are arguably not integral dimensions and should be represented in a Manhattan space where only cuboids are convex.

Unimodality

Having a probabilistic conceptual space allows to specify criteria of natural concepts not only in terms of the topology of a conceptual region but also the shape of its probability distribution. It is well established that humans are strongly biased towards unimodal probability distributions, that is, distributions which have not more than one local maximum. Categories, where two extreme values are more likely than medium values, such as in a U-distribution, are extremely difficult to learn and represent [3].

A simple generalisation of unimodality in n-dimensional spaces is *quasi-concavity*. A probability density f is quasi-concave iff for each $t > 0$, $\{x|f(x) \geq t\}$ forms a convex region [2]. Building on quasi-concavity as a generalization of unimodality, I thus propose the following criterion of conceptual naturalness in probabilistic conceptual spaces: *A natural concept, represented in a conceptual space CS with a probability density f, is a region $S \subseteq CS$ such that f is quasi-concave within S, meaning that for every $t > 0$, $\{x \in S|f(x) \geq t\}$ is a convex region.* The criterion formally captures the metaphor that natural concept carve nature at its joints, where joints are areas where the probability density is comparably low. One the other hand, the criterion is obviously related to the convexity assumption from conceptual spaces and thus connects two traditions of thinking about naturalness, namely convexity in conceptual spaces with the idea that concepts capture peaks of natural covariation.

References

1. Bechberger, L., Kühnberger, K.U.: Formalized conceptual spaces with a geometric representation of correlations. In: Kaipainen, M., Zenker, F., Hautamäki, A., Gärdenfors, P. (eds.) Conceptual Spaces: Elaborations and Applications, pp. 29–58. Springer International Publishing, Cham (2019).
2. Dharmadhikari, S., Joag-Dev, K.: Unimodality, convexity, and applications. Elsevier (1988).
3. Flannagan, M.J., Fried, L.S., Holyoak, K.J.: Distributional expectations and the induction of category structure. Journal of Experimental Psychology: Learning, Memory, and Cognition **12** (1986).
4. Gärdenfors, P.: Induction, conceptual spaces and AI. Philosophy of Science **57**(1), 78–95 (1990).
5. Gärdenfors, P.: Conceptual Spaces: The Geometry of Thought. MIT Press, Cambridge, MA (2000).
6. Jäger, G.: Natural color categories are convex sets. In: Logic, language and meaning, pp. 11–20. Springer (2010).