

category-markers in a host of contexts. Their immutability and consistency give them the power to support the most stable systems of classification.

What Mutability Is Not

Of course, features differ on dimensions other than mutability, and those other dimensions can also influence the impact of a feature on different tasks. As a means of clarifying our notion of mutability, we distinguish the mutability of a feature in a concept from its category centrality, its diagnosticity, and its salience. These dimensions are all intimately related; nevertheless, they are different both conceptually and empirically. We start with their conceptual differences.

Category Centrality

Concepts and categories are, to a large extent, flip sides of the same coin. Roughly speaking, a concept is an idea that characterizes a set, or category, of objects. We construe the terms “concept” and “category” to refer to two different perspectives that a subject can take on a class of objects, what Tversky and Kahneman (1983) call the “inside” and the “outside” views. The inside view regards the internal structure of a concept, its features and what binds them together. The outside view regards some or all of the instances that are believed to be included in the category referred to by the concept. We take mutability to be determined by centrality in a concept’s feature space—in an inside view of the concept. Other measures, like naming and variability judgments, are largely determined by the boundaries of a category’s instance space—in an outside view of the concept. Because the inside and outside views of a class of objects are usually compatible, measures of conceptual and category structure usually coincide. As we will see, however, situations can be constructed to make them diverge.

Diagnosticity

The term “diagnosticity” has multiple senses. Sometimes it is used to refer to the informational value of a feature for one category relative to a set of categories. If one’s task is to categorize different kinds of apples, color may prove highly diagnostic because it excels at distinguishing some types of apples from others. This sense of diagnosticity is well-captured by the likelihood ratio $P\{C_1|F\}/P\{C_2|F\}$ which states that the informational value of a feature F for a category C_1 is high in proportion to the probability of the category given that we know the feature relative to the probability of some other category C_2 given the feature. When we want to know how much evidence F provides for C_1 alone, the likelihood ratio is monotonically related to $P\{C_1|F\}$, or what has been called cue validity (Rosch, 1978). Tversky (1977) demonstrated how a feature’s weight for determining similarity can change as a function of its informational value within a set of categories.

Informational value is clearly not identical to mutability because some attributes have high informational value but are mutable. For example, “having bones” has little informational value for robin because many things have bones but are not robins. However, it is an immutable feature of robin; imagining a robin that does not have bones but nevertheless