Journal of Experimental Social Psychology 52 (2014) 1–8

Contents lists available at ScienceDirect

Journal of Experimental Social Psychology

journal homepage: www.elsevier.com/locate/jesp

# Latent scope bias in categorization

Abigail B. Sussman [a,*], Sangeet S. Khemlani [b], Daniel M. Oppenheimer [c]

[a] University of Chicago Booth School of Business, 5807 S. Woodlawn Ave., Chicago, IL 60637, USA
[b] Intelligent Systems Section, Code 5515, Navy Center for Applied Research in Artificial Intelligence, Naval Research Laboratory, 4555 Overlook Ave. SW, Washington, DC 20375, USA
[c] UCLA, Anderson School of Management, 110 Westwood Plaza, Los Angeles, CA 90095, USA

## HIGHLIGHTS

- We investigate how people categorize exemplars given incomplete information.
- We define scope as the number of distinct features category membership implies.
- Results show bias for grouping exemplars in categories with narrower latent scope.
- Preferences extend to verbal and visual categorization tasks.

## ARTICLE INFO

*Article history:*
Received 22 January 2013
Revised 30 July 2013
Available online 7 December 2013

*Keywords:*
Narrow scope bias
Categorization
Stereotyping
Person perception
Causal reasoning

## ABSTRACT

Categories often have unobservable diagnostic features. For example, if a person is a lawyer, one might expect him to be both well dressed and knowledgeable about the law. However, without observing the person in a courtroom, one cannot tell whether or not he is knowledgeable about the law. How might we categorize the well-dressed person before we know whether or not he possesses a particular category feature? Two studies showed that, all else equal, individuals prefer to group exemplars into categories that specify fewer unobserved *and unobservable* features — i.e., those that have a narrower latent scope — to those with a broader latent scope. In Experiment 1, participants were more likely to classify novel exemplars as part of a social category that had a narrower latent scope in a verbal task. Experiment 2 demonstrated that the scope bias generalizes to contexts in which category structure is never explicitly specified.



## Introduction

The process of generalizing knowledge from a known category to a novel instance is central to the way we perceive the world, and it has permeated intellectual debate since Plato (Statesman, 261e et seq.). The mechanisms by which we categorize individuals constrain key components of social perception, such as stereotyping, impression formation, and even recall of information about others (e.g., Cantor & Mischel, 1979; Cohen, 1981; Higgins, Rholes, & Jones, 1977; Klein, Loftus, Trafton, & Fuhrman, 1992; Macrae & Bodenhausen, 2000; Stangor, Lynch, Duan, & Glass, 1992; Tajfel, Billig, Bundy, & Claude, 1971), as well as broader aspects of judgment and decision-making (for a review, see Murphy, 2002).

Given that people fit into many different categories, one tradition in person perception research has attempted to discern which categories are used for prediction (e.g., Crisp & Hewstone, 2007; Kunda, Miller, & Claire, 1990; Macrae, Bodenhausen, & Milne, 1995). Rather than making inferences from multiple possible categories, people tend to infer attributes based on the most likely category (Malt, Ross, & Murphy, 1995). It is therefore critical to understand how individuals determine the most likely category.

In laboratory studies on categorization, participants typically have complete information about which relevant features a putative category member possesses – participants are told that the member either possesses or does not possess a feature. However, this design is not paralleled in everyday life, where knowledge about an exemplar's features is frequently unknown or uncertain. The uncertainty complicates an already difficult categorization task. What strategies do people use to overcome informational limitations? Although categorization under uncertainty has received attention (e.g., Griffiths, Hayes, & Newell, 2012; Molden & Higgens, 2004; Murphy & Ross, 1994, 2005; Ross & Murphy, 1996; Verde, Murphy & Ross, 2005), there has been little study of how the structure of the category affects how people attempt to overcome missing or uncertain information.

Studies in explanatory reasoning suggest that people's knowledge of a category's causal structure may drive their categorization judgments. More than a decade of research has shown that both causal and explanatory reasoning play key roles in categorization (e.g., Ahn, Kim,

* Corresponding author at: University of Chicago Booth School of Business, 5807 S. Woodlawn, Chicago, IL 60637, USA.
*E-mail addresses:* asussman@chicagobooth.edu (A.B. Sussman), skhemlani@gmail.com (S.S. Khemlani), daniel.oppenheimer@anderson.ucla.edu (D.M. Oppenheimer).

0022-1031/$ – see front matter © 2013 Elsevier Inc. All rights reserved.
http://dx.doi.org/10.1016/j.jesp.2013.11.010

Lassaline, & Dennis, 2000; Murphy & Medin, 1985; Lombrozo, 2009; Rehder, 2003a, 2003b; Rehder & Hastie, 2001; Sloman, Love, & Ahn, 1998; Waldmann, Holyoak, & Fratianne, 1995). If explanatory reasoning and categorization recruit the same cognitive mechanisms, then an examination of the processes that underlie the generation and evaluation of explanations might help account for performance on categorization tasks along with consequences for stereotyping and inference.

A recent analysis of the role of causal structure in explanatory reasoning explored how people determine the best explanation for a set of observations when information is incomplete (Khemlani, Sussman, & Oppenheimer, 2011). The researchers identified a narrow *latent scope* bias. Latent scope describes the number of effects for which an explanation could potentially account, regardless of whether or not the effects are observable. People appear to prefer explanations with narrower latent scope. For example, consider the following:

A causes X and Y.
B causes X, Y, and Z.
We observed X; no information is known about Y or Z.
Which is more likely: A or B?

Khemlani et al. (2011) found that people prefer Explanation A in cases like this, because Explanation A causes fewer unobserved effects, even though these unobserved effects would not have been known even had they been present. In other words, when information is missing, people prefer explanations that make no predictions about items that are both unobserved and potentially unobservable (i.e., are not known to be either present or absent given the available evidence). In the context of categorization, the latent scope bias suggests that when the status of a set of features is unknown, people may prefer to align exemplars with categories that specify fewer features altogether, i.e., those that have narrower latent scope.

In this paper, we briefly review findings on the latent scope bias in explanatory reasoning and then provide empirical support for a parallel bias in the categorization of novel exemplars given limited information.

## Latent scope in explanatory reasoning and categorization

The latent scope of an explanation can be thought of as the number of distinct effects for which the explanation can potentially account. An explanation's scope is *latent* because the possible effects that it can describe may not necessarily materialize or be observable. Explanations that could account for fewer effects have narrower latent scope than those that account for many effects. For instance, contrast two explanations for why someone might dye his hair and then shave his head: he dislikes his new hair color, or he is going through a mid-life crisis. The first explanation has narrower latent scope; it can only explain how someone might behave with respect to his hair. The second explanation has a broader latent scope: going through a mid-life crisis could also account for a wide range of other behaviors, many of which may never materialize. If we assume equal proportions of people who dislike their hair color and who are going through a mid-life crisis, then there is no normative reason to prefer one explanation to the other.

However, Khemlani et al. (2011) show that individuals exhibit a strong bias in favor of the explanation with narrower latent scope. The effect was robust even when base rates favored the broad scope explanation, and it persisted in more naturalistic domains in which the structure of the explanation was not made explicit. Khemlani et al. (2011) also ruled out several explanations for the effect, including the possibility that participants interpreted the absence of information about an effect to mean that the effect was not present.

In explanatory reasoning, latent scope is defined as the number of effects the explanation could account for. In categorization, we extend the notion of latent scope to refer to the number of distinct features that category membership implies. The scope is latent because not all features of an exemplar may be observable when inferring category membership. For example, suppose that you observe someone who is knowledgeable about medications, and you have to decide whether she belongs to the category of pharmacists or physicians. If you do not know whether or not the person also has the ability to write prescriptions, and assume equal base-rates of pharmacists and physicians, would you be more likely to categorize the person as a pharmacist or a physician? The "pharmacist" category has narrower latent scope, since it specifies only one distinct features (e.g., is knowledgeable about medications), whereas the "physician" category has broader latent scope, because it specifies at least two distinct features (e.g., is knowledgeable about medications and has the ability to write prescriptions).

We next describe two studies that show that, parallel to the findings in explanatory reasoning, people prefer to place exemplars in categories with narrower latent scope. We first describe a test of the narrow latent scope hypothesis that used a verbal social categorization task before detailing an investigation of the bias in a visual category learning task. Both methods of examination reveal a robust narrow latent scope bias in categorization.

## Experiment 1

Experiment 1 tested the narrow latent scope hypothesis by first presenting participants with verbal descriptions of a variety of social categories and category members. Participants then performed a classification task. Problems relied on fictional categories to ensure that variations in category structure, rather than prior knowledge, caused any observed variations in judgment. The narrow latent scope bias predicts that people should tend to believe that the exemplar is a member of the category that specifies a more limited set of features. The effect of the bias is accordingly to leave fewer unobserved features unknown.

### *Method*

#### *Participants*

Forty-nine participants were recruited through Amazon.com's Mechanical Turk platform (for a discussion on the validity of results from this platform, see Paolacci, Chandler, & Ipeirotis, 2010) and participated in the study for monetary compensation.

#### *Design and procedure*

Participants were presented with a series of questions. In each question, participants received information about two otherwise unfamiliar categories. They were told explicitly that the categories had approximately equal numbers of category members. For example, participants were told about people who belonged to the Tokolo tribe:

*In the jungles of the Amazon about half of the Tokolo tribe members are hunters, and the other half are spear fishermen. Both hunters and spear fishermen carry spears, but spear fishermen also carry nets.*

In this example, the category of hunters is defined by fewer specific features (carrying spears) and is therefore considered to have narrower latent scope than the category of spear fishermen (who carry spears *and* nets). After participants were told about each category, they were informed about a specific person, for example, "You come across a tribesman who has a spear, but you don't know whether or not he also has a net." They were then asked to choose which category the person was more likely to belong to given the two alternatives.

Participants saw eight problems in total – four experimental and four control – each with a distinct set of categories (see Appendix A for problem forms and Appendix B for category descriptions). Across all problems, information about certain features was known and presented to the participant, and information about other features was unavailable. In the latter case, the fact that this information was unavailable was explicitly stated to the participant to clarify that the absence of information about a feature did not indicate the absence of the feature itself.