

NORTHWESTERN UNIVERSITY

Comparison and Alignment in Categorization

A DISSERTATION

SUBMITTED TO THE GRADUATE SCHOOL IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

for the degree

DOCTOR OF PHILOSOPHY

Field of Psychology

By

Francisco Maravilla

EVANSTON, ILLINOIS

September 2020

Abstract

Many theories of categorization have included an intuitive role for our ability to detect and judge similarity. Yet, this important role of similarity processing has been disputed. This research adopts a model of similarity processing through structure mapping (Gentner, 1983) to explore its role in similarity processing and categorization. *Relational categories*, organized around shared structure rather than overlapping surface features, provide an ideal arena in which to test this claim. If structural alignment is central to both similarity and categorization for these categories, then factors known to affect similarity judgments should affect categorization judgments in the same way. The three experiments reported here explore three manipulations of alignability: promoting alignment between possible category members, promoting alignment between same-category members, or disrupting alignment of a category member and possible category members. The evidence is ultimately inconclusive. While manipulating alignment affects both similarity and categorization, it does not do so consistently, nor is it clear that this is the only explanation for the observed differences. Implications and future directions are discussed.

Acknowledgments

This work would not be possible without the continual support of my adviser, Dr. Dedre Gentner, whose passion and drive for cognitive science never ceases to inspire me to be a better scientist. This work would also not be possible without the invaluable commentary and advice of my committee members, Dr. Sid Horton and Dr. Lance Rips, whose scholarly expertise and continual patience know no bounds. I also want to recognize the contribution of Dr. Nina Simms, whose willingness and eagerness to mull over experimental design and theory greatly improved the quality of my work.

Dedication

This thesis is dedicated, as it always would be, to my parents.

Contents

| | |
|--|----|
| List of Figures and Graphs..... | 7 |
| Introduction | 8 |
| Is Similarity the Basis for Categorization? | 9 |
| Arguments for Similarity in Categorization | 10 |
| Arguments against Similarity in Categorization..... | 11 |
| Similarity is Too Flexible..... | 11 |
| Similarity is Not Sufficient..... | 15 |
| Models of Similarity: Comparison through Structure Mapping | 21 |
| Structure Mapping Theory..... | 22 |
| Structural Alignment..... | 23 |
| Structural Alignment and Categorization | 28 |
| Two Roles for Structural Alignment in Categorization | 29 |
| Forming Categories through Alignment | 30 |
| Co-categorizing Representations through Alignment | 34 |
| Structural Alignment and Relational Categories | 36 |
| What Are Relational Categories?..... | 37 |
| Alignment and Relational Category Learning | 38 |
| Summary | 40 |

| | |
|------------------------------------|-----|
| Overview of Experimental Work..... | 41 |
| Analysis | 45 |
| Experiment 1 | 46 |
| Method | 47 |
| Results..... | 50 |
| Discussion..... | 52 |
| Experiment 2 | 55 |
| Method | 56 |
| Results..... | 59 |
| Discussion..... | 65 |
| Experiment 3..... | 67 |
| Method | 69 |
| Results..... | 71 |
| Discussion..... | 72 |
| General Discussion..... | 74 |
| Conclusion..... | 77 |
| Appendix A..... | 91 |
| Appendix B..... | 102 |

List of Figures and Graphs

| | |
|--|----|
| Figure 1: Sample stimuli from Medin, Goldstone & Gentner, 1991..... | 26 |
| Figure 2: Sample stimuli from Sagi, Gentner, and Lovett, 2012..... | 27 |
| Figure 3: Two passages exemplifying the same logical fallacy: Appeal to Authority..... | 42 |
| Figure 4: Experiment 1 results..... | 50 |
| Figure 7: Experiment 2 test phase results..... | 60 |
| Figure 8: Experiment 2 seen vs. unseen results..... | 61 |
| Figure 9: Example responses to Experiment 2 pre-task..... | 63 |
| Figure 10: Experiment 3 results..... | 70 |

Introduction

Categorization – the means by which we organize and label the things in our world into concepts – is one of the fundamental processes through which humans make sense of the world. Many explanations of this cognitive faculty have made a bold but intuitive claim: categorization is directly tied to our ability to detect and judge similarity. However, evidence suggests that this strong claim cannot be accepted as is; cases exist where similarity and categorization lead to different conclusions. For instance, Rips (1989) showed that people may sometimes incorporate additional content like causal structure into category judgments.

The research presented here aims to preserve the intuitive appeal of the similarity approach while clarifying and delineating its role in categorization. An important component of this clarification is considering the distinction between representational content and cognitive process. Instead of the strong argument that similarity directly determines categorization, I argue that the cognitive products of similarity and category judgment are related in their underlying process. Specifically, that comparison through *structural alignment* – a key component process of Structure-Mapping Theory (Falkenhainer et al., 1989, 1989; Gentner, 1983, 2003, 2010; Gentner & Markman, 1997a) – is involved both in similarity judgments and categorization. To the extent that structural alignment is involved in categorization, manipulations that influence alignability should influence not only similarity judgments but also categorization decisions.

The reported experiments focus specifically on *relational categories*, whose members share common relational structure without necessarily sharing common features. These

categories pose a special difficulty for models of similarity based on shared features. It is argued that in the case of relational categories, alignment is involved a) when a novel category is formed from exemplars and b) when deciding whether a potential new member belongs to a category. There are two main hypotheses:

1. Both similarity *and* categorization should be affected by alignability: making structural alignment easier will increase relational responding for both similarity judgments and co-categorization tasks, while making alignment more difficult will decrease both.
2. When surface and relational similarity compete, people will show higher same-category responding on categorization tasks than on similarity tasks.

These hypotheses are explored in three experiments, each with specific hypotheses related to their design, but guided by the preceding framework.

Before discussing the current research, however, it is important to properly frame the argument by reviewing the argument over similarity in categorization, the structure-mapping process, and how it may be applied to categorization.

Is Similarity the Basis for Categorization?

This section briefly describes arguments for and against similarity as a basis for categorical knowledge representation. It broadly reviews approaches to categorization based on similarity, and two main criticisms of these approaches: that similarity is too flexible to ground

categorization, and that factors other than similarity bear more importance on category judgments.

Arguments for Similarity in Categorization

Similarity has played a strong role in theories of categorization for decades of cognitive research (see Murphy, 2002 for a comprehensive review). The influence of this approach can be seen even in theories not explicitly built around similarity. For instance, the classical view can be construed as stating that all members of a concept must be alike in sharing a set of summary features that are both necessary and sufficient to define the class (E. E. Smith & Medin, 1981). More explicitly, prototype models (Posner & Keele, 1968; Rosch, 1973, 1975; Rosch & Mervis, 1975) account for *family resemblance* (Wittgenstein, 1953) by organizing categories around similar summary representations. Finally, according to *exemplar theories*, an object is part of category A and not category B if it is more similar to individual items that have been stored as instances of category A than it is to those that have been stored as instances of category B (Brooks, 1978; Medin & Schaffer, 1978).

These summaries gloss over important differences between these theories; however, all three share the fundamental assumption that categorizing objects in the world involves some similarity between those objects and stored category representations. This central role of similarity has not, however, been universally accepted. Some critiques argue that similarity is either a vacuous construct, too unconstrained to play a central role in the actual categorization process, while others argue that shared similarity cannot be the necessary nor sufficient basis for categorization.

Arguments against Similarity in Categorization

There have been several arguments proposed against the role of similarity in categorization (see Goldstone, 1994 for a review). The arguments described in this section take two different approaches. The first set takes the strong position that similarity is entirely vacuous and thus should be abandoned as a basis for categorization while. The second set argues the position that similarity, at least as commonly described, is simply unnecessary to the categorization process. Reviewing these objections provides a context for the hypotheses of the current experimental work. This also highlights the point that many previous attempts to account for similarity in categorization focus on the outcome of similarity processing and not the process itself.

Similarity is Too Flexible

One long-standing argument, made famously by Goodman, maintains that similarity is too flexible and unconstrained to guide categorization (1972). This argument holds that any two objects could be similar in a near infinite number of ways, while categories are much more fixed¹. For instance, consider the similarities between a violin and a frog. They both make distinctive noises, they both have relatively symmetrical bodies, interact with humans, are found only on one planet in this solar system (that we are aware of), do not emit light, do not fly, are lighter than a locomotive engine, etc. Of course, the defender of similarity might

¹ Exactly how fixed is a matter of contention; however, this is an implicit premise of this argument.

respond that yes, while there is a near infinite set of *logically possible* properties that any two objects could share, we are interested in *psychological reality*; the process by which humans determine similarity constrains the set of psychological properties. Goodman, however, would deem this response inadequate.

Instead he adopts a strong version of the skeptical argument, claiming that "...since every two things have some property in common, this will make similarity a universal and hence useless relation. That a given two things are similar will hardly be notable news if there are no two things that are not similar," (1972, p. 443). He further claims that saying two things are similar is meaningless unless one can specify *with respect to what property* they are similar. Returning to the frog and violin example, Goodman would argue that the sentence "A frog is similar to a violin," would be rendered meaningless unless someone specifically added "they are both marvelous instruments of sound." This of course poses a problem for similarity. If describing two things as similar ultimately amounts to indicating the source of the similarity, then similarity becomes a cipher: meaningless and empty without a specific property that provides it meaning. In Goodman's view, all the work is done by property attribution and not similarity – to meaningfully categorize something one must indicate that it is a member of category A because it is most similar to A items *with respect to producing sound* – to stick with the previous example. In this case, similarity contributes nothing; we can simply say this object is a member of category A *because it produces sound*. Thus, as Goodman (1972) puts it, "I suspect the best we can do is to say that all [things] that are *a*'s must be alike in being *a*'s. That has the solid ring of assured truth but is hardly electrifying...The words 'alike in being' add nothing; similarity becomes entirely superfluous," (p. 439).

Indeed, there seems to be some empirical evidence to support Goodman's claim, largely arising from featural models of similarity (Barsalou, 1982, 1983; Goldstone et al., 1991; Medin & Shoben, 1988; Roth & Shoben, 1983). One hallmark of similarity judgments is that they can be dramatically influenced by context. In one example, people judged *racoons* and *snakes* as fairly dissimilar outside of any additional context; however, when the word *pets* was placed directly above them, participants rated the pair as more similar (Barsalou, 1982). In related work, Roth & Shoben (1983) gave participants a broad category exemplified in two different sentences. For instance, for the category *beverages* participants would see one of the following sentences: "During the midmorning break, two secretaries gossiped as they drank the *beverage*," or "Before starting his day, the truck driver had the *beverage* and a donut at the truck stop." After reading the first sentence, participants rated tea as more similar to coffee than milk; however, after reading the second sentence, participants rated tea as *less similar* to coffee than milk. Medin & Shoben (1988) extended these results from object concepts to adjective/noun combinations. They gave participants triads of adjective/noun combinations and asked them to select which two were the most and least similar. For example, participants might see a three-way comparison like the following: *grey hair-white hair*, *white hair-black hair*, and *grey hair-black-hair*. Participants tended to select *grey hair* as more similar to *white hair* than *black hair*; however, when the noun was changed to *cloud* (despite keeping the same adjectives), the pattern reversed. These results could be taken as support for Goodman's claim that similarity is meaningless without being specified with respect to a specific property. Context specifies the relevant properties, reducing similarity to a property highlighted within the context. If this kind of similarity determined category membership, any two things that

were similar in a given context would necessarily be categorized together, something that does not happen.

In addition to being variable by context, similarity also seems to be variable across development. Several studies indicate that early in learning children's similarity judgments are more holistic and perceptually-based than they are with greater expertise (Gentner, 1988; Gentner & Rattermann, 1991; Keil & Batterman, 1984; Kemler-Nelson, 1989; L. B. Smith, 1989). When comparing things, adults (and domain experts in general) tend to focus on conceptual commonalities, whereas children (and novices) tend to respond based on overall similarity or surface similarity. Linda Smith (1989) offered a model of perceptual classification in which children transition from holistic similarity judgments to judgments that selectively weight relevant dimensions. Another trend in the development of similarity is the well-documented relational shift from attention to object similarity to a focus on similarity based on relational commonalities (Gentner, 1988; Gentner & Rattermann, 1991; Gentner & Toupin, 1986; Kotovsky & Gentner, 1996). For instance, Gentner (1988) found that when given the comparison "a cloud is like a sponge," a typical 5-year-old might respond that "They are both round and fluffy," while an adult (or a 9-year-old) would note the relational commonality: "They both hold and release water." Gentner and colleagues proposed that this shift stems largely from increases in children's conceptual knowledge (Gentner & Rattermann, 1991). There may also be contributions from developmental increase in executive function and working memory (Richland et al., 2004; Thibaut et al., 2010; see Simms et al., 2018 for a review). These findings indicate that what counts for subjective similarity shifts over the course of development, though this is far less surprising and damaging to the similarity account than other objections.

Taken together, the results described in this section make clear that subjective similarity can vary based on many factors, including the current context, how those objects interact, the salience of specific properties they share, as well as one's prior knowledge of the domains. Further, the factors that influence perceived similarity change over the course of learning and development. Category representations, however, do not show the same variability across contexts and features. Unfortunately, this is not the only criticism of an approach to categorization based on featural similarity.

Similarity is Not Sufficient

In addition to the argument that similarity is too vacuous or over-flexible to determine, categorization, many have argued that categorization cannot be accounted for by similarity alone. For example, consider a classic thought experiment proposed by Murphy and Medin (1985). They describe a scenario where one witnesses a person jumping into a swimming pool with their clothes on. How one interprets this event will depend on contextual knowledge and causal theories about the behavior much more than it will on similarity between the event and others like it. For instance, classifying the person as *intoxicated* will be more likely if there is also music and general revelry. At the same time, "jumps into a pool fully clothed", they argue, is likely not part of the standard concept for intoxicated people. Murphy and Medin use this example to illustrate that it is not the similarity between the instance and the category concept (or its exemplars) that determines classification, but rather whether the example fits one's *theory* of drunken behavior. As Rips put it, "categorization and category learning are special cases of inference to the best explanation" (1989, p. 53).

Much research has utilized a tension between observable appearances and theory-based knowledge to demonstrate the divergence between theory and similarity (Carey, 1985; Gelman & Markman, 1986; Keil, 1989, 1995; Rips, 1989). In one experiment, children were told scientists had discovered a new animal that looks exactly like a raccoon, but actually contains the same internal organs of a skunk, has skunk parents, and skunk children (Keil, 1989). When asked to classify the animal, younger children were likely to use surface similarities, classifying the animal as a raccoon; older children were more likely to classify the animal as a skunk instead. This suggests a developmental trend toward theory-based categorization over categorization based on perceptual similarity. This is congruent with previously discussed evidence of a relational shift over the course of learning and development (Gentner & Medina, 1998).

Rips (1989) conducted a study analogous to Keil's task, in which adult participants were told stories about transformations of natural and artifact kinds. He designed clever instances where surface features contrasted with underlying explanatory causes. For instance, participants were introduced to a novel animal called a *Sorp*. Sorps appear more or less like birds, consuming seeds and berries, having two wings, two legs, covered in bluish-grey feathers and building nests made of twigs and other fibrous plant material in the high branches of trees. In the *accident* condition, a particular sorp had made its nest near a toxic waste disposal site, and before long began to change its appearance, sprouting a new pair of wings made of a transparent membrane, developed an iridescent outer carapace, and two additional pairs of legs with adhesive pads, such that the sorp was able to hold onto and climb smooth surfaces. The sorp abandoned its nest and began subsisting on the nectar of flowers. However, as in

Keil's (1989) scenarios, the mutated sorp mated with a normal sorp and produced normal sorp offspring. In the *essence* condition, by contrast, participants were introduced to another novel animal, the *Doon*. Doons are very much like sorps (i.e. like birds) and undergo the same transformation as in the accident condition; however, the transformation was described as a natural metamorphosis that happens to *all* doons. Furthermore, doons give birth to normal doons, that follow the same life cycle and begin their lives looking like birds. Participants in each condition were asked to provide three ratings of the animals described: whether the animal was *more likely to be* a bird or an insect, whether it was *more typical of* a bird or an insect, and whether it was *more similar to* a bird or an insect. Participants' ratings in both conditions showed a divergence between their categorization and similarity ratings. In the accident condition, participants tended to categorize the afflicted sorp as more likely to be a bird but rated it as more similar and typical of insects. The opposite trend was observed in the essence condition; the doon was categorized as an insect, but more similar and typical of birds. These findings were taken as support that knowledge about explanatory causes seems to influence categorization while surface features influence similarity ratings.

This evidence is bolstered by similar examples found in nature. For instance, whales share many surface similarities with fish, yet are classified as mammals. However, additional research has taken a more skeptical position towards these results. (Ahn et al., 2000; Goldstone, 1994; Hampton, 1998, 2001, 2007; Hampton et al., 2007; Kalish, 1995). One possible explanation is that different methods of assessment could influence the weight given to surface-level attributes vs. deeper information. Goldstone (1994) argues that similarity and typicality could be assessed along a continuum of questions that range from perceptually- to

conceptually-grounded². As Goldstone describes it, (referring also to Keil's raccoon-skunk studies),

Rips and Keil essentially use variations of Question 2 ['Which species is this animal more similar to?']. However, there is no *a priori* reason to think that Question 2 reveals "true similarity." The fact that Question 2 contains the word 'similar' does not guarantee that it provides evidence about what psychologists refer to as "similarity." In Rips' and Keil's experiments, there may very well be a strong task demand to interpret 'similar' as 'visually similar,' but it is doubtful that subjects in their everyday life, only adopt a similarity measure tapped by Question 2 ['Which species is this animal more similar to?'] (Goldstone, 1994, p. 143).

Goldstone reports unpublished work in an informal setting where students made similarity and category judgments in a metamorphosis task like Rips'; however, they were asked to consider similarity in a slightly different manner (e.g., "Which species is this animal more like, taking into consideration *all* of the information that you have available?"). Using this method, the two judgments converged, contrary to the results reported by either Rips or Keil.

Kalish (1995) also found dissociations between typicality and categorization, but found that the dissociation was attenuated when participants were prompted to adopt an explicitly biological perspective (e.g., "Biologically speaking, how typical is ...?"). Relatedly, Ahn and

² Consider the difference between "Which of these species does this animal look more like?" vs. "Which species is this animal more likely to belong to?" We will return to this distinction when describing proposed work.

colleagues report that people weight causes more than effects when making both similarity and category judgments (Ahn et al., 2000), although this weighting was more pronounced for category judgments than for similarity judgments.

Hampton and colleagues (2007) used Rips' original metamorphosis problems to investigate the possibility that people may use surface features to infer deeper causal structure. For example, a camel's hump allows it to last without water for a long time; a bird's low bone density allows it to fly. This could explain the dissociation between similarity, typicality and categorization. Each judgment emphasizes different conceptual components. Typicality judgments place a relatively heavy weight on appearance over deeper properties, category membership judgments involve causal or structural elements that may cause other properties, while similarity judgments could rely either on surface-level properties or on the causal systems of which they are a part. Hampton and colleagues (2007) set out first to replicate the original dissociation between similarity, typicality and categorization. As in the original experiment, the creatures underwent some sort of transformation that in some cases was a normal part of maturation, and in others was the result of a mutation. However, the scenarios were also modified slightly to decrease ambiguity and increase experimental control: the paragraphs were shortened, the connection between the features and animal categories was made more explicit, and the difference between each condition limited to one causal clause (e.g. "One day, [as a result of toxic contamination of its environment/as a result of natural developmental processes], the animal began to change"). Even with these changes, they found a similar pattern of dissociation: in the mutation condition, the creatures were judged to be more typical of their new category after their transformation but were rated as more likely to still be

members of their original species, while in the maturation condition, both category judgments and typicality ratings were closer to the other species after the transformation.

On the face of it, this evidence reinforces a dissociation between categorization and similarity (Rips, 1989). However, Hampton and colleagues (2007) also explored whether the within-subjects design of the original experiment may have led to a contrast effects, or to demand characteristics that led participants to respond differently to the two measures. As Goldstone (1994) argued, participants may have felt that responding in the same or a similar way for every item would seem uncooperative or unusual for an experiment. In a follow-up between-subjects design, the dissociation disappeared, with most participants responding both that post-mutation animals had switched categories and that they were now typical of the new category.

These findings have important implications for categorization models. First, they suggest that classification based on purely theory-based information may not be as common as earlier experimental evidence (Murphy & Medin, 1985; Rips, 1989; Rips & Collins, 1993) may have indicated. Second, they argue that though causal information is important in making category judgments, so too is surface-level similarity, as it often indicates that critical structural changes have taken place. Thus, a model of categorization must be able to incorporate both kinds of information.

The theoretical and empirical work described in this section argues against the claim that similarity explains categorization. Evidence in favor of these objections is compelling – but evidence against these objections has also been found. Before making firm conclusions, however, it is necessary to consider in more depth what we mean by similarity and how it is

computed in the mind. Many of the objections detailed in the previous section concern a general lack of specificity with regards to what is meant by similarity. As Goodman frames it, “[similarity], has indeed, its place and uses, but is more often found where it does not belong, professing powers it does not possess,” (1972, p. 437). Thus, it is important to consider both where similarity *should* belong, and what powers it actually has. The following section briefly reviews influential models of similarity processing and introduces Structure-Mapping Theory (Falkenhainer et al., 1989; Gentner, 1983, 2003). Using this framework, it is argued that while similarity may not always be a necessary condition for categorization, elements of the structure-mapping process are involved in categorization.

Models of Similarity: Comparison through Structure Mapping

Many models have been proposed to account for processing and perception of similarity. These have included ones that compute similarity as the distance between concepts represented as points in high-dimensional space (Coombs, 1952; Shepard, 1974, 1974), as well as ones that map similarity as a weighted function of the difference between shared and unique features between two objects (Tversky, 1977; Tversky & Gati, 1978). Many of the arguments against similarity and categorization reviewed in the previous section were responses to models like these. However, some theorists have argued that a framework for similarity that includes not only object features, but also relational structure is better able to capture the relationship between similarity and categorization (Gentner & Medina, 1998; Goldstone, 1994; Medin et al., 1993). This third class of model assumes that concepts are structured representations and considers the importance of comparison and alignment in determining similarity (Gentner,

1983; Gentner & Markman, 1997a; Holyoak & Thagard, 1989; Hummel & Holyoak, 1997; Larkey & Love, 2003). This work will focus on one of the most prominent structural models, Structure-Mapping Theory (Gentner, 1983), as a framework from which to respond to these objections. The following section describes this framework in more depth, arguing that representing similarity as a process operating over structured representations makes a link between categorization and similarity more defensible. Adopting the structure-mapping process makes it possible to apply similarity processing to theory-based categories. This motivates the present experimental work which investigates whether the *process* of structural alignment, as opposed to the *product* – perceived similarity – is involved in learning and applying categorical knowledge.

Structure Mapping Theory

Feature-matching models work well for explaining surface similarity; however, these models are unable to account for the full range of similarity perception. Specifically, feature overlap cannot explain analogical similarity. For instance, understanding the similarity expressed in the analogy “A battery is like a reservoir,” has less to do with the features shared by batteries and reservoirs and more to do with the shared relation: batteries and reservoirs both store a resource. Structure-Mapping Theory (SMT) provides a theoretical framework that can apply to analogies like these (Gentner, 1983) as well as to other kinds of similarity. Instead of representing concepts as vectors or sets of features, SMT assumes structured representations³

³ Object-level attributes are represented as single-argument predicates (e.g. RED(x)). First-order relations are represented as multiple-argument predicates (e.g. STRIKE(x, y) or BETWEEN(x, y, z)) while second-order (and higher) relations take other relations as arguments (e.g. CAUSE[COLLIDE(x, y), STRIKE(y, z)]).

that predicate over object-level arguments. A great deal of evidence has demonstrated the utility of SMT in understanding domain-general similarity processing (Christie et al., 2016; Gentner & Markman, 1997a; Jones & Love, 2007; Markman & Gentner, 1993a; Paik & Mix, 2006; Sagi et al., 2012, inter alia). Adopting structure mapping as a model of similarity processing makes a link between categorization and similarity more defensible.

One factor that makes structure-mapping theory a good candidate is its specific process-level claims. Structure-mapping theory has been formalized in a computational model, the Structure-Mapping Engine (SME) (Falkenhainer et al., 1989). SME functions in a local-to-global fashion in the mapping process, first finding all possible *local matches* between the individual representational elements of two potential analogs. These matches are then combined into *kernels* that enforce structural consistency. Kernels are then combined into an overall mapping, preserving the largest and most deeply connected structures (capturing systematicity), using structural evaluation scores. SME has also been incorporated into several broader cognitive models (Forbus et al., 2017).

Though mapping may sometimes be an intentional, goal-directed process, it can also occur as a result of perceptual matches in the environment. The mapping process includes two sub-processes: *structural alignment* and, if appropriate, *inference projection*. For the present purposes, only the alignment process will be discussed in depth since it is most relevant to the current research.

Structural Alignment

Structural alignment is concerned with the identification of shared structure between two concepts. It is guided by tacit constraints favoring *systematic* mappings with maximal (or near-

maximal) *structural consistency* (Falkenhainer et al., 1989; Gentner, 1983; Markman, 1997). Detailing the constraints of structural alignment not only provides a partial response to the objection that similarity is over-flexible or unconstrained, but also provides a set of expectations against which categorization may be measured. To the extent that alignment is involved in similarity processing and categorization, interventions with these constraints in mind should affect the results of both processes.

Structural Consistency

Structurally consistent matches must satisfy two constraints. The first of these constraints is *parallel connectivity*: if two predicates correspond with one another, their arguments must in turn correspond with each other, and play identical roles in the representation. For example, when Spellman & Holyoak (1992) asked people to draw analogies between the then-contemporary Operation Desert Storm and World War II, they found that despite variation in the preferred mappings, people generally maintained structural consistency within those mappings. That is, those who matched George Bush with FDR typically went on to pair the United States during Desert Storm with the United States in WWII. Those who matched George Bush with Winston Churchill went on to map the United States during Desert Storm to Britain during WWII.

The second constraint is *one-to-one correspondence*. This constraint requires that each element of one representation may match with, at most, one element of the other. One-to-one correspondence has not been a constraint in all models of analogical similarity processing, however. Holyoak and Thagard's (1989) ACME model of analogical in fact allowed many-to-one matches as well. Holyoak and colleagues argued that people naturally make one-to-many

mappings (e.g., Spellman & Holyoak, 1996). Markman (Gentner & Markman, 1997b; Markman, 1997) argued to the contrary, that results seeming to show one-to-many mappings in fact reflect that analogical mapping can result in two (or perhaps three) them structurally consistent alternate interpretations. This dispute was resolved only after a carefully conducted follow-up (Krawczyk et al., 2005) concluded that the one-to-one constraint does indeed hold. Thus, there is some consensus across the field of analogical reasoning that the mapping process can produce a small set of structurally consistent alternative mappings. However, which mapping is accepted depends on various evaluative factors. (See Appendix A for more details.)

Systematicity

In addition to structural consistency, alignment is guided by *systematicity* (Falkenhainer et al., 1989; Gentner, 1983). All else being equal, people tend to prefer, larger, more deeply connected systems of matches over local, unrelated matches in the mapping process. This means that when available, people will prefer systems of relations containing higher-order constraining relations, such as causal relations (Clement & Gentner, 1991). Gentner, Rattermann and Forbus (1993) asked people to rate both the similarity and soundness of pairs of stories with shared events. Soundness – the extent to which inferences could be made from one story to the other – was used to assess the relational similarity between the two stories. Half the stories shared higher-order constraining relations linking the events into a system, the other half did not. The systematicity principle predicts that stories with shared higher-order structure should be rated more sound than those with just surface-level matches. Participant ratings confirmed this prediction. Further, and perhaps surprisingly, pairs that shared higher-

order relations were also rated as more *similar* overall than those with shared surface features alone. Higher ratings for both similarity and soundness indicate a general preference for coherent, systematic mappings, as well as mappings with inferential potential.

There is also evidence that comparison through structural alignment




Figure 1: Sample stimuli from Medin, Goldstone & Gentner, 1991

can highlight systematic relational structure (Markman & Gentner, 1993b). Participants received a one-shot mapping task over pictures portraying causal scenes. The experimenter indicated an object in one picture and asked the participant to indicate the corresponding object in another. Each pair contained a cross-mapping where an object in the scene played a different role, despite sharing identical surface features with another. For example, in one picture a woman was pictured feeding a bird, while in the other, a woman *received* food from a delivery person. This presents an appealing surface match in tension with the relational match connected to the causal structure of the scene; the woman in the first picture and the delivery person in the second are both the agents of the *giving* in the scene. The key manipulation was whether participants had an explicit opportunity to compare the two scenes and rate their similarity before the mapping task. When participants compared, they were more likely to select the relational match (e.g., woman → delivery person) over the object match (e.g., woman → woman). Markman and Gentner concluded that comparison exerts a notable effect on what kinds of commonalities are relevant for similarity – and more specifically, that people tend to favor interpretations based around common relational structure.

Structural alignment not only leads to the detection of systematic relational similarity. Alignable structure can also influence how people conceptualize parts of the objects being compared. For example, Medin, Goldstone & Gentner (1993) asked participants to compare sets of drawings where some objects had ambiguous properties. As seen in Figure 1, object B can be interpreted either as two triangles with a line on top, or three triangles: two upright and one inverted. Participants compared either A and B or B and C, then listed the similarities and differences between the objects. On average, participants listed more properties consistent with the compared object (e.g. “two triangles” in an A-B comparison vs. “three triangles” in a B-C comparison), than comparison-inconsistent properties. This indicates that shared structure has a marked effect on how concepts are represented during similarity processing.

Evidence from Sagi, Gentner and Lovett (2012) indicates that shared systematic structure influences not only the processing of similarity, but also the processing of difference.

They showed that people are faster to discriminate two different things if they are highly dissimilar but, paradoxically, faster to identify *specific* differences when the two items are nearly identical except for *alignable differences* – those occupying corresponding positions in their respective relational structure. Participants compared two images and performed a same/different discrimination task and a difference identification task. In each block they




Figure 2: Sample stimuli from Sagi, Gentner, and Lovett, 2012

received pairs of shields that were either highly similar with one alignable difference, or very dissimilar. For instance, A and B (Figure 2) have one alignable difference (the color of the central circle), while A and C have many differences. As expected, participants were faster to differentiate highly dissimilar pairs, but faster for the alignable pairs. This dissociation supports the idea that constraining systematic structure is involved in both similarity as well as difference.

In sum, during structural alignment the commonalities and differences between two concepts are determined through the alignment of structured representations. Guided by the constraints of structural consistency and systematicity, structural alignment determines the most optimal match. Structure-mapping theory provides a process-specified theory of how similarity is processed and determined. The findings reviewed in this section demonstrate how this framework explains constraints on the processing of similarity and includes relational structure as well as featural overlap – thus permitting application to theory-based categories. The following section argues that the sub-process of structural alignment may be applied to categorical knowledge representations in several ways. If this is so, these categorization processes must also be subject to the constraints that guide the alignment process, laid out above.

Structural Alignment and Categorization

This section applies structure-mapping principles to the recognition and extension of categories. To preview, the first part of this section considers how structural alignment may operate on categorical knowledge representations, focusing on two areas: how categories are

formed, and how categories are *applied*. The second part of this section reviews evidence that comparison through structural alignment can influence the learning and application of *relational categories*, which may serve as an ideal test case for the claim that structural alignment is involved in the learning and extension of categories.

Two Roles for Structural Alignment in Categorization

Extrapolating from the preceding evidence on similarity processing, there are at least two ways structural alignment may be involved in categorization: the *formation* of new categories and the *application* of categorical representations to newly encountered potential cases. We might assume that, when forming a new category, and in the absence of any top-down indication of category structure, people align potential category members to each other in order to determine shared diagnostic elements of the category. Similarly, once someone has a well-structured representation for a category, we might assume that they would also align newly encountered category candidates with that representation. Alignment from representation to instance would be involved both in evaluating the category membership of a new instance and in updating the category representation if the new instance belongs to the category. The following sections provide evidence in favor of the view that each of these operations involves structural alignment, either between category members, or between summary category representation and evaluated cases. This evidence grounds the current experimental work.

Forming Categories through Alignment

Structural alignment facilitates the formation of categories by aiding in the recognition of shared commonalities among category members – especially weighting shared systems of commonalities. Indeed, some evidence indicates that alignment could play an important role in category learning (Gentner & Namy, 1999; Lassaline & Murphy, 1998; Namy & Gentner, 2002). For example, Lassaline and Murphy (1998) provided evidence that categorization involves alignment between potential category members and category representation. Their goal was to contrast an alignment model of categorization with feature-based models (Rosch & Mervis, 1975; E. E. Smith & Osherson, 1984; Tversky, 1977), attribute/value models (Estes, 1994; Medin & Schaffer, 1978) which may be construed as including some degree of alignment. Feature models assign objects to a category by considering the overlap of shared features between an instance and a concept. Attribute/value models assume that rather than total feature overlap, representations are compared according to feature dimensions. For instance, Medin and Shaffer's context model (1978), determines similarity by counting overlapping features within separate dimensions. Each model has different predictions for categories with different structures. Feature models include no role for dimensional alignment and predict that overall feature overlap will aid learning. Attribute/value models predict that only alignable *matches* on category-relevant dimensions facilitate learning, and that alignable *differences* on category-irrelevant dimensions have no effect on categorization.

Lassaline and Murphy (1998) gave participants a category learning task involving categories with underlying structure. Each category consisted of four drawings of birds which varied on four dimensions: head pattern, wing pattern, body pattern, and tail pattern. Each

dimension could have one of eight values: white, spotted, striped, grey, checkered, scalloped, circles with dots, and black. Category structure varied across experiments to test different models. In the alignable category structure, categories were determined by dimension-value relations (e.g. dotted head *and* black back). In the non-alignable structure, category membership was still determined by features in the same diagnostic dimensions, but values were not connected to dimensions (e.g., having a head or back that was either dotted or black).

In each experiment participants performed a supervised category learning task where they placed each drawing into one of two categories and received feedback on their choice, which was repeated until they showed no errors. They then performed a timed test block without feedback. The key prediction concerned which category structures would show faster category learning. In experiments 1 & 3 participants took significantly longer to reach criterion for the non-alignable match category than for the alignable category and produced more errors during learning. The attribute/value model would predict a learning advantage for alignable categories over non-alignable categories; however, a feature model would predict no difference, since both categories share the same overlap of features. Experiments 2 & 4 contrasted the non-alignable feature match category with another “no-match” condition where each category included one exemplar that did not include critical features on the diagnostic dimensions (though those features were still included in non-diagnostic dimensions). Alignment and Attribute/value models would both predict no difference between these categories since only matches on the *same dimension* count toward similarity. Alignment and featural models, however, would predict faster learning for the non-alignable feature match category versus the no-match category since both models allow for features on non-diagnostic dimensions to

influence categorization. Indeed, this pattern of results was found, with participants learning the non-alignable match category faster than the no-match category and producing fewer errors in the learning phase. Taken together the results of these experiments support a structural alignment model of category learning. People learn categories faster when diagnostic features occur as part of alignable structure, compared to categories without alignable differences. Matches that are not part of alignable structure do facilitate category learning, however, when compared to categories with exemplars that contain no matches.

The Sequential Generalization Engine (SAGE, Forbus et al., 2017) provides a computational model of how structural alignment may be applied in this way. SAGE uses SME (Falkenhainer et al., 1989) to model category learning as progressive abstraction over exemplars. SAGE stores structured descriptions that are generated by generalizing across encountered exemplars using SME. In addition to this set of generalizations it also maintains a set of unincorporated exemplars. When the system receives a new potential case, it is compared using SME, first to the set of existing generalizations, and then to each of the unincorporated exemplars. If it matches with one of the prior generalizations, it is assimilated into that generalization. If it matches with one of the unassimilated exemplars, they both form the basis of a new generalization. Otherwise, the new exemplar is stored with the remainders. As it employs SME, this process is bound by the constraints of structure-mapping theory, including the systematicity principle, which is used to compute a structural evaluation score for each match. The model includes a threshold parameter which determines how conservative the system will be in its generalization, with extremely high values leading to no abstraction (since

the system would require *exact* matches between exemplars) and/or extremely promiscuous generalization (i.e., a global generalization containing every exemplar).

Kuehne et al. (2000) used a precursor of SAGE, SEQL, to model the results of a human categorization study by Ramscar and Pain (1996). Ramscar and Pain investigated the role of structural similarity in categorization by having participants sort stories adapted from (Gentner et al., 1993). The stories varied in their featural and relational overlap. Each variation was derived from an original base story and shared degrees of first-order relations, second-order relations, and object attributes. Participants were given stories one at a time, were asked to familiarize themselves with the stories. They were then asked to “group the stories into categories that seem the most natural and appropriate to you.” While participants could choose any potential grouping (in fact 10 types of groupings were identified by Ramscar and Pain), there were 5 that exceeded 3% of responses. The most common grouping (79.5%) included the base story with literally similar and analogically similar variants and grouped the story variants with only shared surface features together. The next most common grouping (8%) preserved the same groups but kept the base story separate. The next most common groupings were based on shared object matches (5%) or first-order relations (4%). As Ramscar and Pain (1996) noted, this pattern suggests a close connection between analogical processing and category learning. Kuehne and colleagues (2000) used the SEQL model to approximate these grouping results by feeding the generalization engine structured representations of the stories and setting the model to different generalization thresholds. At low threshold values the model showed hyper-generalization, creating a generalization that included every example, while at high values the model was hyper-discriminatory, forming no generalization. However,

at threshold values from approximately 0.85 to 0.98, the types of generalizations showed the same pattern as the human results.

Co-categorizing Representations through Alignment

Structural alignment may also be involved when applying stored category representations through co-categorization. This includes both determining whether a newly encountered instance is a member of a known category as well as updating a category representation to account for distant exemplars. Developmental evidence indicates that children are better able to co-categorize objects into a common category when they are easier to align.

Shared categorical language often serves as a cue to co-categorize through symbolic juxtaposition (Davidson & Gelman, 1990; Gelman, 1989; Gelman & Markman, 1986; Gentner & Medina, 1998; Gentner & Rattermann, 1991). This can be especially useful as children are learning and updating conceptual representations. For instance, Gelman and Markman (1986) found that 4-year-old children were more likely to import knowledge from one creature to another when they were given common labels (e.g., 'bird). However, children are also better able to extend categories when explicitly told to align same-category objects. Gentner and Namy (1999) gave 4-year-olds sets of color drawings of real objects. Each set consisted of two same-category standards and two alternatives: a perceptual match, and a category match. For instance, children might be shown a bicycle and a tricycle as the standards (each sharing the perceptual feature of adjacent circles in the form of the wheels), and choose between a category match (a skateboard, another form of transportation) and a perceptual match (a pair of glasses with the lenses in the same configuration as the wheels). Children were assigned to

either the compare or no-compare group and either the word or no-word group. The no-compare group only saw one of the standards, while the compare group saw both. In the word group the objects were given a common label (e.g., “This is a *blicket* and this is a *blicket.*”), while in the no-word group the experimenter presented the objects without naming them (e.g., “See this one, and this one?”). Children were more likely to choose the same-category match when they received the labels; however, they were also more likely to choose the same-category match when they were told to compare. Importantly, the children that received a label, but did *not* compare two standards, only selected the same-category match at chance. These findings indicate that even though common labels can function as invitations to form categories, this may largely be due to the alignment of multiple category exemplars.

This is not limited to object concepts, but also more abstract relations. In an analogous task Christie and Gentner (2010) gave 3- and 4-year-olds labels for novel spatial patterns and asked them to extend those labels to one of two alternatives: an object match and a relational match. For instance, a child might see a pair of pigs facing each other (i.e., demonstrating the spatial relation SYMMETRICAL(pig₁ , pig₂)). The experimenter would label each arrangement of animals using a novel count noun (e.g., “Look, this is a jiggy”). The child would then be shown two alternatives and asked to choose which had the same label (e.g., “Can you tell me which of these two is a jiggy?”). One alternative was a featural match including one pig in a different spatial relation, while the other alternative would be two different animals in the same relation. When children only saw one example of the relation, or two examples sequentially they showed a general preference for the surface match; however, when they compared two instances, they were much more likely to prefer the relational match. This experiment is

especially illustrative since it demonstrates both applications of alignment in the same task. Alignment seems to help the children to *form* a category referred to by the relational label, but the use of the label in reference to a new case helps children *apply* that category representation better and select the relationally similar case.

To summarize, the evidence described in this section supports the view that structural alignment operates not only in category formation, but also as category representations are applied in co-categorization. The importance of alignment and comparison is especially evident when learning more abstract relational categories. The final section considers an ideal case where these claims have been tested, and that is the focus of the current research: relational categories.

Structural Alignment and Relational Categories

In recent years researchers have increasingly recognized that categories show a range of representational content, from *ad hoc* categories organized around single features or contexts (e.g., *red things*, or *things to take on a picnic*), to basic level categories that share perceptual and relational content, to abstract, relational categories that share few overlapping features (Gentner, 2005; Gentner & Kurtz, 2005; Markman & Stilwell, 2001). Indeed, relational categories serve as an ideal arena in which to test the claim that alignment is involved in categorization.

What Are Relational Categories?

The study of relational categories as a subtype of category structure is relatively new. The most accepted definition takes relational categories as “a category whose membership is determined by a common relational structure, rather than by common properties” (Gentner & Kurtz, 2005, p. 3). In the case of a basic-level category like *cats*, category members share both common features (e.g. four-legged, whiskers) as well as relations (e.g., eats fish, hunts mice). Members of relational categories, on the other hand, share only relational commonalities – and not necessarily features. For instance, in the case of *predators* the only requirement for class-inclusion is attacking or hunting something else⁴ (usually for sustenance).

The type of relational commonalities that define a given relational category determines whether it falls into one of the two most widely-accepted subtypes: *schema-governed categories* and *role-governed categories* (Goldwater et al., 2011; Markman & Stilwell, 2001). Whether the category denotes a relation itself, or the role in which members participate largely determines subtypes of relational categories. Schema-governed categories specify entire relational systems (Gentner & Kurtz, 2005). These can range from single concepts (e.g. *support*, *reciprocity*), to sets of events and their related arguments (e.g. *wedding*, *robbery*). Role-governed categories are closely related to schema-governed categories. They typically denote the participants in a particular role in that schema (Gentner & Kurtz, 2005; Markman & Stilwell, 2001). For example, the schema-governed category *robbery* might include three role-governed

⁴ Typical predators might be animals that prey on other animals, (e.g., sharks, lions, or wolves), but the category can include much more far-ranging and abstract exemplars (e.g. corporations, conmen, or computer hackers). Indeed, Kurtz and Gentner (2001) found that when asked to generate exemplars of relational categories (e.g. *barrier*), people often first listed concrete members and then went on to list abstract members.

categories: *thief*, *victim*, and *stolen goods*. Not all arguments need be explicitly mentioned for every robbery; however, we recognize that they are part of the relational structure of the schema-governed category *robbery*. For example, if we hear “The thief stole many valuable paintings,” we expect there was a prior owner who was the victim of the robbery – even if they are never talked about. The sole requirement for something to be considered a member of the *thief* category is participation in the agent role (x) of the relation $STEALSFROM(x, y, z)$.

Though the preceding examples indicate that relational categories may vary in which relation constrains the category, they are all nonetheless organized around shared relational content. This is what makes them especially interesting for the present purposes. Indeed several researchers have considered how alignment functions with relational categories.

Alignment and Relational Category Learning

Relational categories have received increased attention largely due to their unique psychological behavior compared to other entity categories. They are slower to be learned by children (Gentner, 2005), more likely to be described in terms of ideals rather than prototypes (Goldwater et al., 2011), and are generally more difficult to learn (Kurtz, 2015). However, evidence has shown that comparison through structural alignment is one way that relational category learning can be made easier.

For instance, Kurtz, Boukrina and Gentner (2013) found that comparison facilitated the learning and transfer of schema-governed categories. Participants completed a feedback learning task where they were learning to classify different arrangements of rocks into three categories. Participants in the single-item condition saw each arrangement one-by-one and

were asked to place it in one of the three categories. Participants in the comparison condition saw two arrangements side-by side and were asked to classify one arrangement and then the other. After the learning phase, participants then performed two transfer tasks: one where they classified newly encountered items mixed in with the ones they had seen before, and a far transfer task where they applied the same category labels to mobiles with roughly the same spatial configuration as the rocks. The comparison group outperformed the single-item control group on both the near transfer and the far transfer task, suggesting that alignment aids in the learning and transfer of schema-governed categories.

Goldwater, Bainbridge and Murphy (2016) found similar evidence for role-governed categories. They created novel stimuli where three objects performed different roles in a schema event. For instance, one object might chase another object (each filling the role of chaser and pursued respectively), while the third object observed. After watching a series of videos where the objects performed the actions, participants were given a classification task where they were shown one object and asked to select which of the two objects best went with it to form a category. There were three kinds of test triads: *role triads* where the standard and its match performed the same role, *thematic triads* where the standard and its match performed reciprocal roles, and *conflict triads* where participants had to choose between both potential matches. Participants generally showed a preference for role matches in conflict trials, but across experiments common labels increased this preference. When the role relations were made more abstract, participants failed to group same-role items if the specific relations differed (e.g., objects that prevented different actions). However, participants were better able to identify same-role objects if they compared multiple films beforehand.

The evidence presented here suggests that structural alignment improves the learning and transfer of relational categories. This may provide a computational framework for the formation of category-level generalizations. However, the bulk of this research has been focused chiefly on relational category learning alone. Additional research must consider the role of alignment in what have typically been considered categorization operations, and how well this approach extends to the full breadth of relational categories.

Summary

So, what may be salvaged from the hypothesis that similarity is the basis of categorization? What must be abandoned? Given the literature reviewed in the previous sections, it is clear the strong version of this hypothesis cannot be accepted. Similarity is not *sufficient* for categorization to occur; certainly, neither overall similarity nor featural overlap (the default kinds of similarity that have typically been assumed in critiques of the similarity approach) are *necessary* for categorization. One might then conclude that the intuition that categorization is closely related to similarity, which has formed the basis of many influential approaches in the field, is simply an overgeneralization of the observation that *some* categories contain highly similar members. However, there may be an alternative conclusion. The research reviewed in the previous sections suggests that the strong hypothesis in fact conflates a distinction between process and product that must be analyzed further. Similarity judgments do not determine categorization judgments; but it is suggested that both judgments draw on the same computational processes. One candidate for this process is comparison through structural alignment, which plays a key role in the assessment of relational similarity as well as overall

similarity. It is argued that rather than relying on the outcome of a similarity judgment, people use the underlying process of structural alignment to identify same-category matches.

As summarized above, structural alignment has already been applied to the study of relational category learning and these categories may be especially useful in the study of how people apply categorical knowledge. The current experiments consider whether comparison through structural alignment plays an essential role in categorization using one kind of relational category – logical fallacies – as a test case. The basic claim is that structural alignment processes are integrally involved in categorization as well as in similarity judgments. The next section summarizes the three experiments presented in this paper and provides an overview of the specific hypotheses.

Overview of Experimental Work

The experiments presented here use informal logical fallacies (e.g., *strawman* or *post hoc ergo propter hoc*) to investigate the shared role of alignment in similarity and categorization. A basic assumption of this work, then, is that logical fallacies function as schema-governed relational categories that are likely relatively unfamiliar to participants. As a result, participants are unlikely to have pre-stored category representations for them. That is, individual passages featuring different characters and plot can be classified as instances of the same fallacy insofar as they satisfy certain relations between the premises, or the premises and the state of the world. For example, consider the passages in Figure 3. While each passage has different surface features, involving different characters, and come to conclusions specific to each scenario, both

arguments involve a premise that uses authority to prove its conclusion (Mrs. Blavatsky in one case, the King in the other).

A: *“Karen’s mother is convinced there are ghosts haunting their home, causing noises in the middle of the night. Karen has tried telling her that there’s likely a more rational explanation than spirit energy. Her mother, however, maintains that it must be ghosts because her upstairs neighbor Madame Blavatsky has examined the house, and she must know.”*

B: *“Robert and Liam are engaged in a heated debate about which end should be down when storing eggs, the wide one or the narrow one. Robert argues that the wide end is more stable and therefore better to store the egg on. Liam, however, has heard that the King stores his eggs on the narrow end, and after all, he’s the King.”*

Figure 3: Two passages exemplifying the same logical fallacy: Appeal to Authority

All experiments used a two-alternative forced choice task (2AFC) as the primary outcome measure. Influenced by Rips (1989) and others, the *baseline* condition was used across all experiments. Participants were presented with a paragraph that exemplifies a specific logical fallacy (the standard) and asked to choose between two additional paragraphs, each of which is also an instance of a logical fallacy (the alternatives). In each experiment, one alternative is an instance of the same fallacy as the standard, and the distractor is an instance of a different fallacy. The specific instructions, as well as some aspects of the alternatives and standard, vary across experiments; however, all experiments contrast same-fallacy response between an explicit categorization task and a similarity task. Several factors were considered when designing the instructions for the categorization and similarity tasks. One concern was how the instructions refer to the fallacies. Pointing out their fallacious nature may inadvertently bias participants on the categorization task; failing to point out their status as *kinds* of things could have the opposite effect. In the end it was decided to refer to the passages as errors in reasoning. This had the advantage of referring to them as a kind without specifically referencing fallacies and calling to mind specific instances with which participants may be familiar. In this version of the 2AFC task participants choose between a surface similar distractor and relational

match. It is predicted that, as in the experiments that motivated this research, participants will show higher same-fallacy response for the categorization task compared to the similarity task when in the baseline condition.

Structural alignment may be involved at several points in this task. First, participants are expected to align (or attempt to align) the standard with each of the alternatives as they attempt to differentiate them. Second, when they can, participants will align (or attempt to align) the alternatives with each other. As described previously, evidence of alignment in similarity processing is well-established. However, we would expect that participants will align the standard with each alternative and extract shared relational structure – especially if participants do not have a stored representation for a fallacy category. Finally, if participants do have a stored category representation, we expect they align their stored category representations with both the given standard and the alternatives.

Experiment 1 contrasts the baseline condition with an alignable alternatives condition and a neutral condition. When participants are choosing between alignable alternatives, it is hypothesized that they will show the highest same-fallacy response on both the similarity and categorization task compared to neutral and baseline conditions.

Experiment 2 introduces a pre-task phase where participants either compare two same-fallacy examples or see them sequentially. Based on prior research (Christie & Gentner, 2010; Loewenstein et al., 1999), it is expected that participants who do the comparison task will abstract the common schema, thus facilitating alignment in the subsequent 2AFC task. Thus it is predicted that there will be higher same-fallacy response on both the similarity and

categorization task for the participants in the pre-task comparison group compared to the pre-task sequential group or a group that receives no pre-task manipulation.

Whereas the first two studies involve manipulations predicted to facilitate alignment, Experiment 3 involves a manipulation predicted to hamper alignment. In this study performance on the baseline task is compared to the performance of a group where comparison between the standard and alternatives is made more difficult. When alignment is more difficult, it is predicted same fallacy choice should decrease compared to baseline on both the similarity task and on the categorization task.

All three experiments test the following main hypothesis:

1. If alignment is involved in categorization as well as similarity, then manipulations which promote or hamper alignment should not only affect similarity judgments, but also categorization judgments.

That is, depending on the manipulation, participants in both tasks will either increase or decrease in their likelihood of selecting the same-fallacy alternative compared to their choices in the baseline control condition. However, there is no expectation that participants will be just as likely to select the same-fallacy response for both tasks. Indeed, as mentioned above, previous evidence has indicated that categorization and similarity judgments diverge (Rips, 1989; but see Hampton et al., 2007). Thus, the following hypothesis functions as an empirical generalization of this work:

2. When surface and relational similarity compete in the baseline condition, participants will show higher same-fallacy response on the categorization task compared to the similarity task.

Finally, the logic of the current research hinges on the established role of alignment in similarity processing. In each experiment a technique is imported that has been found to either foster or impede alignment. Thus, for each experiment it is important to keep in mind that if response on the categorization task is unaffected by the alignment manipulation, this could have two explanations. It is possible that alignment in fact plays no role in categorization; however, it is also possible that a manipulation shown to influence alignability on similarity tasks simply does not have the same effect for our materials as it has in prior studies.

Analysis

All data reported in these experiments were cleaned and analyzed using the R environment (R Core Team, 2017). Given the repeated measures nature of the data and the need to account for variance across items, the lme4 package (Bates, Mächler, Bolker, & Walker, 2014) was used to fit generalized linear mixed effects models to predict same-fallacy choice. Logistic regression with a binomial distribution was used to assess the relationship between the predictors and participant choice since the task involves a binary outcome measure. Instruction condition was entered as a single fixed effect in the model. Following recent recommendations (Barr, 2013; Barr et al., 2013), crossed random effects for participants and fallacy type were included in each model. Confirmatory hypothesis testing has been notoriously difficult to interpret using mixed

effects modelling techniques (Luke, 2017). However, as has become recent convention, Satterwaithe approximation for denominator degrees of freedom was used via the lmerTest package (Kuznetsova, Brockhoff, & Christensen, 2017) for all experiments where significance tests are reported. All pairwise comparisons were generated using the emmeans package (Lenth et al., 2020), which allows for the estimation of estimated marginal means.

Experiment 1

Experiment 1 has two purposes. First, to confirm the baseline prediction (Main Hypothesis 2) that under conditions in which surface-similarity conflicts with relational structure, similarity and categorization lead to different outcomes (Rips, 1989). Thus, it is expected that in the baseline task described above, participants will be more likely to select the same-fallacy response on the categorization task compared to the similarity task. The second purpose of Experiment 1 is to test the prediction that increasing the alignability of the alternatives will increase relational response on *both* categorization and similarity tasks.

One way this might occur is by making it easier for participants to align the two alternatives in the 2AFC task. Some inspiration for this comes from previous work by Shao & Gentner (2019). Children in ages ranging from three years old to nine years old were given a perceptual 2AFC task. Given a standard featured two geometric shapes children were asked to select between two alternatives. On test trials the correct match showed the two objects in a symmetrical formation. They found that children didn't perform above chance until they were 8-9 years old; however, when the incorrect match was made highly surface-similar to the correct match by featuring the same shapes, 3-4-year-olds chose the correct match at rates

similar to their 8-9-year-old counterparts. They hypothesized that this manipulation bootstrapped children's relational insight by promoting alignment between the two alternatives. This not only has the potential to highlight the relational difference between them, but also de-emphasizes the role of object matches.

A similar manipulation is employed in this experiment. However, there is an important difference in this experiment. In Shao & Gentner's experiment the incorrect match was never similar to the standard in any way. In this experiment the baseline task always includes a surface-similar distractor. Thus, any increased same-fallacy response in the alignable alternatives condition could be attributed to the simple lack of a highly competitive match. As a result, an additional control condition was included in this experiment in order to account for this alternative explanation.

Method

Participants

155 participants were recruited for this experiment. Participants received paid compensation for their participation (\$12) and were a mix of Northwestern University students and residents of the Chicago metro area. 11 participants were excluded for failing to complete catch trials, leaving a total sample size of 144. Participants were tested on computers and were run alone or in groups of no more than four.

Materials

12 paragraphs were created, each an instance of a different type of fallacy. These paragraphs were designed to serve as the standards against which participants would judge the alternatives. Three sets of alternatives were created for each standard, always including the same-fallacy target, and a different-fallacy distractor. In the *baseline* set the distractor shared surface features with the standard. In the *alignable alternatives* condition the distractor shared surface features not with the standard, but with the same-fallacy match, making it easier for participants to align them. This meant that each alternative was nearly identical except for the final sentence or two where the fallacious conclusion was proposed. Each alternative also came to the same conclusion – though using different fallacious reasoning. In the *neutral distractor* set the distractor shared no commonalities with either the standard or the same-fallacy match. The full set of Experiment 1 stimuli can be found in Appendix A. Four catch trials were included to serve as an attention check. In these triads the distractor also shared no similarities whatsoever with the standard, while the same-fallacy match also shared surface similarities. Thus, participants should have no reason to select different-fallacy distractor in the catch trials.

Procedure

Materials were presented using the Qualtrics survey presentation platform. Participants were randomly assigned to either the categorization or similarity task and one of the distractor groups. Thus, Experiment 1 had a 2(Task: categorization, similarity) x 3(Distractor: baseline, alignable, neutral) between-subjects design. Participants in both task groups received an example with general instructions, and then proceeded to the main task. Participants in the

categorization group were given the following instructions: *“In this study you will see an example of a reasoning error, along with two additional passages. Your task is to say which of the two passages belongs to the same category as the example.”* Participants in the similarity group received the following instructions: *“In this study you will see an example of a reasoning error, along with two additional passages. Your task is to say which of the two passages is the most similar to the example.”* This wording was arrived at in an attempt to strike a balance between the two task conditions. One potential drawback is that similarity participants may be biased toward the same-fallacy response, by virtue of calling the fallacies “reasoning errors.”

For each triad participants were shown the standard with each alternative presented below and asked to click on the paragraph they thought best matched the standard. Side of presentation was left/right counterbalanced so that the same-fallacy match was not always on the same side of the screen. Each triad was presented on a separate screen, with the instructions and the standard on the top, and the two alternatives below. Order of presentation was randomized, and the alternatives were left-right counterbalanced in a semi-random fashion.

Given this design it is hypothesized that participants will show the highest same-fallacy responding in the alignable alternatives group, compared to the other two. This should result in the following pattern of same-fallacy responding: Alignable > Neutral > Baseline. Thus, the fine-grained predictions for Experiment 1 are as follows: A) same-fallacy responding on Neutral > Baseline for both categorization and similarity; B) same-fallacy responding on Alignable > Baseline for both categorization and similarity; and C) same-fallacy responding on Alignable > Neutral for both categorization and similarity. However, same-fallacy responding on these tasks

should differ in the baseline group, where participants in this condition are expected to show higher same-fallacy responding on the categorization task than on the similarity task.

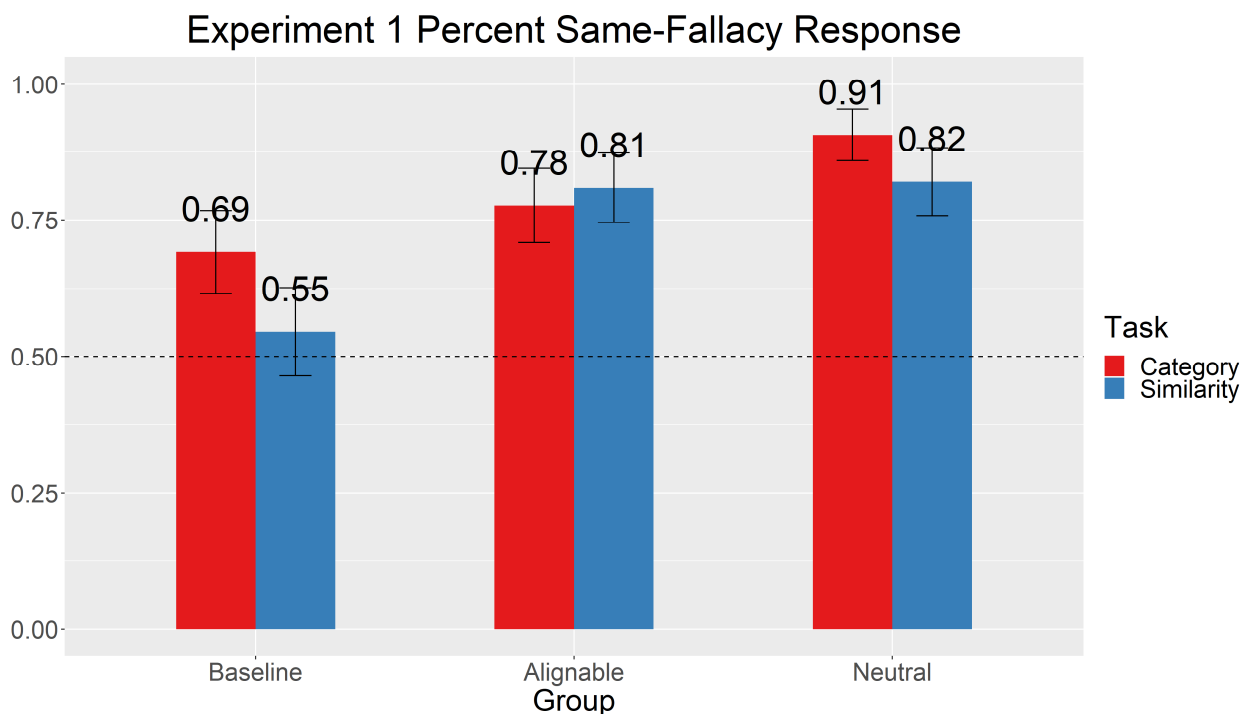


Figure 4: Experiment 1 results

Results

Distractor and Task were included as in a logistic regression model, as well as the interaction term. A main effect was found for Distractor type. Participants in the alignable ($\beta = 1.6$, 95% CI[0.69 - 2.51], $p < .001$), and the neutral distractor ($\beta = 1.67$, 95% CI[0.78 - 2.57], $p < .001$) groups showed higher same-fallacy responding compared to the baseline group. Overall,

Table 1: Experiment 1 model parameters

| <i>Predictors</i> | Response | | |
|------------------------------------|-----------------|--------------|------------------|
| | <i>Log-Odds</i> | <i>CI</i> | <i>p</i> |
| Intercept | 0.24 | -0.48 – 0.97 | 0.512 |
| Alternative (Alignable) | 1.60 | 0.69 – 2.51 | 0.001 |
| Alternative (Neutral) | 1.67 | 0.78 – 2.57 | <0.001 |
| Task (Categorization) | 0.89 | -0.06 – 1.83 | 0.068 |
| Alignable * Categorization | -1.14 | -2.45 – 0.17 | 0.087 |
| Neutral * Categorization | 0.17 | -1.18 – 1.51 | 0.809 |
| Random Effects | | | |
| σ^2 | 3.29 | | |
| τ_{00} Participants | 1.84 | | |
| τ_{00} Fallacies | 0.39 | | |
| ICC | 0.40 | | |
| N Participants | 144 | | |
| N Fallacies | 12 | | |
| Observations | 1728 | | |
| Marginal R^2 / Conditional R^2 | 0.111 / 0.470 | | |

participants that received the categorization task were more likely to choose the same-fallacy response, though this pattern was not significant ($\beta = 0.89$, 95% CI[-0.06 – 1.83], $p = 0.068$). The Distractor x Task interaction was also not significant for the alignable group ($\beta = -1.14$, 95% CI[-2.45 – 0.17], $p = 0.087$) or the neutral group ($\beta = 0.17$, 95% CI[-

1.18 – 1.51], $p = 0.089$).

Post-hoc comparison of the estimated marginal means was used to test the fine-grained predictions listed above. Confidence level and p-value adjustment was done using the Tukey method for comparing a family of three estimates. Confirming prediction A) the neutral group showed higher same-fallacy responding than the baseline group on both the categorization task ($\beta = 1.84$, 95% CI[0.63 - 3.05], $p = 0.001$) and the similarity task ($\beta = 1.67$, 95% CI[0.37 – 2.98], $p = 0.003$). Prediction B) was not fully supported, with the alignable group showing higher same-

fallacy responding only on the similarity task ($\beta = 1.60$, 95% CI[0.28 – 2.93], $p = 0.008$), but not on the categorization task ($\beta = 0.46$, 95% CI[-0.9 – 1.83], $p = 0.93$). Finally, prediction C) was not supported: the alignable group did not show higher same-fallacy responding to either the categorization task ($\beta = -1.38$, 95% CI[-2.77 – 0.12], $p = 0.054$) or the similarity task ($\beta = -0.07$, 95% CI[-1.38 – 1.23], $p = 1.00$). Thus, results only were partly consistent with Main Hypothesis 1 in that participants in the alignable alternatives group showed higher same-fallacy responding than the baseline group on both categorization and similarity tasks. However, these results failed to completely support Main Hypothesis 1 since the alignable alternatives group did not show higher same-fallacy responding than the neutral alternative group.

Finally, post-hoc comparison was used to evaluate the empirical generalization that when surface and relational similarity compete, participants will show higher same-fallacy responding to the categorization task than the similarity task. Contrast of the estimated marginal means in the baseline group found only a trending effect of Task ($\beta = -0.89$, CI[-1.83 - 0.06] $p = 0.07$). This finding does not provide strong support for Main Hypothesis 2. Interestingly, the neutral group did show higher same-fallacy responding on the categorization task than the similarity task ($\beta = 1.05$, CI[0.09 - 2.01] $p = 0.03$).

Discussion

Experiment 1 found mixed support for Main Hypothesis 1. While participants in the alignable alternatives showed increased same-fallacy responding compared to baseline on both categorization and similarity tasks, the alignability manipulation conveyed no greater advantage than the neutral group. Thus, although same-fallacy responding was higher in the alignable

alternatives condition than in the baseline condition, this could have resulted from simply removing the surface distractor. This prediction was based on experiments by Shao & Gentner (2019) using a similar manipulation; however, there may be a readily available explanation for the current pattern of results. The failure to find a benefit from the alignable alternatives manipulation may reflect additional complexity involved in identifying relational categories like logical fallacies, or perhaps from the complexity of the textual materials. The task used by Shao & Gentner (2019) differed in three important ways: the stimuli all involved perceptual similarity rather than conceptual similarity, the task was aimed at 3-9-year-olds and therefore was much simpler, and finally, the design did not include a neutral control group.

The few studies that have shown benefits of alignable differences have been in the arena of visual perception. Shao and Gentner (2019) gave children cards featuring geometric shapes in various colors. The relational match involved symmetrical spatial orientation, and the alignable alternative featured objects in the same shape and color. The stimuli used in the current experiment, by comparison, involve *conceptual* rather than *perceptual* similarity. The relational match involved identifying how the arguments relate to each other and the world. As a result, participants may have been distracted by superficial commonalities and differences, making it more difficult to detect the alignable differences than participants in Shao & Gentner's study. Further, participants may not have read the alternatives in sufficient detail to bootstrap the relation. The two scenarios were presented side by side, in order to promote alignment between them similar to Shao & Gentner (2019). However, participants may have found the closeness of the similarity confusing, and after noticing that the first few sentences or

words were the same, did not completely read the passages, or only skimmed them. Thus, participants may have missed the logic of the scenarios.

Another concern, relating specifically to the baseline condition, is whether the design was under powered at either the participant or item level. The sample size used in Experiment 1 (24 participants per group and 12 items) was selected based on previous simulation work indicating that ML models show reduced Type I error for similar subject/item sample sizes (Luke, 2017). Nonetheless, this was only estimation and could still be underpowered. The question of power will be taken up at length in the general discussion.

Taken together, the findings of Experiment 1 are inconclusive. The findings are partly consistent with Main Hypothesis 1, in that the neutral condition (which lacks a competing surface-similarity match) showed greater same-fallacy responding than the baseline condition for both the similarity task and the categorization task. However, other predictions failed. The alignable-alternatives manipulation did not show the predicted effect: contrary to prediction, participants that chose between highly-alignable alternatives did not show higher rates of same-fallacy responding than the baseline condition for *either* categorization or similarity task. However, a difference was found between the categorization and similarity task for the neutral group. This finding falls beyond the scope of the main hypotheses, but may be interesting nonetheless. It may be a statistical oddity, or it may reflect that the hypothesized divergence between similarity and categorization tasks is obfuscated by the contrast between surface features and relational structure. Categorization participants may have found the task much less ambiguous with the removal of the alignable alternative fallacy. Experiment 2 considers another manipulation that may be more successful.

Experiment 2

Rather than making alignment easier during evaluation, Experiment 2 asked participants to directly align two same-fallacy examples before making their choices. As in previous work (Christie & Gentner, 2010; Loewenstein et al., 1999), participants in Experiment 2 were asked either to compare same-fallacy examples or to view them sequentially before going on to perform categorization or similarity tasks on the baseline set. If within-category alignment helps participants better identify the relational structure of the fallacies, participants that received the comparison manipulation in the pre-task should show higher same-fallacy responding during the test phase than those that received the sequential manipulation. Importantly, this should be true for both categorization and similarity tasks.

There is also an open question of whether the benefit of alignment is limited to only the specific fallacy types that are compared, or if completing the comparison task will make participants more broadly sensitive to relational commonalities. To test this, participants only saw half the fallacy types during the pre-task. It is expected that same-fallacy comparison should increase same-fallacy response for those fallacy types seen during the pre-task phase; however, comparison may induce an overall relational mindset, promoting same fallacy response for the unseen items as well (Brown & Kane, 1988; Richland & Simms, 2015; Vendetti et al., 2014; but see Goldwater & Jamrozik, 2019).

The pre-task manipulation also asked participants to provide qualitative evaluations of the fallacies which could serve both as a manipulation check and lend additional insight into their familiarity with and reasoning around logical fallacies. For instance, if participants used a

specific name for two fallacies compared during the pre-task phase, it is likely they already have a clear representation of the fallacy and thus the comparison manipulation may have been less effective.

Method

Participants

145 Northwestern University undergraduates participated in this study for course credit. 15 failed to complete the task within the allotted time, and 27 failed more than half the catch trials. This left a total sample of 103.

Materials

The test triads were the same triads used in the baseline condition of Experiment 1. An additional set of paragraphs were created for the pre-task phase. These consisted of two additional instances of the fallacy types used as standards at test. Pre-task paragraphs shared no surface similarities with each other.

Procedure

In the pre-task phase participants were randomly assigned to either the comparison or sequential group. Participants saw a random set of 6 of the 12 fallacy types during to pre-task phase. This maintained both a manageable task duration and allowed for the seen/unseen contrast. Participants in the comparison group received the six same-fallacy pairs side-by-side. The sequential group saw the same 6 fallacy types, but instead of paired paragraphs, each

paragraph was presented one-by-one. Thus where the comparison group saw 6 pairs, the sequential group saw 12 individual paragraphs. Each same-fallacy response was separated by 5 other fallacies. The comparison group was asked to *“compare the two passages and using the space provided, describe key parallels between them.”* They were also given an example and a description of possible parallels. The sequential group was also given a task to ensure that even though they weren’t engaging in comparison they were processing the materials deeply. They were asked to *“read each passage carefully and using the space provided, describe the relevant parts of the argument, and evaluate why it is erroneous.”* The test-phase procedure was identical to the baseline condition of Experiment 1.

In addition to examining the contrast between sequential presentation and comparison, it is also important to establish the effect of these manipulations over the baseline task. However, due to experimenter error, a unique baseline condition was not included in the design of Experiment 2. Thus, the decision was made to compare test-phase results from Experiment 2 to the test-phase results of the baseline group in Experiment 1. Thus, the analyses Experiment 2 had a 3(Pre-task group: baseline, comparison, sequential) x 2(Task: categorization, similarity) between-subjects design. While this is sub-optimal, this decision has some methodological merit as this group performed an identical task but did not receive any pre-task manipulation. However, since the Experiment 1 results were collected at a different time of year from a different set of students, they do not represent a randomly assigned sample. Follow-up work is currently underway utilizing this full design with a new set of participants.

The overarching prediction is again that participants will respond in a similar way on both the categorization and similarity tasks to manipulations designed to promote or hamper alignment. Given the design described above, the specific predictions are as follows: A) participants that perform the comparison task will show higher same-fallacy responding than those that perform the sequential task for both categorization and similarity tasks, and B) that participants that perform the comparison task will show higher same-fallacy responding than the baseline group for both categorization and similarity tasks.

Analysis

Table 2: Experiment 2 model parameters

| Results | <i>Predictors</i> | Response | | |
|----------------------------|--|-----------------|--------------|--------------|
| | | <i>Log-Odds</i> | <i>CI</i> | <i>p</i> |
| Pre-task group and Task | Intercept | 0.25 | -0.89 – 1.39 | 0.668 |
| were included as fixed | Pre-task (Comparison) | 1.60 | 0.13 – 3.07 | 0.033 |
| effects in a logistic | Pre-task (Sequential) | 0.62 | -0.99 – 2.23 | 0.449 |
| regression model, as well | Task (Categorization) | 0.96 | -0.64 – 2.56 | 0.239 |
| as the interaction term. A | Comparison * Categorization | -0.50 | -2.64 – 1.64 | 0.645 |
| main effect of Pre-task | Sequential * Categorization | 0.19 | -2.09 – 2.47 | 0.868 |
| group was found, | Random Effects | | | |
| indicating participants | σ^2 | 3.29 | | |
| showed higher same- | τ_{00} Participants | 6.02 | | |
| fallacy responding in the | τ_{00} Fallacies | 0.36 | | |
| comparison group than | ICC | 0.66 | | |
| the baseline group | N Participants | 147 | | |
| collapsing across both | N Fallacies | 12 | | |
| | Observations | 1764 | | |
| | Marginal R ² / Conditional R ² | 0.051 / 0.677 | | |

task types. ($\beta = 1.6$, 95% CI[0.13 - 3.07], $p = 0.03$). A main effect of Task was not found ($\beta = 0.96$, 95% CI[-0.64 – 1.18], $p = 0.24$). The Task x Pre-task interaction was also not found for either the sequential ($\beta = 0.19$, 95% CI[-2.09 – 2.47], $p = 0.87$) or the comparison group ($\beta = -0.5$, 95% CI[-2.64 – 1.64], $p = 0.65$).

As in Experiment 1, post-hoc comparisons of the estimated marginal means were used to assess fine-grained predictions. Prediction A) was not supported: participants in the comparison group did not show higher same-fallacy response than the sequential group on the similarity task ($\beta = 0.98$, 95% CI[-0.87 - 2.83], $p = 0.43$) or the categorization task ($\beta = 0.29$, 95% CI[-1.51 - 2.08], $p = 0.93$). Prediction B) was also not supported: participants in the comparison group did not show higher same-fallacy response than those in the baseline group on either the similarity task ($\beta = 1.6$, 95% CI[1.61 – 3.63], $p = 0.08$), or the categorization task ($\beta = 1.1$, 95% CI[-1.17 – 3.37], $p = 0.59$). Considering these findings along with the overall effect for group

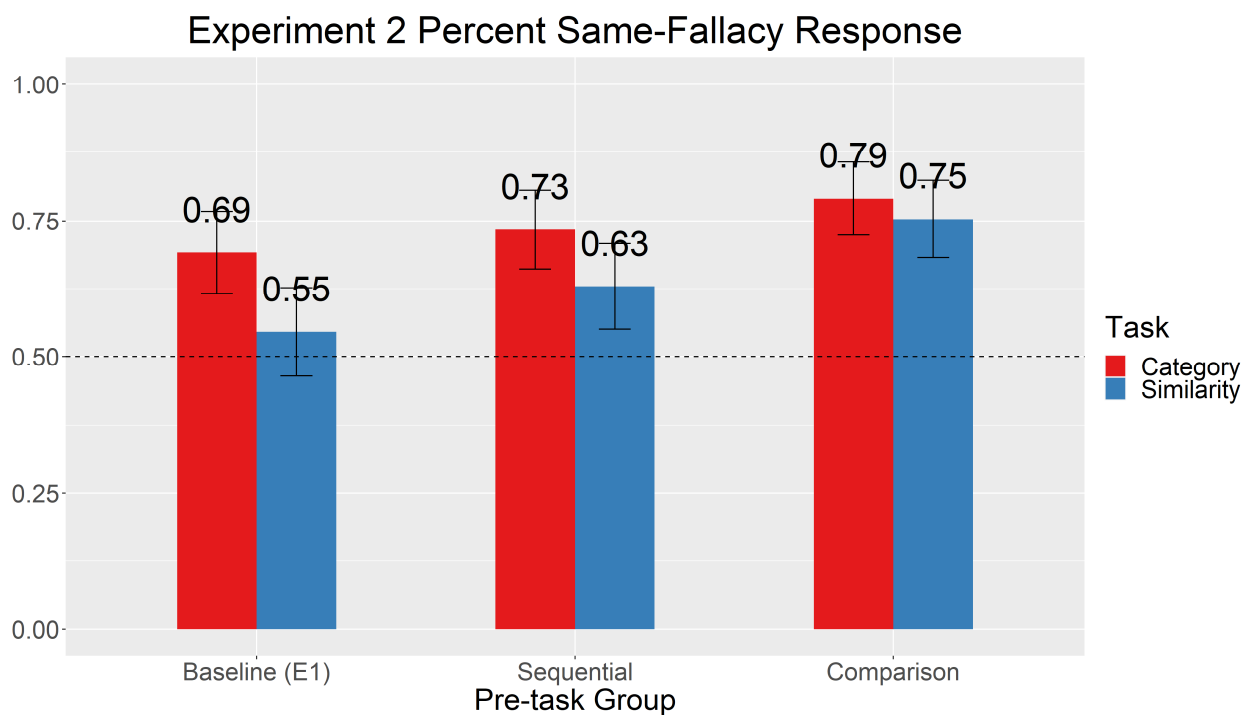


Figure 5: Experiment 2 test phase results

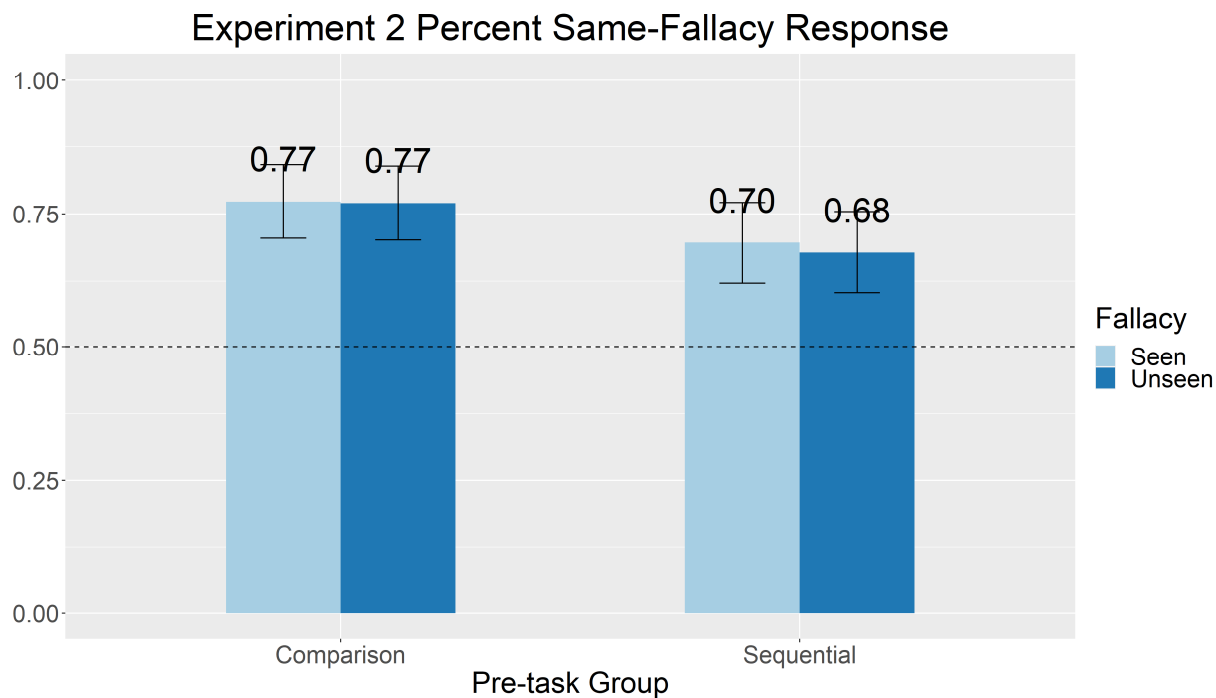


Figure 6: Experiment 2 seen vs. unseen results

suggests that while pre-task comparison had an overall positive effect on same-fallacy response, this effect was present only when collapsing across task type.

In order to address the open question of relational mindset, a separate model was used to predict same-fallacy response at test using group and fallacy presentation during the pre-task phase as fixed-effects. There was no evidence for a relational mindset: participants showed no difference on the items seen during the pre-task compared to those that were only seen at test ($\beta = 0.07$, 95% CI[-0.42 - 0.56], $p = 0.78$).

Qualitative Data

The qualitative data collected during the pre-task were also examined. The full qualitative dataset is awaiting coding by judges blinded to the conditions and hypotheses of the experiment. However, given the quantitative findings of Experiment2, it was deemed useful to take a cursory look at the patterns of response during the pre-task. Given the size of the data

and the amount of data to be analyzed without coders, the results from two fallacies were selected to be hand-coded by the experimenter: *argument from ignorance* and *circular reasoning*. This choice was made to select two fallacies that may span the spectrum of possible relational content in the fallacy. Circular reasoning involves a schematic relation between two premises that support each other, while an argument from authority relies on the relation between a claim and the authority's epistemic position to support the claim.

Each response was coded in one of several ways. Participants may have a) identified the shared relational content (in the comparison group) or responded to the reasoning of the fallacy in a way that showed they understood the problem with the fallacy (sequential group), b) specifically named the fallacy by a well-known name, c) focused on abstract aspects of the passage or passages unrelated to the logical structure, d) identified a problem with the reasoning *unrelated* to the fallacy, or e) failed to put in a good-faith effort to complete the task. Figure 9 provides examples of each of the codes mentioned, all taken from the circular reasoning fallacy. See Appendix B for the full set of coded qualitative data.

Two questions are especially relevant when it comes to the qualitative data: 1) did participants perform the task in the way they were directed? 2) is there any evidence that participants were already familiar with the fallacy type used in the experiment, and could this explain why the manipulation did not affect same-fallacy choice?

| |
|---|
| a) In both passages, one subject uses their argument as evidence rather than actual evidence. For example, Janene believes that warehouses are architecturally valuable because they add to the town's character given their architectural value (backing up their claim by restating their claim). Claudia believes that she must be cursed because such a bad week must be due to a curse (again, explaining her original claim by using her claim as evidence). |
| b) In the first passage, Claudia makes a logical error in claiming that the only possible reason for such a bad week is a curse. She fails to explain other options. In the second passage, Janene uses circular reasoning in an attempt to support her position on why the warehouses should be saved. She claims that they are architecturally valuable because they lend a distinctive character to the town, and the way they do that is by being architecturally valuable. |
| c) Both passages have a person believe something true, and there is a friend or other person in which they share this belief. Both have the person with the belief effect the other in a certain way of thinking, however it is not the same result |
| d) Claudia exhibits a superstitious belief in curses based on her recent experiences. Attributing a series of occurrences to something is also associated with good luck charms, something has been disproved before. Her reasoning that something supernatural must be to blame for all of her woes takes out any human fault on her part and on others. You can respond by demonstrating that these associations are only mental with examples of fake good luck charms and common misconceptions about ladders, umbrellas, mirrors, etc. |
| e) I see no similarities between the passages; one discusses people's opinions about a town, and the other discusses someone who thinks she has been cursed. No relation. |

Figure 7: Example responses to Experiment 2 pre-task

When it comes to question 1), it seems that participants showed fairly variable performance across comparison and sequential groups. The total of a) and b) codes, which indicate successful identification of the fallacy, was tallied up for each fallacy and each task. For *argument from authority* 19 responses showed identification of the fallacy in the comparison group, and 42 in the sequential group. This is to be expected since each participant in the sequential group produced two responses to each passage. For *circular reasoning* the results were a bit different. 15 responses identified the fallacy in the comparison group, but only 9 did in the sequential group. A closer look indicates that this is because many participants fixated on another aspect of one of the passages, which pertained to someone reasoning about a curse.

Participants almost all identified that the erroneous reasoning had to do with superstition regarding curses, and not the circular reasoning. These admittedly exploratory findings suggest there could be fair variability between fallacy types – especially in the sequential condition. When the two passages are not seen at the same time, participants may not home in on the core structure of the fallacy.

There are a few ways to approach question 2). First, using the coding system above, we can look at the number of participants that explicitly named the fallacy in their response. Though not a guarantee, this could indicate that participants were already familiar with the fallacy types before performing the task and so received little benefit from the manipulation. It is also possible that performing the comparison task made participants more likely to recall the name for the fallacy, leading to more explicit mentions in the comparison group than in the sequential group. For *argument from authority* comparison group responses explicitly mentioned the fallacy 3 times, and 4 times in the sequential group. For *circular reasoning* comparison group responses explicitly mentioned the fallacy 10 times, and 4 times in the sequential group. Again, it is hard to extrapolate from this small sample, but it seems there is some degree of by-fallacy variation in the ease of identifying the fallacy name. This may indeed be connected to how schematic certain fallacies are and how easily their structure comes to mind. In an attempt to gather more information over the whole set of qualitative data, a quick text search was conducted over the entire set of fallacies with particular key words that might yield names. These included “fallacy,” “appeal”, “argument”, and “logic.” Each returned response was read in-depth to determine if it did include a fallacy name. This returned 72 out of the 884 responses, representing 8.1% of the total responses. This is most certainly an

undercount, and a full analysis of the full qualitative dataset should be done to ascertain how often and under what circumstances participants were familiar with the fallacies.

Discussion

The chief prediction was that participants who compared during the pre-task would show higher same-fallacy responding on both similarity and categorization tasks than those that completed the sequential task. The finding that collapsing across task type, same-fallacy response was higher in the comparison group compared to the baseline group from Experiment 1 is consistent with Main Hypothesis 1. However, comparing across experiments is less than ideal. Further, while the comparison group showed higher same-fallacy responding than baseline, it was not higher than the sequential group. This is a key prediction, and the failure to find support for it – even for the similarity task – suggests that the pre-task requires additional examination. For example, it could be that asking for explanations during the sequential manipulation also increased same-fallacy responding. In addition, participants did not show higher same-fallacy responding on either the categorization task or the similarity task in either the sequential or baseline conditions. Overall, it appears that the comparison manipulation failed for Experiment 2.

There are several possible explanations for the failure of the comparison manipulation. First, comparing same-fallacy examples in the pre-task phase may not have brought out shared relational structure strongly enough to affect results during the test phase. Additional analysis of the qualitative data from the pre-task phase help clarify this. The cursory look at differences between the comparison and sequential group does not completely determine the success of the pre-task manipulation. It is also possible that despite seeing the items separately,

participants in the sequential group benefitted from the additional exposure to the fallacies during the pre-task phase. Finally, more than simply additional exposure, it is possible that the sequential group managed to compare same-fallacy examples from the pre-task phase into the test phase. This possibility seems unlikely given the length of the paragraphs and the number of cross-category examples between each pair – as well as the failure to find a difference between the seen and unseen fallacies. Qualitative data could be useful here as well. For instance, sequential group participants may have mentioned a fallacy they had seen previously, indicating comparison across items.

Differences in the wording of the instructions between the comparison and sequential pre-task conditions, may have also contributed to the failure to find a difference between comparison and sequential groups. In order to ensure participant engagement and to better match the tasks, participants in the sequential condition were asked to evaluate the reasoning in each fallacy and respond to the argument. However, this meant that in the sequential condition, and *not* the comparison condition, instructions specifically referred to the paragraphs as “errors in reasoning.” This may have focused sequential group participants on the reasoning of each passage, increasing relational response just enough to account for the current pattern of results. A follow-up experiment with modified sequential group instructions is currently underway to address this.

Additional analysis of the qualitative data may also lend insights into *how* comparison vs. sequential presentation affected participants. As mentioned previously, participants may have identified some fallacies by name, or specifically mentioned aspects of their relational structure. For instance, participants in the comparison condition may have listed alignable parts

of the argumentation more often compared to the sequential group. Participants may have also shown different patterns of response for fallacies that have more or less systematic relational structure. To the extent that comparison makes relational structure more apparent, the effect of comparison itself will be limited by this underlying relational structure

There was also no observed difference at test between items seen during the pre-task phase and those that were newly encountered. There are two possible explanations for this: either participants derived no benefit from the pre-task manipulation, or whatever benefit they did derive extended to both seen and unseen fallacy types. It is tempting to conclude that the pre-task manipulation induced a relational mindset. If this is the case, however, it must be true for both the comparison and sequential task.

Experiment 2 asked whether alignment through an explicit comparison task would increase same-fallacy response for both categorization and similarity. The results are inconclusive. There was no clear support for this hypothesis. However, additional work can help address some possible explanations for this pattern of results. Experiment 3 asks the reciprocal question: does disrupting alignment in turn *decrease* same-fallacy response for both categorization and similarity?

Experiment 3

Following the logic of the main hypotheses, if applying relational category knowledge involves alignment, making alignment more difficult should not only make it harder for people to detect relational similarity, but also less likely to correctly co-categorize relational category members. Experiment 3 tests this possibility by contrasting performance on the baseline task with a task

where alignment between the standard and the alternatives is made more difficult through temporal disruption.

Motivation for this manipulation once again comes from experiments showing that simultaneous juxtaposition of same-category examples reinforces relational similarity (Catrambone & Holyoak, 1989; Christie & Gentner, 2010; Kurtz et al., 2001; Loewenstein et al., 1999; Loewenstein & Gentner, 1998). For instance, Kurtz, Miao, and Gentner (2001) presented participants with analogous pictures of scenarios involving heat flow, and gave them tasks that varied in the amount of comparison involved. For instance, some participants interpreted the pictures separately while some viewed the pictures jointly and listed their correspondences. As an outcome measure, participants were given tasks that measured their insight into the relational structure shared between the pictures. Participants that viewed the images jointly *and* compared in order to list correspondences, showed a better understanding of the shared relational structure. Findings like this suggest that when participants cannot rely on alignment with co-present fallacies, they will focus more on surface features shared with the standard rather than relational structure. Consistent with these findings, participants should show lower same-fallacy responding on the similarity task. However, if alignment is also involved in categorization, this should also be the case for participants that receive the categorization task.

Method

Participants

111 Northwestern University undergraduates received course credit for their participation in this experiment. 2 participants failed to complete the task in the allotted time, and 13 failed more than half of the catch trials. This resulted in a total sample size of 97

Materials

All fallacies were the same as those used in the baseline set from Experiment 1.

Procedure

Participants were randomly assigned to the baseline condition or the disrupted condition. The baseline condition was identical to the one from Experiment 1. For each triad in the disrupted condition participants were given the standard and instructed to read it carefully since they would not be able to go back and read it again. When they indicated they were ready to move on the standard disappeared and the two alternatives appeared on the screen. Within each presentation group participants performed either the categorization or similarity tasks. Task instructions were identical to Experiment 1. Thus, Experiment 3 had a 2(Presentation: baseline, disrupted) x 2(Task: categorization, similarity) between-subjects design.

If presenting the standard and alternatives separately successfully disrupts alignment, participants in the disrupted condition should show lower same-fallacy response on both the similarity task and the categorization task than those in the baseline control condition. Finally, Experiment 3 also replicates the baseline task. It is expected that competing surface and

relational commonalities lead to higher same-fallacy responding in the categorization task than on the similarity task, as found in Experiment 1 and prior studies. Thus, the following predictions follow from this design: A) participants in the disrupted condition should show lower rates of same-fallacy responding on both similarity and categorization tasks, and B) in the baseline group, participants should show higher same-fallacy responding on the categorization task than the similarity task.

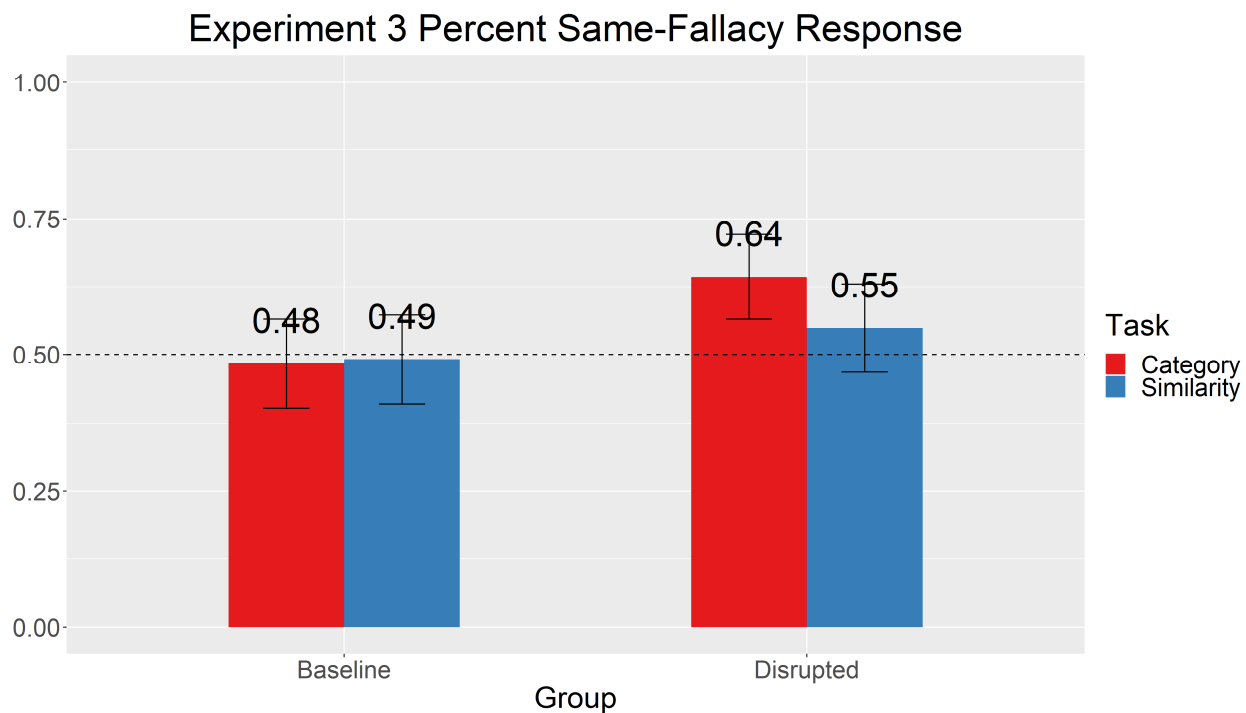


Figure 8: Experiment 3 results

Results

Table 3: Experiment 3 model parameters

| A full-factorial logistic regression model was fit predicting same-fallacy responding including Presentation and Task as fixed effects. A main effect not found for Presentation ($\beta = -0.76$, 95% CI[-2.47 – 0.96], $p = 0.39$) or Task ($\beta = 1.58$, 95% CI[-0.14 – 3.30], $p = 0.39$), nor was an effect found for the | Response | | | |
|--|------------------------------------|-----------------|--------------|----------|
| | <i>Predictors</i> | <i>Log-Odds</i> | <i>CI</i> | <i>p</i> |
| | Intercept | -0.39 | -1.69 – 0.90 | 0.553 |
| | Presentation (Disrupted) | -0.76 | -2.47 – 0.96 | 0.386 |
| | Task (Categorization) | 1.58 | -0.14 – 3.30 | 0.071 |
| | Disrupted * Categorization | 0.54 | -2.11 – 3.19 | 0.689 |
| | Random Effects | | | |
| | σ^2 | 3.29 | | |
| | τ_{00} Participants | 9.84 | | |
| | τ_{00} Fallacies | 0.50 | | |
| | ICC | 0.76 | | |
| | N Participants | 97 | | |
| | N Fallacies | 12 | | |
| | Observations | 1176 | | |
| | Marginal R^2 / Conditional R^2 | 0.063 / 0.774 | | |

Presentation x Task interaction ($\beta = 0.54$, 95% CI[-2.11 – 3.19], $p = 0.69$). This indicates that participants in the baseline and disrupted groups showed similar rates of same-fallacy responding.

Again, pairwise comparison of the estimated marginal means was used to assess specific hypotheses. Prediction A) was not supported as participants in the disrupted condition did not show lower same-fallacy response than those in the baseline control condition on either the similarity task ($\beta = 0.76$, 95% CI[-0.96 – 2.47], $p = 0.39$) or the categorization task ($\beta = 0.22$, 95% CI[-1.62 – 2.05], $p = 0.82$). Contrary to prediction B), participants in the baseline control condition did *not* show higher same-fallacy response on the categorization task than on the

similarity task ($\beta = 1.58$, 95% CI[-0.14 - 3.30], $p = 0.71$). However, participants in the disrupted condition, *did show* higher same-fallacy response on the categorization task than the similarity task ($\beta = 2.12$, 95% CI[0.26 - 3.98], $p = 0.03$). This was not predicted – though may reflect the failure of the disruption manipulation. Participants in the disrupted condition may be responding as expected in the baseline group. If so, this would make the lack of a Task effect in the baseline group more puzzling.

Discussion

The results of Experiment 3 did not reveal the expected difference between the baseline condition and the disrupted condition predicted by Main Hypothesis 1. This was true both for the similarity and categorization tasks. In principle, the problem could lie with the baseline task. Given the low level of same-fallacy responding in this condition, it would have been difficult to demonstrate a further decrease as a result of disrupted presentation. But this explanation is unlikely given the baseline results from Experiment 1 (0.69 for the categorization task and 0.55 for the similarity task). It is important to consider then, whether the disruption manipulation operated as expected. A core premise of Main Hypothesis 1 is that manipulations that affect alignment should demonstrate the same effect for the categorization task, as the similarity task where they are expected to have an effect. The finding that participants in the disrupted group did not show decreased same-fallacy responding on the similarity task suggests that the manipulation may have been ineffective.

Why did the disruption manipulation fail? One possibility is that placing the standard and the alternatives on separate pages did not sufficiently disrupt comparison. Participants'

representations of the standard may have still been sufficiently active that they were available for alignment with the alternatives despite the temporal disruption. A second possibility is that they were able to encode the standards well enough at their first encounter that this was not an issue. Given the complexity and length of the stimuli, this seems unlikely; however, it nonetheless remains a possibility. A third possibility is that participants may not have needed the co-present standard to identify the same-fallacy match. They may have already been familiar with the fallacies used in this experiment, and thus simply aligned the alternatives with stored category representations. This also seems relatively unlikely given the results of the previous experiments. A fourth possibility is that participants in the disrupted condition, realizing that they would not have another chance to read through the standard before making their choice, studied the standard more carefully than those in the baseline condition. That is, they may have engaged in additional self-explanation or deeper encoding than the baseline group. If so, then it is possible that the disruption worked contrary to expectations.

Additional work should focus on disentangling these possibilities. In order to ensure disruption, the difficulty of maintaining each passage in memory across the disruption could be increased. This could be achieved by including a filler task or increasing the time between the presentation of the standard and the alternatives. Participant familiarity with the fallacy types could be assessed by contrasting familiar relational categories with unfamiliar relational categories. This could be done by verifying participant classification of some types of fallacies before they are used at test, using relational categories that more familiar than logical fallacies, or specifically teaching participants the category structure before test. Of course, each of these approaches brings drawbacks that must be carefully considered in follow-up studies. Ultimately

there is still much work to be done investigating how alignment may be disrupted without also introducing additional memory or task demands.

In sum, the results of Experiment 3 are inconclusive. While the findings do not support either Main Hypotheses 1 or 2, it is unclear whether this results from a failure of the manipulation, or a more general failure of this line of prediction.

General Discussion

Let us begin by summarizing the results of all three experiments. Experiment 1 predicted that alignable alternatives would promote alignment between the alternatives, allowing participants to notice differences in relational structure with the standard. Contrary to predictions, this manipulation *did not* lead to the highest same-fallacy response in the alignable alternatives group compared to baseline and neutral groups for either task. Even though the alignable alternatives group showed higher same-fallacy response on the similarity task compared to the baseline group, it was not higher than the neutral group. Finally, there was also no evidence that participants showed higher same-fallacy response on the categorization task than on the similarity task when surface and relational similarity competed in the baseline task,

Experiment 2 predicted that comparing multiple same-fallacy alternatives before test would strengthen participants' representation of the relational structure shared by each fallacy. The results of this experiment are mixed. As predicted, when collapsing across task type the comparison group did show higher same-fallacy response in the comparison group compared to the baseline group from Experiment 1; however, the comparison group showed no higher same-fallacy response compared to the sequential group.

Finally, Experiment 3 predicted that making alignment more *difficult* should *decrease* same-fallacy choice compared to baseline. Again, this was *not* found for either the similarity task or the categorization task. As in Experiment 1, it was also predicted that the baseline group, when surface and relational similarity are in competition, show higher same-fallacy response on the categorization task than on the similarity task. However, this difference was *not* found in the baseline group.

Unfortunately, this pattern of results does not fully support either of the main hypotheses:

1. Both similarity *and* categorization should be affected by alignability: making structural alignment easier will increase relational responding for both similarity judgments and co-categorization tasks, while making alignment more difficult will decrease both.
2. When surface and relational similarity compete, it is expected that people will show higher same-fallacy response on categorization tasks compared to similarity tasks.

Both Experiment 1 and Experiment 2 found that manipulations that increased alignability increased same-fallacy response. However, additional results make it difficult to draw a clear conclusion based on these results. Even though the alignable alternatives group was higher than the baseline group in Experiment 1, it did not demonstrate any advantage over the neutral condition, which only removed the alluring surface match. On the one hand, this indicates that both categorization and similarity judgements are influenced by alluring surface matches; however, on the other hand it also indicates that alignability – at least of the two alternatives –

does not influence these judgments over and above removing the distractor. This conclusion is only confirmed by comparisons within each task group.

Clear evidence in favor of Main Hypothesis 2, which was meant to confirm previous findings, was also not found. Experiment 1 found a trending effect toward higher same-fallacy choice in the categorization task than the similarity task when surface similarity and relational similarity were in competition at baseline. This condition was replicated in Experiment 3; however, in this experiment there was no difference between the categorization and similarity groups.

What are some possible explanations for this pattern of results? One possibility is the question of sufficient power. While the practice of observed power calculation has drawn considerable criticism from statisticians (Hoenig & Heisey, 2001), it is nonetheless worth discussing whether the current design provided adequate power to test the claims. As described previously the by-subjects and by-items design of these experiments was approximated from previous work estimating Type I and Type II error for generalized linear mixed-effects models (GLMMs) (Luke, 2017). These estimates were used since it can be difficult to estimate power for GLMMs given the nested variance structure for these models. However, simulation methods from Ecology and Evolution are now being adapted to estimate power in other fields (Green & MacLeod, 2016; Johnson et al., 2015). Follow-up studies should adopt these methods and ensure that an adequate sample is collected at both the subject and item level. It may be, however, that these results are indeed adequately powered. If so, then it is important to consider the implications these data have on the categorization and similarity literature.

Conclusion

Many approaches to categorization and conceptual representation have made an intuitive appeal to similarity as a potential factor that undergirds categories. As reviewed in this paper, that has been a controversial claim, which empirical evidence has argued against. The research presented here took up a weaker version of this claim in an attempt to preserve the intuitive appeal of the similarity approach: that structural alignment as a cognitive process – and not similarity as a product – is involved in categorization. It was argued that relational categories, whose members share common relational structure rather than common features, could serve as an ideal test case for this claim. In the three experiments presented here, manipulations that have been shown to influence similarity processing were also expected to influence categorization. Unfortunately, the results largely did not reflect this. However, the results also did not find the expected divergence between categorization and similarity that has been found in the broader literature. So, what is the takeaway?

This work carries some methodological implications along with it. Central to the logic of this thesis was the expectation that manipulations that have been shown to affect certain cognitive faculties in one domain, will have the same and predictable effect in a new domain. All experiments reported here attempted to export a methodology that was well-studied to the study of logical fallacies as relational fallacies. The failure of these manipulations serves as a reminder that though we think of cognitive faculties as core and abstract components of the mind, they are nonetheless subject to the demands of current context. Too often it is expected that for every person and in every situation, cognition functions in regularized and measurable

ways. The failure of these manipulations by no means indicate that the project of quantifying higher-order cognition is hopeless; rather, we as cognitive scientists must be judicious and methodical in how we extend our theories in different domains.

Of course, another possibility is that the core premise of this thesis fails. While it is tempting to conclude that structural alignment plays no role in categorization, this would be an overinterpretation of these results. If, however, this is indeed the case, it will be key for future research on this topic to consider the intuitions that give rise to this claim. The motivation behind this work, and arguably much of the previous work in similarity in categorization, is the sense that as people move through the world their latent curiosity and cognitive capacity leads them constantly to search for meaning. A long history of research in analogy and similarity suggests that part of the human condition is this constant thirst for common structure. It remains to be seen whether this also applies to the domain of categorization.

References

- Ahn, W., Kim, N. S., Lassaline, M. E., & Dennis, M. J. (2000). Causal status as a determinant of feature centrality. *Cognitive Psychology*, *41*(4), 361–416.
<https://doi.org/10.1006/cogp.2000.0741>
- Barr, D. J. (2013). Random effects structure for testing interactions in linear mixed-effects models. *Frontiers in Psychology*, *4*. <https://doi.org/10.3389/fpsyg.2013.00328>
- Barr, D. J., Levy, R., Scheepers, C., & Tily, H. J. (2013). Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language*, *68*(3), 255–278.
<https://doi.org/10.1016/j.jml.2012.11.001>
- Barsalou, L. W. (1982). Context-independent and context-dependent information in concepts. *Memory & Cognition*, *10*(1), 82–93. <https://doi.org/10.3758/BF03197629>
- Barsalou, L. W. (1983). Ad hoc categories. *Memory & Cognition*, *11*(3), 211–227.
<https://doi.org/10.3758/BF03196968>
- Brooks, L. R. (1978). Nonanalytic concept formation and memory for instances. In E. H. Rosch & B. Lloyd (Eds.), *Cognition and Categorization* (pp. 3–170). Lawrence Erlbaum Associates.
- Brown, A. L., & Kane, M. J. (1988). Preschool children can learn to transfer: Learning to learn and learning from example. *Cognitive Psychology*, *20*(4), 493–523.
[https://doi.org/10.1016/0010-0285\(88\)90014-X](https://doi.org/10.1016/0010-0285(88)90014-X)
- Carey, S. (1985). *Conceptual change in childhood*. MIT Press.

Catrambone, R., & Holyoak, K. J. (1989). Overcoming contextual limitations on problem-solving transfer. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *15*(6), 1147–1156. <https://doi.org/10.1037/0278-7393.15.6.1147>

Christie, S., & Gentner, D. (2010). Where hypotheses come from: Learning new relations by structural alignment. *Journal of Cognition and Development*, *11*(3), 356–373.

<https://doi.org/10.1080/15248371003700015>

Christie, S., Gentner, D., Call, J., & Haun, D. B. M. (2016). Sensitivity to relational similarity and object similarity in apes and children. *Current Biology*, *26*(4), 531–535.

<https://doi.org/10.1016/j.cub.2015.12.054>

Clement, C. A., & Gentner, D. (1991). Systematicity as a selection constraint in analogical mapping. *Cognitive Science*, *15*(1), 89–132. https://doi.org/10.1207/s15516709cog1501_3

Coombs, C. H. (1952). *A theory of psychological scaling*. The University of Michigan Press.

<https://doi.org/10.1037/14785-000>

Davidson, N. S., & Gelman, S. A. (1990). Inductions from novel categories: The role of language and conceptual structure. *Cognitive Development*, *5*(2), 151–176.

[https://doi.org/10.1016/0885-2014\(90\)90024-N](https://doi.org/10.1016/0885-2014(90)90024-N)

Estes, W. K. (1994). *Classification and Cognition*. Oxford University Press.

Falkenhainer, B., Forbus, K. D., & Gentner, D. (1989). The structure-mapping engine: Algorithm and examples. *Artificial Intelligence*, *41*(1), 1–63. [https://doi.org/10.1016/0004-3702\(89\)90077-](https://doi.org/10.1016/0004-3702(89)90077-5)

Forbus, K. D., Ferguson, R. W., Lovett, A., & Gentner, D. (2017). Extending SME to handle large-scale cognitive modeling. *Cognitive Science*, *41*(5), 1152–1201.

<https://doi.org/10.1111/cogs.12377>

Gelman, S. A. (1989). Children's use of categories to guide biological inferences. *Human Development*, *32*(2), 65–71. <https://doi.org/10.1159/000276364>

Gelman, S. A., & Markman, E. M. (1986). Categories and induction in young children. *Cognition*, *23*(3), 183–209. [https://doi.org/10.1016/0010-0277\(86\)90034-X](https://doi.org/10.1016/0010-0277(86)90034-X)

Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. *Cognitive Science*, *7*(2), 155–170. [https://doi.org/10.1016/S0364-0213\(83\)80009-3](https://doi.org/10.1016/S0364-0213(83)80009-3)

Gentner, D. (1988). Metaphor as structure mapping: The relational shift. *Child Development*, *59*(1), 47–59. <https://doi.org/10.2307/1130388>

Gentner, D. (2003). Why we're so smart. In D. Gentner & S. Goldin-Meadow (Eds.), *Language in mind: Advances in the study of language and thought* (pp. 195–235). MIT Press.

Gentner, D. (2005). The development of relational category knowledge. In L. Gershkoff-Stowe & D. H. Rakison (Eds.), *Building object categories in developmental time* (pp. 245–275). Psychology Press.

Gentner, D. (2010). Bootstrapping the mind: Analogical processes and symbol systems. *Cognitive Science*, *34*(5), 752–775. <https://doi.org/10.1111/j.1551-6709.2010.01114.x>

Gentner, D., & Kurtz, K. J. (2005). Relational categories. In W. Ahn, R. L. Goldstone, B. C. Love, A. B. Markman, & P. Wolff (Eds.), *Categorization inside and outside the laboratory: Essays in honor of Douglas L. Medin* (pp. 151–175). American Psychological Association.

- Gentner, D., & Markman, A. B. (1997a). Structure mapping in analogy and similarity. *American Psychologist*, 52(1), 45–56. <https://doi.org/10.1037/0003-066X.52.1.45>
- Gentner, D., & Markman, A. B. (1997b). Structure mapping in analogy and similarity. *American Psychologist*, 52(1), 45–56. <https://doi.org/10.1037/0003-066X.52.1.45>
- Gentner, D., & Medina, J. (1998). Similarity and the development of rules. *Cognition*, 65(2–3), 263–297. [https://doi.org/10.1016/S0010-0277\(98\)00002-X](https://doi.org/10.1016/S0010-0277(98)00002-X)
- Gentner, D., & Namy, L. L. (1999). Comparison in the development of categories. *Cognitive Development*, 14(4), 487–513. [https://doi.org/10.1016/S0885-2014\(99\)00016-7](https://doi.org/10.1016/S0885-2014(99)00016-7)
- Gentner, D., & Rattermann, M. J. (1991). Language and the career of similarity. In S. A. Gelman & J. P. Byrnes (Eds.), *Perspectives on Language and Thought: Interrelations in Development* (pp. 225–277). Cambridge University Press.
- Gentner, D., Rattermann, M. J., & Forbus, K. D. (1993). The roles of similarity in transfer: Separating retrievability from inferential soundness. *Cognitive Psychology*, 25(4), 524–575. <https://doi.org/10.1006/cogp.1993.1013>
- Gentner, D., & Toupin, C. (1986). Systematicity and surface similarity in the development of analogy. *Cognitive Science*, 10(3), 277–300. https://doi.org/10.1207/s15516709cog1003_2
- Goldstone, R. L. (1994). The role of similarity in categorization: Providing a groundwork. *Cognition*, 52(2), 125–157. [https://doi.org/10.1016/0010-0277\(94\)90065-5](https://doi.org/10.1016/0010-0277(94)90065-5)
- Goldstone, R. L., Medin, D. L., & Gentner, D. (1991). Relational similarity and the nonindependence of features in similarity judgments. *Cognitive Psychology*, 23(2), 222–262. [https://doi.org/10.1016/0010-0285\(91\)90010-L](https://doi.org/10.1016/0010-0285(91)90010-L)

Goldwater, M. B., Bainbridge, R., & Murphy, G. L. (2016). Learning of role-governed and thematic categories. *Acta Psychologica, 164*, 112–126.

<https://doi.org/10.1016/j.actpsy.2015.10.011>

Goldwater, M. B., & Jamrozik, A. (2019). Can a relational mindset boost analogical retrieval?

Cognitive Research: Principles and Implications, 4(1), 47. <https://doi.org/10.1186/s41235-019-0198-8>

Goldwater, M. B., Markman, A. B., & Stilwell, C. H. (2011). The empirical case for role-governed categories. *Cognition, 118*(3), 359–376. <https://doi.org/10.1016/j.cognition.2010.10.009>

Goodman, N. (1972). Seven strictures on similarity. In *Problems and projects*. Bobs-Merril.

Green, P., & MacLeod, C. J. (2016). SIMR: An R package for power analysis of generalized linear mixed models by simulation. *Methods in Ecology and Evolution, 7*(4), 493–498.

<https://doi.org/10.1111/2041-210X.12504>

Hampton, J. A. (1998). Similarity-based categorization and fuzziness of natural categories.

Cognition, 65(2), 137–165. [https://doi.org/10.1016/S0010-0277\(97\)00042-5](https://doi.org/10.1016/S0010-0277(97)00042-5)

Hampton, J. A. (2001). The role of similarity in natural categorization. In M. Ramscar & U. Hahn (Eds.), *Similarity and categorization* (pp. 13–28). Oxford University Press.

<https://doi.org/10.1093/acprof:oso/9780198506287.003.0002>

Hampton, J. A. (2007). Typicality, graded membership, and vagueness. *Cognitive Science, 31*(3), 355–384. <https://doi.org/10.1080/15326900701326402>

Hampton, J. A., Estes, Z., & Simmons, S. (2007). Metamorphosis: Essence, appearance, and behavior in the categorization of natural kinds. *Memory & Cognition, 35*(7), 1785–1800.

<https://doi.org/10.3758/BF03193510>

Hoenig, J. M., & Heisey, D. M. (2001). The abuse of power. *The American Statistician*, *55*(1), 19–24. <https://doi.org/10.1198/000313001300339897>

Holyoak, K. J., & Thagard, P. (1989). Analogical mapping by constraint satisfaction. *Cognitive Science*, *13*(3), 295–355. https://doi.org/10.1207/s15516709cog1303_1

Hummel, J. E., & Holyoak, K. J. (1997). Distributed representations of structure: A theory of analogical access and mapping. *Psychological Review*, *104*(3), 427–466.

<https://doi.org/10.1037/0033-295X.104.3.427>

Johnson, P. C. D., Barry, S. J. E., Ferguson, H. M., & Müller, P. (2015). Power analysis for generalized linear mixed models in ecology and evolution. *Methods in Ecology and Evolution*, *6*(2), 133–142. <https://doi.org/10.1111/2041-210X.12306>

Jones, M., & Love, B. C. (2007). Beyond common features: The role of roles in determining similarity. *Cognitive Psychology*, *55*(3), 196–231.

<https://doi.org/10.1016/j.cogpsych.2006.09.004>

Kalish, C. W. (1995). Essentialism and graded membership in animal and artifact categories. *Memory & Cognition*, *23*(3), 335–353. <https://doi.org/10.3758/BF03197235>

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

Keil, F. C. (1995). The growth of causal understandings of natural kinds. In *Causal cognition: A multidisciplinary debate* (pp. 234–267). Clarendon Press/Oxford University Press.

Keil, F. C., & Batterman, N. (1984). A characteristic-to-defining shift in the development of word meaning. *Journal of Verbal Learning and Verbal Behavior*, *23*(2), 221–236.

[https://doi.org/10.1016/S0022-5371\(84\)90148-8](https://doi.org/10.1016/S0022-5371(84)90148-8)

Kemler-Nelson, D. G. (1989). The nature and occurrence of holistic processing. In *Object perception: Structure and process* (pp. 357–386). Lawrence Erlbaum Associates, Inc.

Kotovsky, L., & Gentner, D. (1996). Comparison and categorization in the development of relational similarity. *Child Development*, 67(6), 2797–2822. <https://doi.org/10.1111/j.1467-8624.1996.tb01889.x>

Krawczyk, D. C., Holyoak, K. J., & Hummel, J. E. (2005). The one-to-one constraint in analogical mapping and inference. *Cognitive Science*, 29(5), 797–806.

https://doi.org/10.1207/s15516709cog0000_27

Kurtz, K. J. (2015). Human category learning: Toward a broader explanatory account. In B. H. Ross (Ed.), *Psychology of Learning and Motivation* (Vol. 63, pp. 77–114). Academic Press.

<https://doi.org/10.1016/bs.plm.2015.03.001>

Kurtz, K. J., Boukrina, O., & Gentner, D. (2013). Comparison promotes learning and transfer of relational categories. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 39(4), 1303–1310.

Kurtz, K. J., & Gentner, D. (2001). Kinds of kinds: Sources of category coherence. In *Proceedings of the 23rd Annual Conference of the Cognitive Science Society*, 522–527.

Kurtz, K. J., Miao, C.-H., & Gentner, D. (2001). Learning by analogical bootstrapping. *Journal of the Learning Sciences*, 10(4), 417–446. https://doi.org/10.1207/S15327809JLS1004new_2

Larkey, L. B., & Love, B. C. (2003). CAB: Connectionist analogy builder. *Cognitive Science*, 27(5), 781–794. [https://doi.org/10.1016/S0364-0213\(03\)00066-1](https://doi.org/10.1016/S0364-0213(03)00066-1)

- Lassaline, M. E., & Murphy, G. L. (1998). Alignment and category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 24(1), 144–160.
<https://doi.org/10.1037/0278-7393.24.1.144>
- Lenth, R., Singmann, H., Love, J., Buerkner, P., & Herve, M. (2020). *emmeans: Estimated Marginal Means, aka Least-Squares Means* (1.4.8) [Computer software]. <https://CRAN.R-project.org/package=emmeans>
- Loewenstein, J., & Gentner, D. (1998). Relational language facilitates analogy in children. *Proceedings of the Twentieth Annual Conference of the Cognitive Science Society*, 20, 615–620.
- Loewenstein, J., Thompson, L., & Gentner, D. (1999). Analogical encoding facilitates knowledge transfer in negotiation. *Psychonomic Bulletin & Review*, 6(4), 586–597.
<https://doi.org/10.3758/BF03212967>
- Luke, S. G. (2017). Evaluating significance in linear mixed-effects models in R. *Behavior Research Methods*, 49(4), 1494–1502. <https://doi.org/10.3758/s13428-016-0809-y>
- Markman, A. B. (1997). Constraints on analogical inference. *Cognitive Science*, 21(4), 373–418.
https://doi.org/10.1207/s15516709cog2104_1
- Markman, A. B., & Gentner, D. (1993a). Splitting the differences: A structural alignment view of similarity. *Journal of Memory and Language*, 32(4), 517–535.
<https://doi.org/10.1006/jmla.1993.1027>
- Markman, A. B., & Gentner, D. (1993b). Structural alignment during similarity comparisons. *Cognitive Psychology*, 25(4), 431–467. <https://doi.org/10.1006/cogp.1993.1011>
- Markman, A. B., & Stilwell, C. H. (2001). Role-governed categories. *Journal of Experimental & Theoretical Artificial Intelligence*, 13(4), 329–358. <https://doi.org/10.1080/09528130110100252>

- Medin, D. L., Goldstone, R. L., & Gentner, D. (1993). Respects for similarity. *Psychological Review*, 100(2), 254–278. <https://doi.org/10.1037/0033-295X.100.2.254>
- Medin, D. L., & Schaffer, M. (1978). Context theory of classification learning. *Psychological Review*, 85(3), 207–238.
- Medin, D. L., & Shoben, E. J. (1988). Context and structure in conceptual combination. *Cognitive Psychology*, 20(2), 158–190. [https://doi.org/10.1016/0010-0285\(88\)90018-7](https://doi.org/10.1016/0010-0285(88)90018-7)
- Murphy, G. L. (2002). *The Big Book of Concepts (Bradford Books)*. The MIT Press.
<http://www.amazon.ca/exec/obidos/redirect?tag=citeulike09-20&path=ASIN/0262632993>
- Murphy, G. L., & Medin, D. L. (1985). The role of theories in conceptual coherence. *Psychological Review*, 92(3), 289–316. <https://doi.org/10.1037/0033-295X.92.3.289>
- Namy, L. L., & Gentner, D. (2002). Making a silk purse out of two sow's ears: Young children's use of comparison in category learning. *Journal of Experimental Psychology: General*, 131(1), 5–15. <https://doi.org/10.1037/0096-3445.131.1.5>
- Paik, J. H., & Mix, K. S. (2006). Preschoolers' use of surface similarity in object comparisons: Taking context into account. *Journal of Experimental Child Psychology*, 95(3), 194–214.
<https://doi.org/10.1016/j.jecp.2006.06.002>
- Posner, M. I., & Keele, S. W. (1968). On the genesis of abstract ideas. *Journal of Experimental Psychology*, 77(3, Pt.1), 353–363. <https://doi.org/10.1037/h0025953>
- Richland, L. E., Holyoak, K. J., & Stigler, J. W. (2004). Analogy use in eighth-grade mathematics classrooms. *Cognition and Instruction*, 22(1), 37–60.
https://doi.org/10.1207/s1532690Xci2201_2

- Richland, L. E., & Simms, N. (2015). Analogy, higher order thinking, and education. *WIREs Cognitive Science*, 6(2), 177–192. <https://doi.org/10.1002/wcs.1336>
- Rips, L. (1989). Similarity, typicality, and categorization. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning* (pp. 21–59). Cambridge University Press.
- Rips, L., & Collins, A. (1993). Categories and resemblance. *Journal of Experimental Psychology: General*, 122(4), 468–486. <https://doi.org/10.1037//0096-3445.122.4.468>
- Rosch, E. H. (1973). Natural categories. *Cognitive Psychology*, 4(3), 328–350. [https://doi.org/10.1016/0010-0285\(73\)90017-0](https://doi.org/10.1016/0010-0285(73)90017-0)
- Rosch, E. H. (1975). Cognitive representations of semantic categories. *Journal of Experimental Psychology: General*.
- Rosch, E. H., & Mervis, C. B. (1975). Family resemblances: Studies in the internal structure of categories. *Cognitive Psychology*, 7(4), 573–605. [https://doi.org/10.1016/0010-0285\(75\)90024-9](https://doi.org/10.1016/0010-0285(75)90024-9)
- Roth, E. M., & Shoben, E. J. (1983). The effect of context on the structure of categories. *Cognitive Psychology*, 15(3), 346–378. [https://doi.org/10.1016/0010-0285\(83\)90012-9](https://doi.org/10.1016/0010-0285(83)90012-9)
- Sagi, E., Gentner, D., & Lovett, A. (2012). What difference reveals about similarity. *Cognitive Science*, 36(6), 1019–1050. <https://doi.org/10.1111/j.1551-6709.2012.01250.x>
- Shao, R., & Gentner, D. (2019). Symmetry: Low-level visual feature or abstract relation? *Proceedings of the 41st Annual Meeting of the Cognitive Science Society*, 2790–2796.
- Shepard, R. N. (1974). Representation of structure in similarity data: Problems and prospects. *Psychometrika*, 39(4), 373–421. <https://doi.org/10.1007/BF02291665>

- Simms, N. K., Frausel, R. R., & Richland, L. E. (2018). Working memory predicts children's analogical reasoning. *Journal of Experimental Child Psychology, 166*, 160–177.
<https://doi.org/10.1016/j.jecp.2017.08.005>
- Smith, E. E., & Medin, D. L. (1981). *Categories and concepts* (Vol. 9). Harvard University Press.
- Smith, E. E., & Osherson, D. N. (1984). Conceptual combination with prototype concepts. *Cognitive Science, 8*(4), 337–361. [https://doi.org/10.1016/S0364-0213\(84\)80006-3](https://doi.org/10.1016/S0364-0213(84)80006-3)
- Smith, L. B. (1989). A model of perceptual classification in children and adults. *Psychological Review, 96*(1), 125–144.
- Spellman, B. A., & Holyoak, K. J. (1992). If Saddam is Hitler then who is George Bush? Analogical mapping between systems of social roles. *Journal of Personality and Social Psychology, 62*(6), 913–933. <https://doi.org/10.1037/0022-3514.62.6.913>
- Spellman, B. A., & Holyoak, K. J. (1996). Pragmatics in analogical mapping. *Cognitive Psychology, 31*(3), 307–346. <https://doi.org/10.1006/cogp.1996.0019>
- Thibaut, J.-P., French, R., & Vezneva, M. (2010). The development of analogy making in children: Cognitive load and executive functions. *Journal of Experimental Child Psychology, 106*(1), 1–19. <https://doi.org/10.1016/j.jecp.2010.01.001>
- Tversky, A. (1977). Features of similarity. *Psychological Review, 84*, 327–352.
- Tversky, A., & Gati, I. (1978). Studies of Similarity. In E. H. Rosch & B. Lloyd (Eds.), *Cognition and Categorization* (pp. 1–1978). Lawrence Elbaum Associates.
- Vendetti, M. S., Wu, A., & Holyoak, K. J. (2014). Far-out thinking: Generating solutions to distant analogies promotes relational thinking. *Psychological Science, 25*(4), 928–933.
<https://doi.org/10.1177/0956797613518079>

Wittgenstein, L. (1953). *Philosophical investigations. Philosophische Untersuchungen.*

Macmillan.

Appendix A

The following tables provide the stimuli used in Experiment 1. Of note, each standard, surface-similar distractor, and same-fallacy match was used in the baseline triad of each subsequent experiment. Table 1 provides the fallacies used as the standard in the triad task, and the same-fallacy match, including the type of fallacy for each. Table 2 provides the alternatives participants received, with the alternative fallacy type listed. Participants in Experiment 1 always received the standard and same-fallacy alternative and would receive either the distractor, alignable, or neutral alternative as part of the triad task. Some fallacy types used as standards are also used for alternatives, though these were never paired with the same standard in the triad task.

| Standard Fallacy Type | Standard | Same-fallacy Alternative |
|------------------------------|--|--|
| Accident (false cause) | Stulternious was a dedicated, but not especially talented philosopher. One of the topics that interested him most was the operation of heavenly bodies. Stulternious noticed that every time he went to sleep the Sun would go down and he would see it rise the next day when he awoke. From this he concluded in amazement that his going to sleep caused the Sun to set. | The USDA announced in a press conference recently that it is limiting the amount of margarine that can be sold in supermarkets. A recent study shows that per capita margarine consumption is down in Maine. Divorce rates have also fallen in equal numbers. "The incredible damage that margarine has been causing to the American family is simply unacceptable," the USDA stated. |
| Appeal to ignorance | Gina and Miguel are two astronomy students. Recently they have become better friends after realizing they both believe that it's entirely possible UFOs could have visited earth. One day at a department meeting Professor Lin overhears them discussing this possibility. "I can't believe what I'm hearing!" Lin exclaimed, "Not a single report of UFO sightings has ever been authenticated, therefore they couldn't possibly exist!" | General Warren is testifying in front of congress. The country has been at war for months on end, and this has led to unrest at home. The subject of the committee hearing is possible subversive activity on the homefront. "I believe now, more than ever," testified General Warren, "that we are at no risk of an attack. We are unaware of any ongoing conspiracy from known subversives we have identified. This means there is no chance we could be in danger." |
| Circular reasoning | During Bill's job interview, the interviewer asked him if he had any additional references. Bill said he had a very positive one from his previous employer, Don. The interviewer responded that that was nice, but wanted assurances that Don was indeed a reputable businessman. Bill responded that Don indeed was a reputable businessman. In fact, he could provide him an excellent reference if the interviewer would like. | The head surgeon and the new doctor, Tom, were preparing for the complicated procedure. The head surgeon asked Tom about his previous training, and whether he understood the procedure. Tom mentioned that his training with Dr. Minsky provided him with valuable experience. The head surgeon replied that he hadn't heard of Dr. Minsky, and asked for his credentials. Tom told him that he, Tom, had the greatest respect for Dr. Minsky, so there was no cause for concern. |

| | | |
|-----------------------|---|---|
| Ignoring common cause | Jerry was an educational researcher who was trying to understand what helps children become successful students. For one study in particular, he gathered information about several variables. He found it most interesting that children who did the best in school also were the ones who went to the zoo most often. From this Jerry concluded that there must be something about zoos that makes children become better students. | Mark noticed that the tree in his front lawn was dying. A few weeks later he saw that the flowers he planted near the tree were also dying, causing great concern. He noticed that many of the leaves of the dead tree were piled up around the flowers. From this, he concluded that the leaves from the tree must have killed the flowers. He felt this was most likely, despite news reports that a leak from a local chemical plant had contaminated some of the town's irrigation water. |
| Loaded question | Todd was on trial, accused of the first-degree murder of his wife. While on the stand he was questioned harshly by the prosecuting attorney. At the end of a long barrage of questions, the lawyer asked him: on the night of his wife's murder, how long it had been since Todd had stopped verbally abusing his wife? | Tim was a high school student, who had just completed his math test. He stayed after class gathering his books and noticed the teacher staring at him. Tim was perplexed, wondering what she wanted. After a few minutes, the teacher approached Tim, and asked him what he did to hide the fact he cheated on the test. |
| No true scotsman | Todd was on trial, accused of the first-degree murder of his wife. While on the stand he was questioned harshly by the prosecuting attorney. At the end of a long barrage of questions, the lawyer asked him: on the night of his wife's murder, how long it had been since Todd had stopped verbally abusing his wife? | Ben was a physics student. He believed that all scientists must be atheists like him. He shared this with his adviser, Dr. Rodriguez, who shared that she's a practicing Muslim. He was visibly shocked. She pointed out that his position was illogical, but he was unswayed. "Dr. Rodriguez must not really be a true scientist," Ben thought to himself as he left her office. |
| Appeal to authority | Erin's mother is convinced there are ghosts haunting their home, causing noises in the middle of the night. Erin has tried telling her that there's likely a more rational explanation than spirit energy. Her mother, however, maintains that it must be ghosts because her upstairs neighbor Madame Blavatsky has examined the house, and she must know. | Ben was a physics student. He believed that all scientists must be atheists like him. He shared this with his adviser, Dr. Rodriguez, who shared that she's a practicing Muslim. He was visibly shocked. She pointed out that his position was illogical, but he was unswayed. "Dr. Rodriguez must not really be a true scientist," Ben thought to himself as he left her office. |

| | | |
|-------------------|--|--|
| Is/ought | Larry and Tina are college sophomores who live on the same hall. One night near the end of the term, they discuss how they are preparing for their chemistry final. Larry is appalled to learn Tina is planning to cheat on the exam. "I don't see what the big deal is," says Tina, "people with ambition naturally try to get ahead. I'm an ambitious person, so there's nothing wrong with it." | Robert and Liam are engaged in a heated debate about which end should be down when storing eggs, the wide one or the narrow one. Robert argues that the wide end is more stable and therefore better to store the egg on. Liam, however, has heard that the King stores his eggs on the narrow end, and after all, he's the King. |
| Strawman argument | Teddy is an economics student. In his last macroeconomics class Professor Klint asked students for suggestions on how to combat inequality. Teddy proposed a graduated income tax that was higher for the rich. Professor Klint was surprised by Teddy's idea, wondering why Teddy didn't like rich people. Otherwise, why else would he think the rich should pay more taxes than the poor? | During Thanksgiving dinner, Thom gets into a political argument with his uncle Stephen. Thom thinks the state should eliminate capital punishment since he thinks it's morally unjust to kill someone -- even if they've committed murder. Stephen thinks this is ridiculous. Every well-run society throughout history has had ways to maintain order. A crime as serious as murder can destroy the social order, so capital punishment must be just. |
| False dichotomy | Inez is home alone on a Saturday night. She had previously planned to meet up with friends, but the combination of bad weather and poor communication caused everyone to cancel. As she flipped through the channels on the TV. She Concluded there were only two possibilities: watch boring TV or drink whiskey until she fell asleep. | Inez is home alone on a Saturday night. She had previously planned to meet up with friends, but the combination of bad weather and poor communication caused everyone to cancel. As she flipped through the channels on the TV. She Concluded there were only two possibilities: watch boring TV or drink whiskey until she fell asleep. |
| Tu quoque | Karen is out for brunch with her aunt Gina. Karen ordered a breakfast bacon cheeseburger, with extra bacon and a fried egg on top. Gina mentions to Karen that maybe she shouldn't eat such fatty food, since recent science indicates high-fat diets can cause cancer. Karen responds that Gina eats fatty food all the time, so that clearly can't be true. | Ethan and Carl are grabbing a happy hour drink. After a drink or two Carl reveals that for the last couple weeks he's been cheating on his boyfriend Jacob with someone from work. Ethan chastises Carl saying that infidelity is just morally wrong. Carl is unrepentant. He reminds Ethan that his marriage ended after he himself cheated on his wife; Ethan's own unfaithfulness invalidates his argument. |

| | | |
|------------------------|--|--|
| Counterfactual fallacy | Anisha is distraught. After five years of being together she and her long-term partner Zack decided to break up. Anisha's mother is less than sympathetic. "Ani, you know this your fault dear," she says to her one day, "I hate to say it, but if you had cooked him dinner every night, he wouldn't have left you." | The major looked around at the stunned control room. Everyone in the room looked on in silence as the monitors showed the satellite disintegrating in the Martian atmosphere. Suddenly, he was furious. He was certain the lack of adequate personnel was responsible. With double the staff, this disaster would never have happened. |
|------------------------|--|--|

| Alternative Fallacy Type | Surface-similar to Standard | Alignable Alternative | Neutral Alternative |
|---------------------------------|--|--|---|
| Argument from authority | Tuccius was a dedicated, but not especially talented philosopher. One of the topics that interested him most was the operation of heavenly bodies. His tutor, Agnemor, told him that after careful study he had concluded that going to sleep caused the Sun to set. Because Agnemor was clearly a wiser and more practiced philosopher, Tuccius also concluded that going to sleep caused the sun to set. | The USDA announced in a press conference recently that it is limiting the amount of margarine that can be sold in supermarkets. Apparently, the Vice President believes that excess consumption of margarine was responsible for his messy public divorce. "The incredible damage that margarine has been causing to the American family is simply unacceptable," the USDA stated. | Marzanna and Dominique are discussing the merits of different diets. Dominique has recently adopted an all-fish oil diet and has been trying to convince Marzanna to join her. Marzanna, however is skeptical, since she's worried about maintaining balanced nutrition. Dominique assures her that her favorite lifestyle blogger recommends it, and that's all she needs to know. |
| Argumentum ad populum | Kendra and Karl are two astronomy students. Recently they have become better friends after realizing they both believe that it's entirely possible UFOs could have visited earth. One day at a department meeting Professor McCreary overhears them discussing this possibility. "I can't believe what I'm hearing!" McCreary exclaimed, "Everyone knows that the very idea of UFOs is silly; thus the idea that UFOs have visited Earth is impossible!" | General Warren is testifying in front of congress. The country has been at war for months on end, and this has led to unrest at home. The subject of the committee hearing is possible subversive activity on the homefront. "I believe now, more than ever," testified General Warren, "that we are at no risk of an attack. A recent public poll suggests that few people believe an attack is coming. This means there is no chance we could be in danger." | Silvio is the CFO of a large multinational corporation. In the most recent meeting with investors, he recommended that the firm open a new office in Macau to take advantage of emerging markets. The chairman of the board is skeptical however, arguing that this can't be true. None of their major competitors have opened offices in Asia, and therefore there must be no value in doing so. |

| | | | |
|--------------------------------|---|--|--|
| <p>Argument from ignorance</p> | <p>During Dan's job interview, the interviewer asked him if he had any additional references. Dan said he had a very positive one from his previous employer, Arnold. The interviewer responded that that was nice, but wanted assurances that Arnold was indeed a reputable businessman. Dan responded that he need not worry, nobody seemed to know of any reason to think Arnold was disreputable, so he must be reputable. At the very least, he admired Arnold's intelligence.</p> | <p>The head surgeon and the new doctor, Tom, were preparing for the complicated procedure. The head surgeon asked Tom about his previous training, and whether he understood the procedure. Tom mentioned that his training with Dr. Minsky provided him with valuable experience. The head surgeon replied that he hadn't heard of Dr. Minsky, and asked for his credentials. Tom told him that if Dr. Minsky had done anything disreputable, the surgeon would have heard of him. Therefore, there was no cause for concern.</p> | <p>Marco and Daniel are brothers. Sometimes they amuse themselves by pondering various questions when they go on walks. One day on one of their walks, they began to discuss the existence of the Loch Ness monster. After they had discussed the many legends and supposed sightings of the monster, Marco admitted to some doubt and was still uncertain what to believe. Daniel, however, confidently asserted that because no one has proved that the monster doesn't exist, that it must exist.</p> |
| <p>Argument from force</p> | <p>Martin was an educational researcher who was trying to understand what helps children become successful students. For one study in particular, he gathered information about several variables. He was concerned that if people didn't start coming to the zoo, they wouldn't have any way to develop a love of science and biology. As a result, Martin concluded that visiting the zoo must have a positive effect on students' performance in school.</p> | <p>Mark noticed that the tree in his front lawn was dying. A few weeks later he saw that the flowers he planted near the tree were also dying, causing great concern. He noticed that many of the leaves of the dead tree were piled up around the flowers. The idea that he might be a bad gardener distressed him. He might have to rethink the numerous projects he had planned for the lawn. As a result, he concluded that the leaves from the tree must have killed the flowers.</p> | <p>Tatiana runs a small tech company. She wanted to get employees to bike into the office more often, so she decided to dock their pay if they drive into the office. One employee, Finn, complained that this wasn't fair. "Well," responded Tatiana, "If you drive to work, you earn less money, Sounds like the best option is to invest in the environment and yourself and just start biking in. Seems pretty fair to me."</p> |

| | | | |
|----------------------------------|--|---|---|
| <p>Post hoc ergo propter hoc</p> | <p>Harry was on trial, accused of the first-degree murder of his wife. While on the stand he was questioned harshly by the prosecuting attorney. At the end of a long barrage of questions, the lawyer asked him when he got home that night. From this, the lawyer concluded that since Harry got home at 11 PM and his wife was killed at 11:30, that he was responsible.</p> | <p>Tim was a high school student, who had just completed his math test. He stayed after class gathering his books and noticed the teacher staring at him. After a few minutes, the teacher approached him, and asked him what he did to hide the fact he cheated on the test. She had noticed that he finished in record time and therefore concluded he had cheated.</p> | <p>Eleanor just moved into a new building. She hadn't even had time to unpack when she heard a knock on the door. The building manager, Mr. McClelland was upset, and immediately began a tirade. "What did you do? The building furnace has gone faulty! Just the other day it was working fine. Suddenly you move in and it's broken! This is your fault."</p> |
| <p>Tu quoque</p> | <p>James and Elise were avid sports fans, but this often led to arguments. James claimed that all Wembley fans behave with class at matches. "Then how do you explain the last match against Exeter?" Elise retorted, referencing an incident where drunken Wembley fans were jailed for fighting in the stands. "That doesn't count," James replied, "Exeter did the exact same thing two weeks ago!"</p> | <p>Ben was a physics student. He believed that all scientists must be atheists like him. He shared this with his adviser Dr. Rodriguez, who shared that she's a practicing Muslim. He was visibly shocked. She pointed out that his position was illogical, but he was unswayed. "Someone who holds illogical beliefs can't tell me I'm being illogical," Ben thought to himself as he left her office.</p> | <p>Angelina was incensed. This was the third week she had noticed her daughter's room was a mess and she's had it. When her daughter gets home from school she sits her down and scolds her for not doing her chores. Her daughter is undeterred. "Why should I listen to you? Your room is always a mess! If you can't manage to keep things clean, I don't see why I should have to."</p> |

| | | | |
|----------------------------------|--|--|---|
| <p>Argument from incredulity</p> | <p>Irene's mother is convinced there are ghosts haunting their home, causing noises in the middle of the night. Irene has tried telling her that there's likely a more rational explanation than spirit energy. Her mother, however, responds that given the things she's been hearing she cannot even imagine how else they could happen; therefore, it must be ghosts.</p> | <p>Robert and Liam are engaged in a heated debate about which end should be down when storing eggs, the wide one or the narrow one. Robert argues that the wide end is more stable and therefore better to store the egg on. Liam, however, can't even believe that anyone could store an egg on the wide end, and so concludes the narrow end is right.</p> | <p>Jasmine is starting her second year in arts school. She's excited about a new exhibition of digital media and design. Spencer, a fellow arts student scoffs at the whole idea when she tells him about it. "The best art is made with real media, not artificial devices. It is literally inconceivable to me that anyone would find all that digital nonsense appealing."</p> |
| <p>False dichotomy</p> | <p>Louis and Tara are college sophomores who live in the same dorm. One night near the end of the term, they discuss how they are preparing for their chemistry final. Louis is appalled to learn Tara is planning to cheat on the exam. "I don't see what the big deal is," says Tara, "I have two choices: cheat on the chem final or be entirely embarrassed in front of the class. I hate being embarrassed, so there's nothing wrong with it"</p> | <p>During Thanksgiving dinner, Thom gets into a political argument with his uncle Stephen. Thom thinks the state should eliminate capital punishment since he thinks it's morally unjust to kill someone -- even if they've committed murder. Stephen thinks this is ridiculous. Either you believe in capital punishment, or you believe that criminals shouldn't be punished at all. This means capital punishment must be just.</p> | <p>Uri has recently started working for one of the presidential hopefuls in the next election. He's most excited about the potential for his candidate to be the first female president. Uri's mother is more skeptical, however, since she's not certain of this candidate's stance on some key issues. Uri is appalled. In his mind, you either support this candidate or you're against women's rights. He realizes his mother must oppose women's rights.</p> |

| | | | |
|---------------------|--|---|---|
| Is/ought | <p>Quincy is an economics student. In his last macroeconomics class Professor Carson asked students for suggestions on how to combat inequality. Quincy proposed a graduated income tax that was higher for the rich. Professor Carson remarked that he was surprised by Quincy's argument. The rich currently pay less taxes than the poor, and since that's the way things have been for years, the rich shouldn't pay more taxes than the poor.</p> | <p>Caroline is a 4th grader with a committed group of friends. One day the teacher announces that a new girl, Angela, will be joining their class. Some of Caroline's friends, including her best friend Summer decide to pick on Angela and draw nasty messages on her notebooks. Later, when Caroline suggests they shouldn't be so mean to Angela, Summer argues that Angela is already unpopular. Since they won't change that, they should just keep picking on her.</p> | <p>Karina has been trying to decide whether or not to get laser surgery to correct her eyesight. Her friend Reynold often has a unique perspective. When she asks him what he thinks she should do, Reynold replies that he thinks she should avoid the surgery. After all, her eyesight was naturally bad from birth, and so even though there are procedures to correct it, she should stick with the way she already is.</p> |
| Appeal to authority | <p>Yesenia is home alone on a Saturday night. She had previously planned to meet up with friends, but the combination of bad weather and poor communication caused everyone to cancel. She flipped through the channels on the TV. One of the ads for a high-class whiskey featured a testimonial from her favorite actor, saying how smooth he found this whiskey. As a result, Yesenia resolved to drink whiskey until she fell asleep.</p> | <p>Dr. Grisham has just announced she is running for president. She is planning to run a pro military campaign. During the press conference she was asked if she would support military action in San Seriffe, a tropical nation in a violent civil war. "I believe the answer is clear," she responded, "I have personally consulted the astrological charts, and they show that it is our duty to intervene. "</p> | <p>Jasmine is starting her second year in arts school. She's excited about a new exhibition of digital media and design. Spencer, a fellow arts student scoffs at the whole idea when she tells him about it. "The best art is made with real media, not artificial devices. It is literally inconceivable to me that anyone would find all that digital nonsense appealing."</p> |

| | | | |
|-------------------------------|--|---|---|
| <p>Appeal to ignorance</p> | <p>Stephanie is out for brunch with her aunt Glenda. Stephanie ordered a breakfast bacon cheeseburger, with extra bacon and a fried egg on top. Glenda mentions to Stephanie that maybe she shouldn't eat such fatty food, since recent science indicates high-fat diets can cause cancer. Stephanie responds that she's never heard anything about these so-called "scientific" studies, so that clearly can't be true.</p> | <p>Ethan and Carl are grabbing a happy hour drink. After a drink or two Carl reveals that for the last couple weeks he's been cheating on his boyfriend Jacob with someone from work. Ethan chastises Carl saying that infidelity is just morally wrong. Carl is unrepentant. He tells Ethan that he's never seen a convincing argument that cheating is immoral. Since Carl's never seen it proven, it simply can't be the case.</p> | <p>Vasily and Chandra are two computer programmers working for a small tech company. Vasily has expressed to Chandra that he is genuinely terrified of the coming AI revolution, when computers rise up and destroy humanity. Chandra dismisses his fears as alarmist nonsense. Still, Vasily argues that there's no proof we won't create strong artificial intelligence, and therefore it's all but inevitable.</p> |
| <p>Appeal to consequences</p> | <p>Anne is distraught. After five years of being together she and her long-term partner Jethro decided to break up. Anne's mother is less than sympathetic. "Annie, ending a marriage is a shameful thing," she says to her one day, "We have to avoid that at all costs. Therefore, you should accept that it's your fault and apologize to him."</p> | <p>The major looked around at the stunned control room. Everyone in the room looked on in silence as the monitors showed the satellite disintegrating in the Martian atmosphere. Suddenly, he was furious. If this was a complete failure, they would lose funding for the next mission. Thus, he concluded that this was not a failure at all.</p> | <p>In his most recent publication, historian Weston Adamo covers the history of scientific theories in the 19th and 20th century. His book been controversial. In it, he claims is that if people accept the theory of evolution, it leads to a strong belief in eugenics, which has negative effects on society. Therefore, the theory of evolution must be false.</p> |

Appendix B

The table included in Appendix B contains the set of qualitative responses for the two fallacies, *circular reasoning*, and *argument from authority*. These are the responses that were subject to an exploratory qualitative analysis in order to pilot broader analysis of the full qualitative dataset.

| Fallacy Type | Group | Response |
|-------------------------|--------------|--|
| Argument from Authority | Comparison | Both Dominique and Sigfried are being influenced by outside forces; a lifestyle blogger and the council of Bishops respectively. They place their faith in these influencers despite the information and concerns presented by Marzanna and Oscar. |
| Argument from Authority | Comparison | Both passages describe debates in which one person advocates for a certain perspective solely based on the recommendations of uninformed third parties, and the other person disagrees with this perspective for a logical reason. Dominique based her diet on a lifestyle blogger's opinion, and Sigfried based his wariness of vaccination on the opinion of the council of Bishops; in both cases, the source of the opinions is not qualified to make such judgments. On the other hand, Marzanna is concerned about her nutrition balance and Oscar is concerned about protecting others from harm; both of these arguments are based on logical benefits. |
| Argument from Authority | Comparison | both involve a person agreeing or disagreeing with a measurement based on whether their favorite person or a famous person says so |
| Argument from Authority | Comparison | People in both passages are listening to only one argument/perspective and are using only that to justify their stances. |
| Argument from Authority | Comparison | Marzanna and Oscar are reasonable and logical, whereas Dominique and Sigfried like worshiping false idols and following them blindly. Marzanna is skeptical for a logical reason, but Sigfried is skeptical for a stupid reason. |
| Argument from Authority | Comparison | Both people are making important decisions based on the words of individuals who are not qualified to determine what is best in these situations. For instance, a lifestyle blogger probably does not know very much about nutrition. it would be smarter to make a decision based off of the advice from a trained professional in the field such as a nutritionist or dietician. In the other scenario, Sigfried would be better off listening to a trained medical professional such as a doctor or looking at scholarly research that has been done on vaccinations as the council of Bishops decision is not a reliable source for information on vaccinations. |
| Argument from Authority | Comparison | in these passages, two couples of people are debating over a topic. yet in the end, one person in each couple considers a biased source in their evidence in order to support their point. |
| Argument from Authority | Comparison | Both Marzanna's and Sigfried's decisions were made based not on an analysis of the positives and negatives of making the decision, but on the authority of another that they view to be credible. Of course, the credibility of the lifestyle blogger and the Bishops is not entirely solid, but the people following what they say wholeheartedly believe in what they say. |

| | | |
|-------------------------|------------|---|
| Argument from Authority | Comparison | Only looking to one source for affirmation of what you already believe. |
| Argument from Authority | Comparison | Dominique and Sigfried are both attempting to convince their friend to adopt an unhealthy habit because an untrustworthy source told them to. The all-fish oil diet can be compared to a lack of vaccination. |
| Argument from Authority | Comparison | In both passages there is a debate going on between two people. In both passages, one person has a clear and strong opinion. In the first, Dominique is advocating an all-fish diet. In the second, Oscar is arguing for vaccination. In both passages the other party is not in full agreement, and shows some hesitation. In the first, Marzanna is described as "skeptical". In the second, Sigfried is described as being "more wary". |
| Argument from Authority | Comparison | Dominique and Sigfried both try to convince someone else of doing something that an unreliable resource supports. Dominique wants Marzanna to join her diet because it is supported by Dominique's favorite blogger. Sigfried wants to convince Oscar to not vaccinate their kids because of a declaration from the Bishops. Both sources used are not reliable since a blogger is most likely not a nutritionist and does not know the risks of the diet and the Bishops do not have scientific evidence of the benefits or claimed risks of vaccinations. |
| Argument from Authority | Comparison | Both passages are a conversation between two people about health-related choices. In each passage, one person is skeptical of the claim the other makes. In both passages, someone appeals to the testimony of some authority to back up their claims. |
| Argument from Authority | Comparison | In both of these cases, two people are disagreeing with each other. While one person is doing something for their own sake or the sake of other people, the other disagrees and wants to do something else just because they heard about it from someone else. Sigfried believes that vaccinations should be outlawed because the "council of Bishops recently declared that vaccination is a dangerous practice." Dominique adopted an all fish oil diet just because her "favorite lifestyle blogger recommends it." |
| Argument from Authority | Comparison | Both passages have opposing views between two people. One person feels that something is dangerous for health, while the other person assures it is safe. |
| Argument from Authority | Comparison | In both passages, one of the characters is deciding whether or not to act on a decision. They are using sound logic to make up their mind. However, the second character in both passages jumps in and argues against the first character by using evidence from outside non-practical sources. |
| Argument from Authority | Comparison | Both passages feature people making decisions based on the decisions of an authority figure/someone who has influence in their lives. In the first passage, Sigfried chooses to believe what the council of bishops believes, likely because he looks up to them. The same goes for the second passage, where Dominique believes an all-fish oil diet is beneficial just because her favorite blogger thinks so. |

| | | |
|-------------------------|------------|--|
| Argument from Authority | Comparison | The two groups in each passage have conflicting perspectives on an issue. Both Sigfried and Dominique use niche information to justify their perspectives. |
| Argument from Authority | Comparison | Both people contain two people. In these passages, the two people differ in their opinions on a certain topic. Both passages contain a person who uses a source they believe to be trusted as reasoning for their argument. |
| Argument from Authority | Comparison | In both passages, there is an argument over one subject, vaccinations and diets. Oscar and Marzanna are both worried about over well-being of themselves and others, which Sigfried and Dominique draw their conclusions and arguments about the subjects by using influential figures (bishops and bloggers) as justification. Sigfried and Dominique do not seem to take into consideration the thoughts of the others, and just believe in whatever they hear. |
| Argument from Authority | Comparison | In both scenarios, one party is trying to convince the other party to switch their beliefs, stating other sources backing their points and arguing that since those sources said it they must follow. However, in the first scenario there is a debate about vaccination and an argument for vaccination is for the greater good, while the other is about diet and the argument is centered around good of self. Further, the first source seems to cite a more reputable source than the other. |
| Argument from Authority | Comparison | In both passages there are two people debating certain health topics. In the first one it is whether or not they should vaccinate their children and in the other it is whether they should adopt a certain diet. In both passages there is one person presenting faulty evidence to prove their point to try to convince the other person to join them. |
| Argument from Authority | Comparison | Both passages describe a conflict between two subjects where one subject argues for the widely known point, while the other argues for an alternative point, mainly because someone of high credibility to them claims that. |
| Argument from Authority | Comparison | Both passages have characters that do not agree on an issue and each provide a strong argument. For example, in Sigfried and Oscar's debate on vaccination, Sigfried is against vaccination because his/her religion states that it is a dangerous practice while Oscar is for vaccination because he thinks it is moral and a preventive method. In Marzanna and Dominique's case, Marzanna believes in a balanced nutrition where as Dominique supports an all-fish diet because that's what her favorite lifestyle blogger recommends. Another similarity includes the weak arguments one of the characters present in both sides. Dominique and Sigfried only have developed their beliefs based on what others have said. |
| Argument from Authority | Comparison | Both passages assume that since a person/group of influence told them a fact, it must be true. In the first passage, Sigfried does not want to vaccinate due to Bishops telling him it is "dangerous". In the second passage, Dominique assumes that since her favorite lifestyle blogger recommends it the fish oil diet is healthy/safe. |

| | | |
|-------------------------|------------|--|
| Argument from Authority | Comparison | Both of them use sources that are not qualified to give advice on each respective issue. The council of Bishops has no idea what the actual science is behind vaccinations and in all likelihood, Dominique's lifestyle blogger has no true knowledge of what is healthy for the human body. |
| Argument from Authority | Comparison | Both passages are founded upon the belief that a testimonial from a non-expert is a valid source of evidence and should be trusted. Sigfried alludes to the council of Bishops and how their belief about vaccines is valid despite the council of Bishops presumably not being a medical organization. Additionally, Dominique says that her lifestyle blogger, who is not a nutritionist, knows bet when it comes to a healthy diet |
| Argument from Authority | Comparison | both passages display two people arguing about a topic; one person is steadfast in their belief while the other holds the opposite opinion and is more skeptical; both people who rejected the belief are both concerned about people being harmed. |
| Circular Reasoning | Comparison | They reflect circular arguments. X causes Y because Y is caused by X. |
| Circular Reasoning | Comparison | The first passage argues that many bad events happening in a row must be related, which is not necessarily true. The second passage argues that something is "valuable" because its distinctive and "distinctive" because its valuable. There's no real argument because something that is distinctive doesn't have to be valuable and vice versa. |
| Circular Reasoning | Comparison | Both Claudia and Janene are making circuitous, self-confirming arguments. Claudia argues that she is cursed because bad things are happening and that bad things can happen only because she is cursed. Likewise, Janene argues that building should be saved because it is architecturally valuable and there lends downtown a distinctive character, and that the buildings have a distinctive character because they are architecturally valuable. |
| Circular Reasoning | Comparison | Circular logic, explaining something by using the same something... In the first one, the curse explains the bad events, and the bad events explain the curse. In the second one, the architecture is unique and adds a distinctive character, and adds a distinctive character because the architecture is unique. |
| Circular Reasoning | Comparison | Both passages have a person believe something true, and there is a friend or other person in which they share this belief. Both have the person with the belief effect the other in a certain way of thinking, however it is not the same result |
| Circular Reasoning | Comparison | Both paragraphs surround one friend discussing their strong beliefs with another friend. The friends who are stating what they believe in list three reasons why they believe this way. Claudia says that being fired. losing her phone, and getting her bike stolen are three reasons why she has been cursed. Janene says the warehouses shouldn't be torn down because they are architecturally valuable, lend the town a distinctive character, and add to the towns distinctive character. Similarly, in both stories their arguments are not validated by these reasons and the friend listen questions her beliefs. |

| | | |
|--------------------|------------|--|
| Circular Reasoning | Comparison | Both display someone who's reasoning is cyclical. Claudia believes her week was bad because of a curse and that she was cursed because her week was bad. The warehouse is architecturally valuable because it makes the town distinctive and it makes the town distinctive because it is architecturally valuable. |
| Circular Reasoning | Comparison | In both situations, one person uses circular reasoning: the first statement proves the second, but the first statement is only proven if the second is true. |
| Circular Reasoning | Comparison | In the first passage, Claudia makes a logical error in claiming that the only possible reason for such a bad week is a curse. She fails to explain other options. In the second passage, Janene uses circular reasoning in an attempt to support her position on why the warehouses should be saved. She claims that they are architecturally valuable because they lend a distinctive character to the town, and the way they do that is by being architecturally valuable. |
| Circular Reasoning | Comparison | Both Claudia and Janene use their own belief (that they have been cursed or that the old warehouses are architecturally valuable) to justify their own belief. Nicole and Carmelita question Claudia and Janene's beliefs, respectively. |
| Circular Reasoning | Comparison | The main voices in both passages only have one explanation for their opinions. Claudia believes she had a terrible week because of a curse and nothing else, and Janene says that everything the old warehouses can provide for the town is only because they are architecturally valuable. |
| Circular Reasoning | Comparison | In both passages there is a character adamant about their views. In the first, Claudia about her curse, and the second, Janene about the valuable architecture. |
| Circular Reasoning | Comparison | Both Claudia and Janene are arguing their points through circular reasoning. |
| Circular Reasoning | Comparison | Both passages describe people who use logic to affirm a conclusion that they drew from that logic. They simply go in circles, repeating the same thing without any outside evidence. Claudia claims that the reason she is cursed, is because she has had a terrible week because of a curse. Janene says the reason warehouses add to the town's character is because they are architecturally valuable, but they are apparently architecturally valuable because they add character. |
| Circular Reasoning | Comparison | Both are making circular arguments. Claudia is saying that she knows it's a curse because it has to be a curse, while Janene is saying that the warehouses are valuable because they add to character because they're valuable. Neither women add reasoning behind their claims, but rather, they think the claims should support themselves. |
| Circular Reasoning | Comparison | They both make wrong associations (a curse with terrible week and a warehouse with character). |

| | | |
|--------------------|------------|--|
| Circular Reasoning | Comparison | Both of our protagonists are running their argument in circles. The answer to one question is the formulated next question, that will be fruitlessly answered by a statement version of the original question. Warehouses are architecturally valuable because they lend the town a distinctive character, and simultaneously warehouse are able to lend the town a distinctive character because they are no architecturally. Similarly, Claudia has been having such a bad week that she must be cursed, meanwhile she knows she's cursed because she had a bad week. |
| Circular Reasoning | Comparison | Both passages feature more faulty logic, as people use circular arguments and beg the question to justify a conclusion only using her own personal feelings as evidence for the conclusion. In the first passage, the warehouses are architecturally valuable because they're architecturally valuable. In the second passage, there must be a curse because only these things would happen if there was a curse. The final sentence is critical. |
| Circular Reasoning | Comparison | Both passages use circular logic. Janene says the buildings are valuable because they add to the town's character because they are valuable. Claudia thinks she's cursed because she had a bad week and a week so bad could only be experienced by someone cursed. |
| Circular Reasoning | Comparison | Both Janene and Claudia are not using logic here; Janene believes the warehouses are architecturally valuable so they add to the town's character, but has no reasoning as to why they are architecturally valuable; similarly, Claudia thinks that curses are real and that bad thins will not happen together otherwise. Both Carmelita and Nicole seem more predisposed to logic. |
| Circular Reasoning | Comparison | Both passages deal with some sort of superstition. Janene believes that the old warehouses give the town a distinctive feel, while Claudia believes that her week has been cursed. Both passages also involve a conversation between the two characters and their friends. |
| Circular Reasoning | Comparison | In both passages, two dialogues are occurring between two friends. However, they differ in content-matter and the way the discussions are handled. In the first, both Janene and Carmelita are relatively on the same page. This is not true for passage two, in which Claudia and Nicole disagree on the reason Claudia's week has been so awful. In both, one character (Carmelita and Nicole) attempt to soothe and interact with their friend. However, in the first Carmelita asks questions as to Janene's reasoning. In the second passage. Nicole does not pose questions and only tells her that she is wrong, offering a different point of view. In the first, the character who is upset poses reasonable explanations for her opinions, whereas in the second no reasonable explanation is given. All that is said is that "a terrible week could only be the result of a curse." |
| Circular Reasoning | Comparison | In both passages, one subject uses their argument as evidence rather than actual evidence. For example, Janene believes that warehouses are architecturally valuable because they add to the town's character given their architectural value (backing up their claim by restating their claim). Claudia believes that she must be cursed because such a bad week must be due to a curse (again, explaining her original claim by using her claim as evidence). |

| | | |
|-------------------------|------------|---|
| Circular Reasoning | Comparison | Both passages do describe situations of people jumping to early conclusions. However, I don't really think the reasoning in the passage about Janene and Carmelita is necessarily entirely flawed. It may just need some further clarification. |
| Circular Reasoning | Comparison | I see no similarities between the passages; one discusses people's opinions about a town, and the other discusses someone who thinks she has been cursed. No relation. |
| Circular Reasoning | Comparison | The first passage talks about a girl who is arguing to her friend that an old warehouse downtown should be kept. The second passage talks about a girl who is having a difficult day in her life and believes she is cursed despite her friend telling her that she isn't. The similarities between the two is the expression of the idea of "friendship" and "loyalty". In the first, Carmelita does not understand why her friend Janene wants to keep the building and so decides to hear her out and give her a chance to express herself. In the second, Claudia was having terrible day, so she vents to her friend Nicole who tries to reassure her since she really cannot do anything more. Still she is being a good friend and is staying by Claudia's side. |
| Circular Reasoning | Comparison | Both Janene and Claudia are falling prey to circular reasoning. Janene thinks that the warehouses are architecturally valuable because they give the town character, and that they give the town character because they are architecturally valuable. Claudia believes that she was cursed because she had a bad week and that she had a bad week because she was cursed. Neither person can support their reasoning with anything other than the supposed results and just continue to use circular logic. |
| Circular Reasoning | Comparison | In both passages, the characters appear to be stubborn and to support their claims with no evidence. |
| Circular Reasoning | Comparison | Both discuss concepts that someone thinks can only be explained by one idea. |
| Circular Reasoning | Comparison | The parallels here lie in the way in which Carmelita and Claudia approach their arguments. They both make a claim, in one case that warehouses are valuable and in the other that she is cursed, then these claims are questioned by a friend. Upon the view being questioned, they double down on their respective stances without giving much of a reason as to why they are doing so. |
| Argument from Authority | Sequential | Dominique is relying too much on a single authority with no obvious credentials. A way to respond to this would be to point out all the different scientific experts that disagree with the lifestyle blogger, or even all the other lifestyle bloggers who do so. |
| Argument from Authority | Sequential | Dominique obviously is privy to ethos since she thinks that her favorite blogger is infallible with respect to food preferences. She thinks that her favorite blogger knows more than accepted beliefs. One way to respond to this is to talk with a real doctor. |

| | | |
|-------------------------|------------|---|
| Argument from Authority | Sequential | Dominique does not actually consider the potential negative effects that her new diet could have on her ability to maintain balanced nutrition. She has rationalized her diet choice as an exceptional one because she is emulating the habits of someone she looks up to, i.e. her favorite lifestyle blogger, and assuming that the blogger must have already thought through these potential issues and concluded that they do not apply. A possible response would be to inform Dominique that the posts created by the blogger only showcase positive (and sometimes faked) elements of her life, and that negatives are very likely to have been left out. |
| Argument from Authority | Sequential | Dominique is assuming that because one person said something (her favorite lifestyle blogger) it must be true. Individuals can often stretch the truth or lie, so it is best to consider multiple sources when doing research on something. |
| Argument from Authority | Sequential | The lifestyle blogger is giving advice that she is not educated in the field to give. Dominique should not be relying on her information to make decisions that regard her nutrition and well-being. |
| Argument from Authority | Sequential | Dominique needs to realize that she can't trust everything she hears online. People you look up to can say whatever they want and get away with it, it doesn't mean that it's right. |
| Argument from Authority | Sequential | Dominique is blindly relying on a blogger's opinion that an all-fish oil diet is beneficial. This is not an example of good reasoning since the blogger could recommend unhealthy diets, and Dominique would believe her. |
| Argument from Authority | Sequential | just because Dominique's favorite lifestyle blogger recommends the fish oil diet does not mean it is good for marzanna |
| Argument from Authority | Sequential | Dominique incorrectly assumes that a diet is healthy for her because a lifestyle blogger recommends it. Dominique should recognize that a lifestyle blogger is a bad source of information when it comes to health, and they are often paid to promote such products or ideas. Marzanna is correct in her thinking, as it has been shown that the human body needs a balanced nutrition to remain healthy. Dominique trusting a lifestyle blogger for information on diets is equivalent to trusting an accountant to write a prescription for the best medication to treat an illness. You would want to get information from an expert in the field, rather than a lifestyle blogger. |
| Argument from Authority | Sequential | Just because Dominique's favorite lifestyle blogger recommends it, does not mean it actually works. The lifestyle bloggers credentials and expertise on nutrition are unknown and thus a foundation has not been laid for Marzanna to make a good judgement as to whether the diet is healthy. |
| Argument from Authority | Sequential | The erroneous reasoning is that since one online personality recommended a routine, that it must be okay. One way to respond to this is to draw an example of something that is seen online that is not true, like a false Facebook news update. |

| | | |
|-------------------------|------------|--|
| Argument from Authority | Sequential | Dominique has been influenced to follow a diet that someone else who she respects is doing and as a result, believes it is healthy enough that she is willing to convince her friend to do the same. However, just because someone famous or with influence does something, doesn't mean that it is good or correct. |
| Argument from Authority | Sequential | two girls are trying new dieting methods. one girl wants the other friend to try the diet she is on, but she is skeptical. to prove the diet's merit, the girl claims that a lifestyle blogger recommends it, so it is good. lifestyle bloggers are not trained nutritionists, however, so this is not necessarily accurate. for example, many bloggers promote items that they do not even use or try in order to get money from companies. |
| Argument from Authority | Sequential | Dominique has started a diet that her favorite lifestyle blogger recommended, an all fish-oil diet. Marzanna her friend is skeptical, but Dominique stands firm that because the blogger recommended it it must be true. |
| Argument from Authority | Sequential | Dominique is being influenced by another person's lifestyle blog because they are someone she looks up to. Marzanna is hesitant to join, perhaps because she does not follow the same blogger. Both of them should do their own research before jumping to a diet change, not just do it because someone else does it. |
| Argument from Authority | Sequential | Dominique trusts her favorite lifestyle blogger in her choice of diet, but Marzanna is skeptical. Marzanna could respond by explaining lifestyle bloggers are not necessarily good role models. |
| Argument from Authority | Sequential | Dominique is using the authority of her favorite lifestyle blogger to vouch for the diet. A way to respond is to point out that the blogger could lack credentials and there is no guarantee that the blogger knows anything about nutrition, or that the blogger does not profit off of advertising the diet. |
| Argument from Authority | Sequential | Dominique wants Marzanna to join her in an all-fish oil diet that she claims must be safe because her favorite lifestyle blogger recommends it. This reasoning is erroneous because the lifestyle blogger is not necessarily credible just because they are Dominique's favorite. Marzanna could point out that any blogger could post any information online, whether it is correct or not. |
| Argument from Authority | Sequential | Marzanna believes that an all-fish oil diet is healthy b/c her favorite bloger recommends it. Marzanna thinks it might be an unbalanced diet because it consists of eating only one thing. You could say that an all-fish oil diet is unbalanced because fish oil lacks many of the nutrients required for healthy functioning (i.e., fiber, carbohydrates, minerals, vitamins). |
| Argument from Authority | Sequential | Dominique is saying that Marzanna should trust one source but Marzanna is logical in knowing she should maintain a balanced diet and she shouldn't be listening to Dominique if she's basing her information off of a lifestyle blogger. |

| | | |
|-------------------------|------------|---|
| Argument from Authority | Sequential | Dominique is trying to convince Marzanna to adopt a radically different diet. Marzanna expresses some reservations about the diet but Dominique assures her that it must be fine because her favorite lifestyle blogger recommends it. Marzanna presents a flawed argument because using a stranger on the internet's experience should not be the sole determiner of whether or not something is a good idea. Blindly following opinions on the internet can lead to disaster. Dominique and Marzanna should seek the opinions of health and diet experts to determine if the diet is a good idea for them rather than trusting the opinions of a random individual on the internet. |
| Argument from Authority | Sequential | Dominique is using only one source as evidence. Specifically, a lifestyle blogger is not the most reliable source for good information. Marzanna should look into more evidence about whether a fish-oil diet is good for people and make her decision based on what she finds. |
| Argument from Authority | Sequential | Dominique is trying a new diet that's promoted by her favorite blogger. Due to her obsession with the blogger, she started to adopt the diet as well. For example, if her blogger commits suicide, is she gonna kill herself as well? |
| Argument from Authority | Sequential | Marzanna is skeptical because an all fish oil diet does not seem very balanced. Dominique believes it is good because one blogger that she likes recommends it. However, she does not actually know the ramifications of it, and the blogger is most likely not an expert. Also, one person recommending it is not sufficient evidence that it is the best plan of action for everyone else. Dominique did not do much other research. |
| Argument from Authority | Sequential | Dominique has adopted an all fish oil diet that was recommended to her by an online lifestyle blogger and tries to convince Marzanna to join her by stating that if the blogger likes it, it must be good for you. Marzanna is skeptical because it is pretty much impossible to maintain a healthy diet and eat only fish oil. Also, Dominique is blindly following the advice of an online persona that is probably not certified to give out nutrition advice, and Marzanna sees that it is a scam. |
| Argument from Authority | Sequential | External influences are impacting one's opinion. Start thinking for yourself |
| Argument from Authority | Sequential | Dominique is arguing that her all-fish diet is safe because someone that she thinks lives a good and healthy lifestyle follows it. However, this is illogical because that person could be advertising a bit of a false or glamorized life on the internet, or she could have a completely different body system than another person who this diet may just not work for. Anyone can post anything on the internet and claim it works, so the girls should be wary of the diet. |

| | | |
|-------------------------|------------|---|
| Argument from Authority | Sequential | Marzanna claims that because a lifestyle blogger recommended a certain diet, that it must be good. A response could be that many lifestyle bloggers do not actually back up their claims with research, so you shouldn't blindly follow them |
| Argument from Authority | Sequential | Sigfried is assuming that the council of Bishops is the best source of information about the dangers of a scientific practice. Oscar could respond by pointing out that this is not the case, and that the organizations that are actually experts on the matter agree that vaccination is good, and the council has no real basis for their statement. |
| Argument from Authority | Sequential | There is a conflict between religion and science here in which the authority of the Bishops is questioned. People that are extremely science oriented and those who are overtly religious will never reach an agreement so perhaps they should just agree to disagree. |
| Argument from Authority | Sequential | The error is that Sigfried is allowing a group of non-medical experts (bishops) to advise him on how he should medically care for his children, when he should actually look towards medical professionals for guidance about medical issues. You could respond by asking Sigfried if he would refrain from praying if a doctor told him that it was a dangerous practice. |
| Argument from Authority | Sequential | Sigfried is seemingly relying on only one source for the information to form his opinion, when it is certainly possible that this source could be wrong. |
| Argument from Authority | Sequential | Sigfried is claiming that vaccination should be outlawed due to the "declaration" of the council of Bishops but this is flawed because the council of Bishops are not medical professionals and therefore cannot determine whether or not vaccination is a dangerous practice. |
| Argument from Authority | Sequential | Vaccination does not have anything to do with the church, it is a necessity for people to use vaccinations because they protect the rest of the population. |
| Argument from Authority | Sequential | Oscar is reasoning that they should vaccinate their kids since it is for the better for others around them. Sigfried's reasoning against vaccination because he believes in the credibility of the council of Bishops. Sigfried's reasoning is more unreliable than Oscar's since he is simply basing his opinion off the credibility of another source. |
| Argument from Authority | Sequential | oscar thinks vacines are good whereas sigfried thinks they are bad. |

| | | |
|-------------------------|------------|---|
| Argument from Authority | Sequential | Sigfried and Oscar are considering whether or not they should vaccinate their kids. Oscar provides a much stronger argument, citing how it will help protect their kids and others from harm. However, Sigfried incorrectly argues that vaccinations should be outlawed and cites an unscientific source to back up the claim. It isn't accurate to assume that because the council of Bishops is the best source for information regarding vaccines, but rather the information should come somewhere within the science community when making a decision such as this one. |
| Argument from Authority | Sequential | Both Oscar and Sigfried are not presenting scientific evidence that shows and proved that vaccination is essential for the health of a society. Sigfried's argument is especially weak, as he calls what Bishops have said about this issue--Bishops are not scientists and therefore they cannot declare that vaccination is a dangerous practice. |
| Argument from Authority | Sequential | The reasoning in both arguments is problematic, as they both rely on a moralistic component to support their claim. Sigfried's argument for outlawing vaccinations relies on the say-so of a group of bishops, and relying on that to make a legal determination crosses the boundary between church and state. One way to respond to this claim is to point to evidence of how safe vaccinations are. Additionally, even while I agree with Oscar's claim, it is still a moral value judgement. A better method would be to frame the point not as a moral obligation but as a material benefit for the society, as it will lead to fewer people indisposed or dead due to sickness. |
| Argument from Authority | Sequential | Sigfried is arguing that vaccinations are unnecessary and harmful because a non qualified source told him so, however if a source that has no knowledge on this topic makes a claim, that doesn't mean that it is true or should be trusted, and therefore the child should be vaccinated under Oscars point. |
| Argument from Authority | Sequential | in debating whether or not they should vaccinate their kids before school, oscar believes it is a beneficial moral obligation while sigfried believes it is a dangerous, nonreligious matter that should be outlawed. whereas oscar uses his own ideas that stem from his own observation, sigfried merely follows that of the council of bishops. |
| Argument from Authority | Sequential | Oscar thinks everyone should have to be vaccinated to protect everyone. Sigfried doesn't think so, he actually thinks vaccination is dangerous. They should do more research rather than just listening to one source. |
| Argument from Authority | Sequential | Oscar believes that vaccination is important because it both prevents their kids from harm as well as protecting the people they're around from getting the illness the vaccine could prevent. Sigfried is claiming that because of the beliefs in his religion, vaccines should be outlawed. Since a high-ranking official in the Church says vaccination is dangerous, to Sigfried it must be dangerous. Oscar but especially Sigfried could do more research on vaccines. |
| Argument from Authority | Sequential | Oscar feels that vaccinations are needed for the greater good of people, whereas Sigfried thinks it is against his religion. One could prove that vaccinations are useful and safe with science, thus proving Sigfried wrong. |

| | | |
|-------------------------|------------|--|
| Argument from Authority | Sequential | Sigfried is appealing to the authority of the council of Bishops to inform his decision. One way to respond to this is to point out the lack of authority the council has in the field of science and disease prevention. Thus, the council's advice should not be given the same weight as scientific evidence that supports Oscar's point. |
| Argument from Authority | Sequential | Sigfried and Oscar disagree on whether their children should be vaccinated. Sigfried thinks that, because the council of Bishops declared vaccination dangerous, it should be outlawed. This assumes that the council of Bishops is an authority on medical information and that its declarations should inform laws. However, many religions have different beliefs that are not followed by everyone and not made laws. For example, it is not illegal not to attend mass every Sunday, even though Catholics are expected to do so. |
| Argument from Authority | Sequential | Oscar and Sigfried are arguing about vaccination. Oscar is for them, believing that vaccines are effective at preventing dangerous illnesses. Sigfried is against them, since the council of Bishops declared that they were dangerous. You could respond by drawing up peer-reviewed studies on vaccinations, and argue that a variety of dangerous diseases were eradicated by vaccines and then brought back into existence after anti-vaxx became a thing. |
| Argument from Authority | Sequential | I'm not trying to be offensive, but using the council of Bishops as your reasoning to not vaccinate because they've declared it's dangerous even though every science has pointed towards it being healthy -- and not causing autism -- is a really bad judgement call on Sigfried and represents the classic my religion says I shouldn't so I won't even though this is outside of sinning. Also, I thought it was a law that you need to vaccinate your kids. |
| Argument from Authority | Sequential | Siegfried and Oscar are debating the merits and drawbacks of vaccination. Oscar thinks they should vaccinate their kids because it is the most responsible thing to do, however, Siegfried thinks it is a bad idea based on the opinions of religious officials. Siegfried's issue is following the opinions of people who have no expertise in the area of concern. If Siegfried had heard from his doctor that vaccination is dangerous that would be a valid point of concern, but since he is listening to religious officials presumably with no medical experience or expertise, his argument is significantly flawed. |
| Argument from Authority | Sequential | Sigfried and Oscar disagree on where they should place their trust. Oscar believes in science, Sigfried in religion. It is harder to reason with people when they disagree on a fundamental set of facts. Oscar should try to expose Sigfried to more evidence that is tested and deemed true by scientists. |
| Argument from Authority | Sequential | the fact that Vaccination is claimed as dangerous by the council of bishops doesnt prove it actually is and still doesnt make vaccination illegitimized. Therefore, vaccination is needed. |

| | | |
|-------------------------|------------|--|
| Argument from Authority | Sequential | Sigfried and Oscar are getting their information from different places. Oscar sees that not vaccinating his kids could get others sick. However, Sigfried sees the declaration by the council of Bishops and believes it. He does not know where the council got its info or how they made that conclusion. He is also failing to look at the thousands of studies disproving that fact. |
| Argument from Authority | Sequential | Sigfried and Oscar are debating whether or not to vaccinate their kids. Oscar argues that vaccination is a moral obligation because it protects those who are vaccinated from harm and those who are vulnerable to the illness from harm. Sigfried argues that vaccination should be outlawed because a council of Bishops declared that it is dangerous. However, Bishops are religious figures and not doctors and therefore are not qualified to speak on the dangers of vaccination. In fact, it would make more sense for the Bishops to agree with Oscar's moral argument than to claim that vaccination is dangerous. |
| Argument from Authority | Sequential | A religious background is preventing Sigfried to see more than one approach to vaccinations. One can respond by revealing evidence to the benefits of vaccinations |
| Argument from Authority | Sequential | Sigfried is arguing that he trusts the council of Bishops in that vaccinations are harmful to children and that his should not get them. However, this is flawed logic because there have been countless well-documented, scientific studies that prove Oscars point, that vaccinations are a moral obligation to protect many children as well as one's own. One could respond by showing all of the arguments disproving the council of Bishops' claims. |
| Argument from Authority | Sequential | Sigfried claims that they should not vaccinate anyone because the council of Bishops said it was dangerous, but Oscar claims that they should vaccinate because it protects the most people from harm. Most responses to Sigfried or Oscar would probably match that of the other man's: to Sigfried you could say herd immunity is important, to Oscar you could try to pull up random opinions. |
| Circular Reasoning | Sequential | Claudia assumes that the frequent occurrence of negative events must have a specific cause when in reality they may be random and unrelated. |
| Circular Reasoning | Sequential | Claudia exhibits a superstitious belief in curses based on her recent experiences. Attributing a series of occurrences to something is also associated with good luck charms, something has been disproved before. Her reasoning that something supernatural must be to blame for all of her woes takes out any human fault on her part and on others. You can respond by demonstrating that these associations are only mental with examples of fake good luck charms and common misconceptions about ladders, umbrellas, mirrors, etc. |
| Circular Reasoning | Sequential | Claudia has had a terrible week where she was fired, lost her phone, and was robbed. She believes in mystical forces being involved in her demise while Nicole believes its just unfortunate circumstance. This is similar to those who are superstitious and will do things in for good luck or to prevent bad luck. |
| Circular Reasoning | Sequential | she thinks that she has been cursed because of all of these bad things happening at once but i agree with her friend when she says this kinda stuff just happens. |

| | | |
|--------------------|------------|--|
| Circular Reasoning | Sequential | Claudia tries to find an explanation for why her life is going wrong. She thinks its a curse, but it is most likely due to her mentality through the week. |
| Circular Reasoning | Sequential | A string of bad luck, entirely coincidental, has caused Claudia to find extreme ways to explain her predicament. Unfortunately there is no telling whether the future may be better or worse. She should just stick to her normal self and hope the bad things stop happening. |
| Circular Reasoning | Sequential | Claudia is assuming there must be a connection between the various events in her life. Correlation of bad events does not mean there is a single cause. Terrible weeks can occur whether there is a curse or not. For example, if you have three tests in one week but do not study for any of them, you will have a bad week but only because of your choices not because of a curse. |
| Circular Reasoning | Sequential | Claudia thinks she has been cursed because her week has gone so terrible. Nicole tells her that sometimes bad thing happen. If something good happened and a bad happened on the same week she probably wouldn't think its a curse |
| Circular Reasoning | Sequential | Here, Claudia is insistent that someone has cursed her. She reasons, despite her friend's rebuttal that terrible things of this magnitude are only resultant of some evil curse. A response would be that perhaps this superstitious mindset is leading her naturally toward more negative events as a means of justifying her belief. |
| Circular Reasoning | Sequential | Claudia is looking for something to blame her bad week on, and so she has made up a curse to be able to put all her blame towards. |
| Circular Reasoning | Sequential | Claudia believes she has been cursed because she had a bad week. There are an abundant of factors that could have gone into why her week was so bad so to claim that it was because she was cursed is implausible. There is no evidence to support her "cursed" theory. |
| Circular Reasoning | Sequential | alot of bad things happened to the girl recently and she thinks that she is cursed |
| Circular Reasoning | Sequential | Claudia's reasoning isn't rooted in this world. In reality, Claudia has no way to know if she is cursed, so it is wrong to say that she "must" be cursed. For all we know, her own actions may have led her to being fired and losing her phone. |
| Circular Reasoning | Sequential | Curses are not supported by science (the pseudopsychology of it all is not really valid), just because Claudia is having a bad week doesn't mean that there is some agent in the universe explicitly TRYING to hurt her. Also, hypothetically if there was a curse, there is still a lot of information missing: WHO cursed her? WHEN did it happen and WHY did it happen? |
| Circular Reasoning | Sequential | Claudia claims to be cursed because she had a bad week. One could respond to this by saying that everyone has a bad week every once in a while and sometimes bad coincidences just pile up. |

| | | |
|--------------------|------------|---|
| Circular Reasoning | Sequential | Claudia has had a bad week and she is certain it is because there is a curse on her. Her reasoning is that it must be a curse because such a terrible week could only be the result of a curse. Her argument doesn't hold up due to the logic of curses. They have not been scientifically proven. Several bad things happening at once can be a coincidence. |
| Circular Reasoning | Sequential | Claudia has been experiencing a lot of unfortunate events and to make sense of it she thinks she is cursed. But her friend does not agree. She is not cursed but is just having a bad week but she blames cursing to feel better about the events. |
| Circular Reasoning | Sequential | Janene repeats the same argument that the buildings are architecturally valuable and add to the town's character, rather than adding any new information that would support her argument. |
| Circular Reasoning | Sequential | Two friends discussing the existence of old warehouses find themselves caught in a common reasoning error known as a logic loop: the warehouses are valuable because they lend a distinct character to the town, and they add to the distinction of the town by being valuable. A response would be to ask what type of architecture they exhibit, what could take their place, is there actual historical significance, etc. Carmelita needs to find a way of breaking Janene out of the logic loop. |
| Circular Reasoning | Sequential | Janene uses circular knowledge and keeps referring to the idea that the buildings make the town architecturally valuable without addressing why they do so. |
| Circular Reasoning | Sequential | she thinks that this building gives the town great distinction and that it should be saved for that. I think that she is right. if there are certain things that make a town special then they should save them. |
| Circular Reasoning | Sequential | Carmelita is asking Janene to expand upon her argument to support it. Although she feels the same way, she wants to make sure that Janene's beliefs truly do line up with her own. |
| Circular Reasoning | Sequential | Janene perceives extraneous value in the warehouses, purely for the aesthetic they provide. Janene seems to care more about it than Carmelita. Carmelita should probably side with Janene. |
| Circular Reasoning | Sequential | Janene's argument is circular: buildings that are architecturally valuable should be saved because they are distinctive, and to be distinctive they have to be architecturally valuable. This is truly a non-answer. Janene's definitions have no real predictive value and have no weight behind them. |
| Circular Reasoning | Sequential | Here Janene is making a circular argument that because the warehouses are architecturally valuable, they give the city character, which makes them valuable. Her argument seems rather normative, and is not necessarily rooted in fact. |
| Circular Reasoning | Sequential | Janene never defines architecturally valuable, but it is her main point. Her only other one is that they add to the towns character, by being architecturally valuable. While I see where she is coming from, her friend has a better point, even by saying nothing. |

| | | |
|--------------------|------------|---|
| Circular Reasoning | Sequential | Janene believes that the old warehouse should not be torn down because it gives the town character. One argument against her view point is that that space can be used for something more valuable to the town. If the town were to refuse to renovate then society would eventually surpass them and the town would lack the funding to stay a float |
| Circular Reasoning | Sequential | one girl wants to tear down the warhouses and one girl wants to kepp them there |
| Circular Reasoning | Sequential | In this case, Janene never explained the true value of the warehouses. She is making generalized statements and claims that could possibly be true, but there is no way of knowing their validity until she provides evidence or examples to back up her claims. |
| Circular Reasoning | Sequential | Janene is using her claim as her reasoning (using character and the idea of architectural value almost synonymously), therefore she hasn't really answered Carmelita's question (and therefore cannot quantify or really even support her claim) |
| Circular Reasoning | Sequential | Janene has a circular argument that the warehouses are valuable because they give the town character and that they give the town character because they are valuable. One could respond that there must be other reasons why they give the town character other than being architecturally valuable. |
| Circular Reasoning | Sequential | Janene and Carmelita are discussing whether the city should tear down a building. Janene says the building has distinctive character, but when asked to describe what she means she claims it's architecturally valuable. This is not explaining herself or answering the question. |
| Circular Reasoning | Sequential | Janene thinks the buildings should be saved because they are architecturally valued and add character because they are valuable. But she does explain what makes them valuable and instead she just answers questions with the same answers without elaborating. |