Teleology and generics

David Rose, Siying Zhang, Qi Han, Tobias Gerstenberg
{davdrose, syzhang6, qihan96, gerstenberg}@stanford.edu
Department of Psychology, Stanford University, USA

Abstract

Generic statements, such as "Bees are striped" are thought to be a central vehicle by which essentialist beliefs are transmitted. But work on generics and essentialism almost never focuses on the type of properties mentioned in generic statements. We test the hypothesis that teleological properties, what something is for, affect categorization judgments more strongly than behavioral, biological, or social properties. In Experiment 1, participants categorized properties as being either behavioral, biological, social, or teleological. In Experiment 2, we used the top four properties from each group to describe a generic noun or a specific individual. Participants then categorized creatures that had one of their properties transformed. We found that changes to teleological properties had the strongest impact on categorization judgments. In Experiment 3, we also found that teleological properties mattered more in an induction task. We suggest that teleological properties play this privileged role in categorization because they are treated as essential properties.

Keywords: teleology, essentialism, generics, transformations, induction.

Introduction

Looks aren't everything. This adage isn't merely an admonishment of vanity. It's also a valuable guide in categorization. Of course, appearances generally serve us well in organizing and cataloging the wide variety of things we find in the world. However, appearances change. Many objects lose and acquire new properties. Determining whether something persists across changes requires more than consulting its appearance (Rips et al., 2006). One important kind of property that we trace across changes to a thing is its essence: the property that an individual possesses that makes it a member of its kind (e.g., Gelman et al., 2003; Keil, 1989; Medin & Ortony, 1989).

If essential properties play a role in categorization, then this raises the question of how we come to learn about essences. Some psychologists maintain that generic statements, such as "bees are striped", serve as a vehicle for transmitting essentialist beliefs (e.g., Rhodes et al., 2012). Research has focused on what generics mean (see, e.g., Carlson, 1977; Leslie, 2007; Tessler & Goodman, 2019) and on the consequences that hearing generic statements have for how categories are represented (e.g., Brandone & Gelman, 2009; Gelman et al., 2010; Cimpian & Markman, 2009; Cimpian & Cadena, 2010; Rhodes et al., 2012,?; Gelman et al., 2010). Our focus is on whether the property type in a generic or specific statement matters for how things are categorized (see also Noyes & Keil, 2019, 2020).

Properties in generic statements

Gelman et al. (2010) used a variety of properties in generic statements (see also Rhodes et al., 2012). For instance, participants were told that "Zarpies have stripes on the bottom of their feet" and that "Zarpies hop over puddles". While having striped feet says something about Zarpies' biology or appearance, hopping over puddles says something about their behavior. Noyes & Keil (2019) investigate generics involving either biological (e.g., "Vawns can hold their breath really long"), social (e.g., "Vawns value nature"), or neutral (e.g., "Vawns can pick apples quickly") properties. They find that only generic statements with biological properties increase participants' acceptance of statements that they take to reflect essentialist beliefs. But sinceNoyes & Keil (2019) assigned the different properties to the three categories themselves, we don't know how participants viewed them.

What properties are privileged on the route from generics to categorization? One possibility is that any property which is viewed as bearing a principled connection to the category is privileged (Prasada & Dillingham, 2006). But this view, and related work, doesn't tell us which properties might be viewed as essential properties (Korman & Khemlani, 2020). For instance, according to this view, having four legs bears a principled connection to the dog category while being brown bears a mere statistical connection (Prasada & Dillingham, 2006). If a dog no longer had four legs, we might still think it is a dog. But the loss of essential properties should make us less inclined to think something retains category membership. Knowing whether a property bears a principled connection to a category doesn't, on its own, tell us whether that property is an essential property. We hypothesize that *teleological properties* – those that concern what something is for – carry more weight in categorization because they are viewed as essential properties. Why think this?

Children and adults accept teleological explanations for the existence of a broad range of living and non-living natural things (e.g., Bloom, 2007; Kelemen, 1999; Kelemen & Rosset, 2009; Kelemen et al., 2013; Lombrozo & Carey, 2006; Lombrozo et al., 2007; Lombrozo & Rehder, 2012; Foster-Hanson & Lombrozo, 2022). Teleological thinking also pervades people's judgments about existence and persistence:

Table 1: **Experiment 1**. Percentage of participants who judged that a given property belonged in that category (e.g., 86% of participants judged that "jump" is a behavior). The four properties in bold are the ones for which participants were the most sure that they belonged to the corresponding category. We use the short labels in square brackets in the figures below. *"Pollinate flowers" was a top-rated item but we used it only as an example in the introduction of the experiments.

Behavior	Biology	Purpose	Social
jump (86%)	**warm blooded** [blood] (84%)	pollinate flowers* (62%)	**share food with group members** [share] (82%)
swim (84%)	**pointy ears** [ears] (80%)	**enable decomposition** [decompose] (62%)	**cooperate with group members** [cooperate] (78%)
chew (76%)	**long legs** [legs] (80%)	**purify water** [purify] (58%)	**follow the dominant group member** [follow] (66%)
run (74%)	**hair** (80%)	**make honey** [honey] (54%)	**pair bond** [bond] (66%)
swallow (66%)	sharp teeth (76%)	**aerate soil** [aerate] (52%)	dance before mating (54%)
fly (60%)	tail (74%)	recycle nutrients in soil (48%)	nomadic (54%)
urinate (44%)	spots (72%)	enable nitrogen fixation (48%)	sing (30%)
salivate (44%)	small nostrils (72%)	catch and kill insects (46%)	mark territory (28%)
smell (36%)	large eyes (70%)	produce oxygen (40%)	store resources (16%)
digest slowly (18%)	claws (70%)	eat animal carcasses (22%)	build shelter (12%)

people think that a collection of parts forms a whole when those parts serve a collective purpose and that a thing persists through changes to its parts when it preserves its purpose (Rose, 2015, 2017, 2019, 2020; Rose et al., 2020). Perhaps teleology plays this central role in explanation and judgments of existence and persistence because we essentialize categories in terms of teleology. Indeed, recent research suggests that we do associate essences with a kind of telos.

Teleological essentialism

People say the purpose of a bee is to make honey and pollinate flowers, and that a spider is for making webs to catch insects (Rose & Nichols, 2019). If a bee is operated on to look like a spider, people categorize the creature as a bee when it still has the bee telos and as a spider when it has the spider telos. Judging that something persists across radical change when it preserves its purpose even extends to artifacts and non-living natural kinds (Rose & Nichols, 2020). Children also essentialize living kinds in terms of teleology (Rose et al., 2022). For both children and adults, changes in what a thing is for influence categorization judgments (though see e.g., Neufeld, 2021; Joo & Yousif, 2022). This work, however, hasn't focused on whether teleological properties might be privileged over other kinds of properties. And it hasn't investigated whether this might be so in the context of generic or specific statements.

While other work has focused on the role of functional features in categorization (e.g., Ahn, 1998; Lombrozo & Rehder, 2012), our focus here is on the role of generics and essentialism. Teleological essentialism has the potential to explain when and why functional features play a special role in categorization. Moreover, our methodological approach to the question of whether behavioral, biological, social, or teleological properties carry more weight in categorization is different. In prior work, the researchers assigned the properties to different categories themselves (e.g. Noyes & Keil, 2019). For example, "Vawns can hold their breath really long" was a biological property. We conducted a norming study first in which we ask participants to classify the different properties, which we then use in subsequent experiments. Additionally, while prior work has focused on induction tasks – categorizing a novel thing based on its properties – we look here also at transformation tasks – categorizing a thing that underwent some transformation – which are a stronger test of essentialist thinking.

Experiment 1: Property classification

All experiments, data, analyses, and links to pre-registrations are available here: `https://github.com/cicl-stanford/teleology_generics`. The goal of this pre-registered experiment was to identify properties that participants view as biological, behavioral, social, or teleological.[1]

Methods

Participants 50 participants who met our pre-registered inclusion criteria were recruited (*age*: M = 36, SD = 9; *gender*: 15 female, 33 male, 2 no response/other; *race*: 1 Asian, 1 Asian Indian, 2 Black, 1 Latino, 42 White, 4 no response/other *ethnicity*: 5 Hispanic, 43 Non-Hispanic, 2 no response/other) through Amazon Mechanical Turk. Participants received $1 as compensation.

Materials We created a list of 40 properties that we expected to be viewed as either biological, behavioral, social, or teleological (see Table 1 for full set of properties). We selected 10 properties for each category. The properties selected as candidates for biological, behavioral, and social properties are the same kinds of properties included in typical work on generics (see e.g., Rhodes et al., 2012; Gelman et al., 2010; Noyes & Keil, 2019). Teleological properties were drawn from or inspired by work in Rose & Nichols (2019) and Rose & Nichols (2020).

Procedure The experiment was programmed using PsiTurk (Gureckis et al., 2016). Participants were first told that they landed on a new island and have discovered 42 new things.

[1] Even though these categories aren't mutually exclusive in general (e.g. a property could be both social and behavioral), we treated them this way here because in later experiments, participants are asked to assign each of four properties to one unique category that fits best.

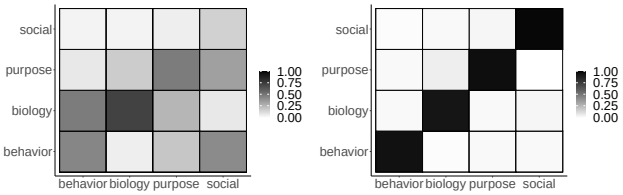

(a) Experiment 1 (b) Experiment 2

Figure 1: Proportion of expected property categorization (x-axis) versus actual property categorization (y-axis).

For each thing, they discover some feature of it and their task is to categorize that feature into one of four types: biology, behavior, purpose, or social. They were then given two examples involving familiar things and how their properties might be categorized. Participants then proceeded to the test trials. Here they were told that they came across a new thing on the island and decided to give it a name. Each name was a made-up name. Then they were told that after observing it, they noticed that it has some property where the property listed was a property from the list in Table 1. They were then asked which category the feature belonged to with "biology", "social", "behavior", or "purpose" to choose from.

Design Participants were given all 40 items from Table 1 and two additional items that served as attention checks. All items were presented in a random order. For each feature, participants could categorize them as belonging to biology, behavior, purpose, or social. The order in which the response options were listed was randomized between participants.

Results and discussion

Figure 1a shows expected versus actual property categorization. For the items expected to be categorized as biological, participants largely viewed them as biological. Similarly for purpose. Participants were unsure how to view the items in the social category and divided over whether putative behavioral properties were behavioral or biological.

Our main goal was to select properties to use in generic statements. Figure 2 shows the top four properties listed in each category which we selected for subsequent experiments.

Experiment 2: Categorization after transformation

How can we assess what role different property types play in the context of generics? A classic test of essentialist thinking involves judging that something persists across transformation (Keil, 1989; Rose & Nichols, 2019). Though transformations are never used in experiments aimed at documenting the role of generics in facilitating essentialism, given how central they are for assessing essentialist thinking, we make use of them in Experiment 2.

Having identified different property types, our question is: which properties carry more weight in categorization when predicated in a generic statement?

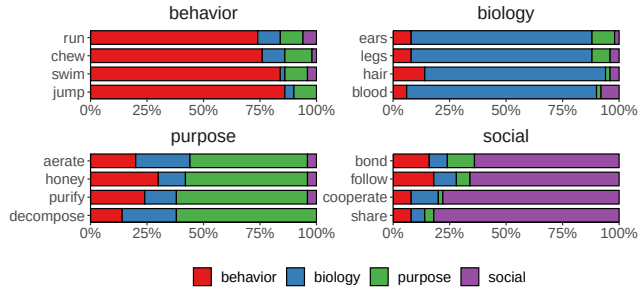

Figure 2: **Experiment 1**: Percentages of top four properties for each category.

Methods

Participants We recruited 100 participants who met our pre-registered inclusion criteria (*age*: M = 34, SD = 11; *gender*: 57 female, 38 male, 5 no response/other; *race*: 20 Asian, 13 Black, 62 White, 5 no response/other *ethnicity*: 10 Hispanic, 87 Non-Hispanic, 3 no response/other) through Prolific. Participants received compensation at a rate of $11 an hour.

Materials We created four creatures – Vulpans, Zerps, Jigs and Xans – based on the properties in Figure 2. Each had the same rank order property from each category: Vulpans had the top ranked property from each category, Zerps had the second ranked property, Jigs had the third, and Xans the fourth. Pictures were included for each creature (see Figure 3).

Procedure The experiment was programmed using jsPsych (de Leeuw, 2015). An example of the procedure for the generic condition is shown in Figure 3. Participants in this condition were given a bare plural generic (e.g., Vulpans) followed by four properties (see Figure 3a). They were also shown three exemplars of the creature below the generic statement. Participants categorized each property as being either behavioral, biological, social, or related to the thing's purpose. Each property needed to be placed in a unique category. Participants' selected categories in a table at the bottom of the page. This was repeated on a new page for a different creature. They then viewed a table summarizing how they categorized the properties of the two creatures.

On the test screen (see Figure 3b), participants were told that the first creature they encountered was exposed to toxic waste and underwent some changes. They were told that one of its properties changed and that it had the property of the same type from the second creature of the pair. For example, if one of the properties predicated to the first creature was "warm blooded" and a participant categorized this as biological, and if one of the properties predicated to the second creature was "has hair" and this was categorized as biological, then if the property changed by exposure to toxic waste was biological, participants would be told that the creature wasn't warm blooded like the first creature but instead had hair like

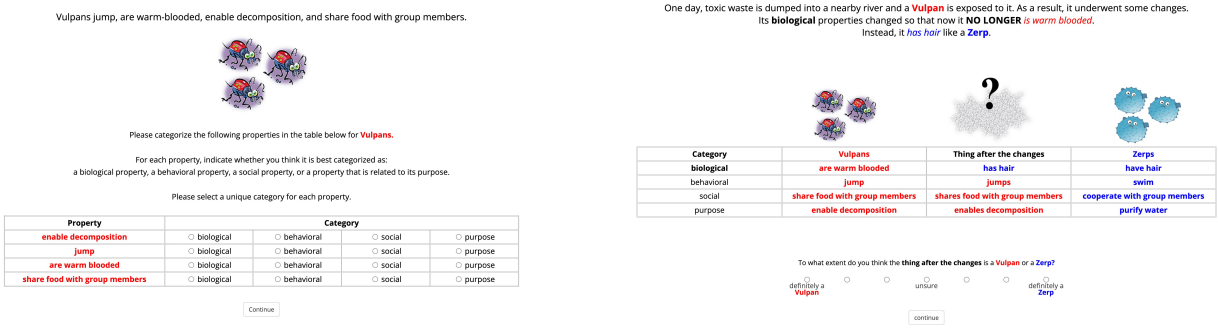

(a) Categorization task

(b) Judgment task

Figure 3: **Experiment 2**: Example trial from the generic condition. (a) *Categorization task*: Participants categorize the properties in the table. Each property can be included in only one category and each property must be assigned a unique category. This is then repeated for a second creature. (b) *Judgment task*: Participants judge to what extent the transformed creature is like one or the other category. In this example, the property that the participant categorized as being biological from the first creature is replaced with the property that the participant categorized as biological from the second creature.

the second creature. A table summarizing how each property was categorized for each creature was provided at the bottom of the screen and included a column indicating the properties of the creature after the changes. After learning about the change, participants then made ratings of category membership on a 7pt scale with 1 anchored with the name of the first creature and 7 anchored with the name of the second creature.

Design Participants were given all four creatures presented in pairs. The order of the creature pairs was randomized. Participants completed eight trials that involved a property from a creature being transformed into a property of the same type from the other creature in the pair. Property type transformation order was randomized within creature pairs. In addition to the property type transformation being manipulated within participants, we also manipulated, between participants, whether the properties predicated to the creature were in generic or specific form (we used a definite singular, e.g., "This Vulpan ...", and depicted a single individual below the statement).

Results

Figure 4 shows participants' category judgments based on the type of property that was transformed and whether a generic or specific statement was used. The main finding: teleological properties carry more weight in categorization (for each contrast, 95% of the posterior difference of the difference between purpose and other properties excluded 0). They do so whether they are predicated in a generic or specific statement (see Table 2). Indeed, whether properties are predicated in generic or specific form seems to make no difference, .06 $[-.14, .25]$.

There are two potential issues with our selection of properties. First, in Experiment 1, we gave people instructions on how to categorize properties and provided examples. That may have influenced them in unintended ways. But, as shown in Figure 1b, even without instruction – or indeed even with no description of the properties – people categorize them in very similar ways.

Second, some of the properties might be more diagnostic of category membership than others. For instance, "makes honey", probably only brings one thing to mind: a bee. But "has hair", probably brings many things to mind. Perhaps teleological properties are simply more diagnostic of category membership than other properties. Of course, if we essentialize in terms of teleology, then teleological properties should be more diagnostic of category membership. If essences do anything, they should be diagnostic of category membership. That said, perhaps our teleological properties are diagnostic in a way that unfairly advantages them.

We probed large language models to see what probabilities they assigned to completions. The basic idea is that if a language model is relatively certain how a sentence with a given property continues, then this property is diagnostic of

Table 2: **Experiment 2**: Posterior distributions of the difference between purpose and other properties for both the generic and specific condition. *Note*: CI = credible interval.

contrast	median	lower 95% CI	upper 95% CI
generic condition			
purpose - behavioral	0.54	0.22	0.90
purpose - biological	0.40	0.07	0.73
purpose - social	0.92	0.57	1.24
specific condition			
purpose - behavioral	0.93	0.57	1.25
purpose - biological	0.63	0.32	0.98
purpose - social	0.87	0.54	1.20

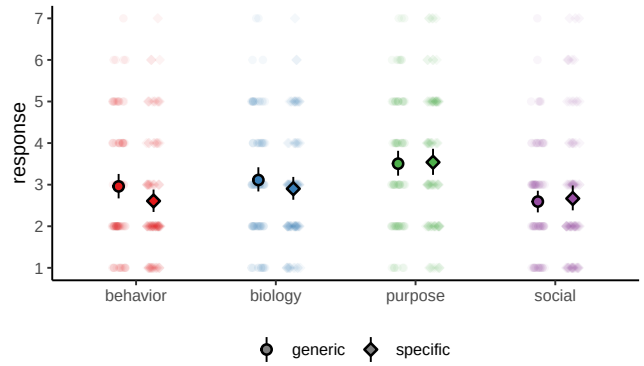

Figure 4: **Experiment 2**: Category judgments based on what property was transformed in the generic (circle) and specific (diamond) conditions. In all figures, higher ratings indicate that the original thing has changed categories after the property change. Large shapes are means with 95% bootstrapped confidence intervals. Small points are individual responses.

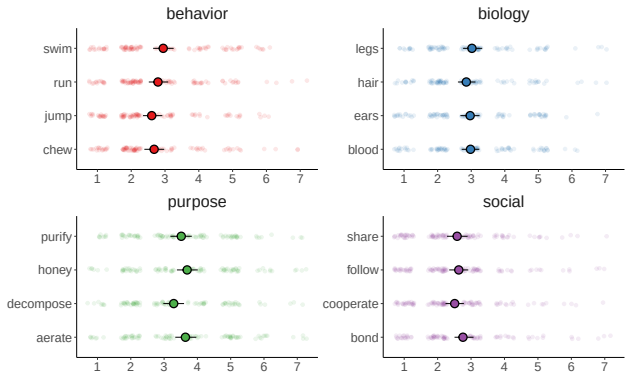

Figure 6: **Experiment 2**: Categorization ratings for each property across all property types. Higher ratings indicate that the original thing has changed categories after the property change. Large points are means with 95% bootstrapped confidence intervals. Small points are individual responses

its kind. We probed BERT-base, BERT-large and RoBERTa-large. Each model was given the following prompt: "An animal that [property] is a ..." [property] was replaced with each property used in Experiment 2 (see Figure 2). For each property, we extracted the top five completions and their associated probabilities. Figure 5 shows the results.

No property type was more diagnostic than the others. But there are some outliers. In particular, one teleological property stands out: "making honey". This seems to point toward bees. There was also a prominent behavioral property: "swimming". In fact, "swimming" was roughly as diagnostic for fish, as "making honey" was for bees. If mere diagnosticity mattered for categorization, one would expect that "swimming" and "making honey" have a greater influence on categorization than other behavioral or teleological properties. But as can be seen in Figure 6, this wasn't the case. "Making honey" didn't have a bigger effect than other teleological properties, even though those other properties were less diagnostic. Similarly, even though "swimming" was more diagnostic than other properties, it didn't appear to be carrying more weight. Indeed, across all property types, there wasn't a single property that affected categorization judgments more strongly, which is what one would have expected if mere diagnosticity was driving categorization judgments.

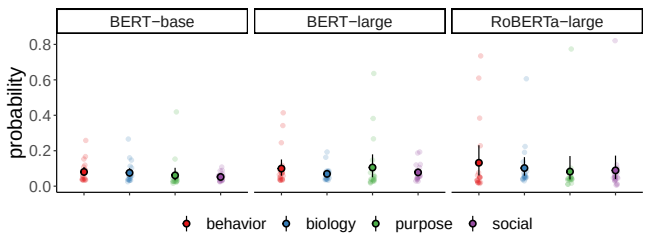

Figure 5: **Experiment 2**: Probabilities of completion assigned by language models separated by property type.

Discussion

Teleological properties carry more weight when a creature undergoes a transformation, regardless of being predicated in a generic or specific statement. And this isn't simply because they are more diagnostic of category membership (see Figure 5 and Figure 6). Categorization judgments were generally below the midpoint. This isn't surprising since only one feature out of four was changed in the transformation. The important point, though, is that changing teleological properties had a stronger impact on categorization judgments than changing any of the other properties.

Another important consequence of exposure to generics – and the corresponding essentialist tendencies they are supposed to induce – is that they facilitate generalization. In our final experiment, we ask whether teleological properties carry more weight in a categorization task that focuses on induction.

Experiment 3: Categorization in induction

Here we ask whether teleological properties carry more weight in categorization in the context of an induction task.

Methods

Participants We recruited 100 participants who met our pre-registered inclusion criteria (*age*: M = 34, SD = 14; *gender*: 48 female, 48 male, 4 no response/other; *race*: 14 Asian, 6 Black, 76 White, 4 no response/other *ethnicity*: 14 Hispanic, 85 Non-Hispanic, 1 no response/other) through Prolific. Participants received compensation at a rate of $12 an hour.

Materials & Design The materials and design were the same as in Experiment 2.

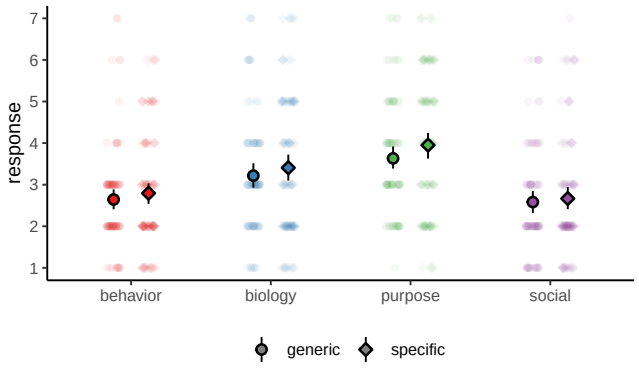

Figure 7: **Experiment 3**: Effect of property type on categorization ratings in the generic (circle) and specific (diamond) conditions.

Procedure The experiment was programmed using jsPsych (de Leeuw, 2015) and the procedure was the same as in Experiment 2. The only difference was that on the test page, participants didn't read about a creature getting exposed to toxic waste and undergoing transformation. Instead, they were told that one day, they came across a creature that had all of the properties of the first creature in the pair except for one. That property was one from the second creature in the pair. Participants were then asked whether this new creature belonged to the first or second creature's category and made ratings on the same 7pt scale.

Results

Figure 7 shows participants' categorization judgments. Teleological properties continue to carry more weight when categorizing in an induction task. They do so whether they are predicated in a generic or specific statement. As with transformations, whether properties are predicated in generic or specific form seems to make no difference $-.09$ $[-.29, .10]$.

General Discussion

Generic statements such as "Bees are striped" have been argued to be a central vehicle by which essentialist beliefs about categories are transmitted (see e.g., Rhodes et al., 2012; Gelman et al., 2010). But what properties take front seat on the ride from generics to categorization? We hypothesized that teleological properties – properties that capture what a thing's purpose is – may play this privileged role. To test this, we asked people in Experiment 1 to categorize a range of properties that are used in typical work on the role of generics in categorization. We then took the top four properties from each of our categories and developed experimental materials. In Experiment 2, we tested which property type carried more weight in categorization. Transformation cases, we suggested, provide a central test of essentialist thinking. We thus used a transformation task and found that teleological properties carried more weight in categorization. It made no difference whether those properties were predicated to a generic noun or a specific individual. In Experiment 3, we found that teleological properties also carried more weight in an induction task where participants categorized a novel creature.

When drawing on the kinds of properties typically used in research on the effects of generics in essentializing categories (e.g., Rhodes et al., 2012; Gelman et al., 2010; Noyes & Keil, 2019), we find that teleological properties carry more weight. This isn't merely because teleological properties are more diagnostic of category membership (see Figure 5 and Figure 6).

The fact that we found that whether properties were predicated to a generic noun or specific individual made no difference suggests that what matters in essentializing is the properties predicated and not whether they are predicated to a generic noun. Assuming that transformation tasks provide good evidence of essentialist thinking, and given that generics make no difference in these tasks, it may be that whatever role generics play in categorization, it isn't one of facilitating essentialism. This isn't to deny that generics play a role in categorization. But not all categorization involves essentializing. And when it comes to categorizing based on essences, generics may be less important. Instead, what matters for whether we essentialize is whether the property is one that we are inclined to essentialize. As our findings suggest, teleological properties seem to play this role.

Our findings contrast with the view that generics facilitate essentialism (see e.g., Rhodes et al., 2012; Gelman et al., 2010). And they contrast with the most prominent view of psychological essentialism – which has dominated the literature for the last 30 years – biological essentialism (e.g., Gelman & Wellman, 1991; Hirschfeld et al., 1999; Gelman et al., 2003; Keil, 1989). Our findings build on a different proposal concerning essences, that of teleological essentialism (e.g., Rose & Nichols, 2019). They also cohere with work on teleological explanation (e.g., Kelemen, 1999; Lombrozo & Carey, 2006; Foster-Hanson & Lombrozo, 2022) and work showing that teleology plays a central role in judgments of composition, persistence (e.g., Rose et al., 2020), causation (Rose, 2017) and classic tests of essentialists thinking (e.g., Rose & Nichols, 2019).

Some recent work challenges teleological essentialism by suggesting that changes in a thing's telos lead people to infer that a thing's insides have changed. And it's insides – not teleology – that matter in essentialist categorization (e.g., Neufeld, 2021; Joo & Yousif, 2022). It isn't clear how that proposal could explain the current set of findings which indicate that teleological and not, for example, biological or behavioral properties play a greater role in categorization. But even if changes in teleology lead to inferences that insides have changed, that is entirely compatible – and indeed expected – on a teleological view of essentialism since insides are relevant to the realization of functions. The important point for present purposes, however, is that teleological properties carry more weight in categorization, whether predicated to generic nouns or specific individuals. Teleological properties might be viewed as essential properties.

Acknowledgments

We thank that Causality in Cognition Lab (CiCL), the Markman Lab and FriSem participants for valuable feedback on the project. TG was supported by a grant from the Stanford Institute for Human-Centered Artificial Intelligence (HAI). DR was supported by a Stanford Interdisciplinary Graduate Fellowship.

References

Ahn, W.-k. (1998). Why are different features central for natural kinds and artifacts?: The role of causal status in determining feature centrality. *Cognition*, *69*(2), 135–178.

Bloom, P. (2007). Religion is natural. *Developmental science*, *10*(1), 147–151.

Brandone, A. C., & Gelman, S. A. (2009). Differences in preschoolers' and adults' use of generics about novel animals and artifacts: A window onto a conceptual divide. *Cognition*, *110*(1), 1–22.

Carlson, G. N. (1977). *Reference to kinds in english: A dissertation* (Unpublished doctoral dissertation). Graduate Linguistic Student Association, Department of Linguistics.

Cimpian, A., & Cadena, C. (2010). Why are dunkels sticky? preschoolers infer functionality and intentional creation for artifact properties learned from generic language. *Cognition*, *117*(1), 62–68.

Cimpian, A., & Markman, E. M. (2009). Information learned from generic language becomes central to children's biological concepts: Evidence from their open-ended explanations. *Cognition*, *113*(1), 14–25.

de Leeuw, J. R. (2015). jspsych: A javascript library for creating behavioral experiments in a web browser. *Behavior Research Methods*, *47*(1), 1–12.

Foster-Hanson, E., & Lombrozo, T. (2022). What are men and mothers for? the causes and consequences of functional reasoning about social categories. In *Proceedings of the annual meeting of the cognitive science society* (Vol. 44).

Gelman, S. A., et al. (2003). *The essential child: Origins of essentialism in everyday thought*. Oxford University Press, USA.

Gelman, S. A., Ware, E. A., & Kleinberg, F. (2010). Effects of generic language on category content and structure. *Cognitive psychology*, *61*(3), 273–301.

Gelman, S. A., & Wellman, H. M. (1991). Insides and essences: Early understandings of the non-obvious. *Cognition*, *38*(3), 213–244.

Gureckis, T. M., Martin, J., McDonnell, J., Rich, A. S., Markant, D., Coenen, A., ... Chan, P. (2016). psiturk: An open-source framework for conducting replicable behavioral experiments online. *Behavior research methods*, *48*(3), 829–842.

Hirschfeld, L., Gelman, S., Atran, S., & Douglas, M. (1999). How biological is essentialism. *Folkbiology*, 403–446.

Joo, S., & Yousif, S. R. (2022). Are we teleologically essentialist? *Cognitive Science*, *46*(11), e13202.

Keil, F. C. (1989). *Concepts, kinds, and cognitive development*. MIT Press.

Kelemen, D. (1999). Why are rocks pointy? children's preference for teleological explanations of th natural world. *Developmental psychology*, *35*(6), 1440–1459.

Kelemen, D., & Rosset, E. (2009). The human function compunction: Teleological explanation in adults. *Cognition*, *111*(1), 138–143.

Kelemen, D., Rottman, J., & Seston, R. (2013). Professional physical scientists display tenacious teleological tendencies: Purpose-based reasoning as a cognitive default. *Journal of experimental psychology: General*, *142*(4), 1074.

Korman, J., & Khemlani, S. (2020). Teleological generics. *Cognition*, *200*, 104157.

Leslie, S.-J. (2007). Generics and the structure of the mind. *Philosophical perspectives*, *21*, 375–403.

Lombrozo, T., & Carey, S. (2006). Functional explanation and the function of explanation. *Cognition*, *99*(2), 167–204.

Lombrozo, T., Kelemen, D., & Zaitchik, D. (2007). Inferring design: Evidence of a preference for teleological explanations in patients with alzheimer's disease. *Psychological Science*, *18*(11), 999–1006.

Lombrozo, T., & Rehder, B. (2012). Functions in biological kind classification. *Cognitive psychology*, *65*(4), 457–485.

Medin, D., & Ortony, A. (1989). Psychological essentialism. *Similarity and analogical reasoning*, 179–195.

Neufeld, E. (2021). Against teleological essentialism. *Cognitive Science*, *45*(4), e12961.

Noyes, A., & Keil, F. C. (2019). Generics designate kinds but not always essences. *Proceedings of the National Academy of Sciences*, *116*(41), 20354–20359.

Noyes, A., & Keil, F. C. (2020). There is no privileged link between kinds and essences early in development. *Proceedings of the National Academy of Sciences*, *117*(20), 10633–10635.

Prasada, S., & Dillingham, E. M. (2006). Principled and statistical connections in common sense conception. *Cognition*, *99*(1), 73–112.

Rhodes, M., Leslie, S.-J., & Tworek, C. M. (2012). Cultural transmission of social essentialism. *Proceedings of the National Academy of Sciences*, *109*(34), 13526–13531.

Rips, L. J., Blok, S., & Newman, G. (2006). Tracing the identity of objects. *Psychological Review*, *113*(1), 1.

Rose, D. (2015). Persistence through function preservation. *Synthese*, *192*(1), 97–146.

Rose, D. (2017). Folk intuitions of actual causation: a two-pronged debunking explanation. *Philosophical Studies*, *174*, 1323–1361.

Rose, D. (2019). 14 cognitive science for the revisionary metaphysician. *Metaphysics and Cognitive Science*.

Rose, D. (2020). Mentalizing objects. *Oxford Studies in Experimental Philosophy*.

Rose, D., Jamarillo, S., Nichols, S., & Horne, Z. (2022). Teleological essentialism across development. *Proceedings of the Annual Meeting of the Cognitive Science Society*(44).

Rose, D., & Nichols, S. (2019). Teleological essentialism. *Cognitive science*, *43*(4), e12725.

Rose, D., & Nichols, S. (2020). Teleological essentialism: Generalized. *Cognitive science*, *44*(3), e12818.

Rose, D., Schaffer, J., & Tobia, K. (2020). Folk teleology drives persistence judgments. *Synthese*, *197*(12), 5491–5509.

Tessler, M. H., & Goodman, N. D. (2019). The language of generalization. *Psychological Review*, *126*(3), 395–436.