The Dissertation Committee for Micah Balser Goldwater

certifies that this is the approved version of the following dissertation:

Role-Governed Categorization

Approved:

____

Arthur B. Markman, Supervisor

____

Catharine H. Echols

____

Zenzi M. Griffin

____

Jeffrey Loewenstein

____

David M. Schnyer

Role-Governed Categorization

by

Micah Balser Goldwater, B. A.

Dissertation

Presented to the Faculty of the Graduate School of

The University of Texas at Austin

in Partial Fulfillment

of the Requirements

for the Degree of

Doctor of Philosophy

The University of Texas at Austin

December, 2009

Acknowledgments

I have many people to thank for the help they have given me throughout the writing of this thesis, and my graduate career more generally. I first want to thank my thesis advisor Art Markman. “Role-governed categories” is his idea. With his mentoring, I’ve been able to base a program of research on it.

I also give special thanks to Cathy Echols. She accepted me into the graduate program, and has enabled my research on cognitive development, not discussed in this thesis.

I thank the remaining members of my committee: Jeff Loewenstein, David Schnyer, and Zenzi Griffin for their helpful comments and discussion that were instrumental in shaping this thesis.

Additionally, I’d like to thank my superiors, Brad Love, Greg Murphy, Ken Kurtz, Dedre Gentner, Brian Ross, and my peers, Eric Taylor, Jon Rein, David Landy, Aaron Hoffman, Ross Otto, Tyler Davis and Marc Tomlinson for unofficial, yet very insightful discussions regarding this line of research.

Finally, I would like to acknowledge Hunt Stilwell, who, with Art, first conceived of “role-governed categories.” This line of research wouldn’t exist without the year or so we were able to work together. Hopefully, we’ll work together again.

Role-Governed Categorization

Micah Balser Goldwater, PhD.

The University of Texas at Austin, 2009

Supervisor: Arthur B. Markman

Theories of categorization typically assume that categories are represented by some set of features that describe the properties of category members. However this view of category representation is incomplete. This dissertation lays out a framework for category representation, following Markman and Stilwell (2001), that creates a taxonomy of categories based on different components of relational structures. Relational categories are categories of entire relational systems while, role-governed categories, are represented as the roles in these systems. Lastly, thematic-relation categories group entities together that play complementary roles within a system.

Four experiments are presented in support of this framework. They contrast thematic-relation categorization with role-governed categorization. Thematic-relation categorization entails categorizing objects together that play different roles within a domain, while role-governed categorization entails categorizing two entities that play the same role across domains. When the two are put in direct conflict, people prefer to form a thematic-relation category because within-domain connections are easier to find than

across-domain connections. The purpose of the four experiments is to examine ways to boost the preference for role-governed categorization, thus revealing underlying processes.

Here, role-governed categorization is facilitated in two ways. Experiment 1 re-frames the question of category formation as novel word extension. Natural role-governed categories have labels while thematic-relation categories do not. This pattern is reflected in the measured behavior as novel labels are extended across members of role-governed categories more readily than across members of thematic-relation categories.

By claiming relational structures are critical to category representation, the framework described in this dissertation predicts that role-governed categorization and analogical reasoning share underlying mechanisms. Experiments 2-4 examine how making an analogy between the members of role-governed categories facilitates forming such categories. When making an analogy, people align the relational representations of a pair of domains, putting entities into correspondence by role, ignoring featural dissimilarities. When analogical comparison is induced, the rate of role-governed categorization is shown to double as compared to a baseline with no such analogical processes. The thesis concludes by outlining several future lines of research generated by unifying the fields of analogy and concept learning.

Table of Contents

Introduction

Investigating the representation of categories is one of the larger areas of research in Cognitive Psychology. While there have been many variants, most theories posit that categories are represented by collections of features that describe category members, e.g., a bird is an animal with wings and a beak. These are *feature-based* categories (e.g. Rosch, 1973). Some representational approaches posit that features are discrete and unbound (e.g. Tversky, 1977), some posit that features are bound by causal theoretical knowledge (e.g. Murphy & Medin, 1985), and others represent features as dimensions, or values along dimensions in mental spaces (e.g., Rips, Shoben & Smith, 1973). However, these theories all share the idea that there are feature-based categories, and the featural information is subordinate to the category label. This means that the features are contained within the category representation, and the label refers to this representation (Markman & Stilwell, 2001).

More recently, there have been arguments that this view of categories does not account for the breadth of natural categories, despite its success explaining results from artificial category learning experiments, (e.g. Gentner & Kurtz, 2005; Goldwater, Markman, & Stilwell, *under review* 2009; Markman & Stilwell, 2001). Many categories are defined, not by featural information subordinate to their label, but by their *role* in larger relational representations superordinate to their labels, i.e. they are *role-governed* categories. For example, *private* and *general* name different ranks in a military hierarchy.

They are defined by their place in the larger hierarchy whose existence allows the individual rank concepts to be induced.

By positing a different form of category representation, the proposal for role-governed categories makes at least two types of (related) predictions. Each of these is discussed in depth below. First, different entities will be categorized together than previously thought. Second, different mechanisms will underlie categorization judgments. For example, if someone notices that a sling is *supporting* a dislocated shoulder, and a tripod is *supporting* a camera, that the sling and the tripod could be categorized together, as each is playing the *supporter* role. In addition, the mechanism of forming such categories could be shared with analogical comparison, which entails aligning relational structures by putting objects into correspondence by role (e.g., Gentner & Markman, 1997).

This dissertation has the following structure: First, I review Markman and Stilwell's theoretical account, discuss its explanatory and predictive power, and compare it to other research on categorization. Second, I review my recent work investigating the representation of role-governed categories empirically (Goldwater, et al. 2009; Goldwater, Markman, Trujillo, & Schnyer, *in prep*, 2009). Third, I present my dissertation experiments that investigate the process of role-governed categorization by examining under what circumstances objects will be categorized by their common relational role. In the fourth and final chapter, I discuss the results and future lines of research working within this framework of category representation.

1. Theoretical Perspectives on Category Representation

In this section I lay out my working model of category representation. This view of representation posits that there are multiple kinds of categories defined by their representational format (including role-governed and feature-based). I first briefly describe all of them, and then go into more detail about each, with a focus of how they relate to role-governed categories. In addition, I motivate this framework for category representation by discussing how it provides a new perspective to some classic problems in categorization, and make several novel predictions.

To place the idea of role-governed categories into perspective, first consider the simple relation *x visits y* illustrated in Figure 2. This figure assumes that people represent relationships among elements in their environment. Proposals for knowledge representation often use structured relational representations to account for these relations (see Markman, 1999, for a review). In a structured representation, there is an explicit connection between elements, so that the scope of a predicate is determined by the arguments that it takes. In the relation x *visits* y, *visits* describes a relation between two elements x and y (which serve as the arguments to the relation). The elements x and y in this case are variables, and so they can be bound to a variety of different objects in different settings in order to allow the cognitive system to represent different instances of the event using the same relation.

Markman and Stilwell (2001) pointed out that the arguments to a particular relation form a class of objects, namely those objects that play that role in some event, situation, or structure. If items that play that role were somehow important to an individual, then a category label could be attached to that role and used as a category. So, a *guest* could be defined as the individual that visits. That is, the category would name the first argument to the relation x visits y in Figure 2. Similarly, a *barrier* might name a class of items, such that x *blocks* y, and a *shield* might name a class of items so that x *protects* y. On this view, then, whenever someone creates a relational system, they open up the possibility for a category that names the roles within that system. Evidence for this view is provided in Chapter 2.

Viewing knowledge as relational structures predicts the existence of multiple kinds of categories distinct from role-governed (as seen in Figure 2): *relational categories, thematic categories* and *feature-based categories*. Here, I will briefly discuss each, before giving a more full treatment to them all in addition to a complete discussion of role-governed categories in the following sections.

A relational category is one that refers to an entire relational structure (Markman & Stilwell 2001, and see Gentner & Kurtz 2005 for a very similar discussion). For example, in x *visits* y, x is the role-governed category *guest* and y is the role-governed category *host*, but a *visit* refers to the relational system *x visits y*, and thus is a relational category. Another good example is the distinction between the relational category *trip*, and the associated role-governed category *destination.* A trip is the whole relational system of travel, relaxing, the destination, etc. *Destination* picks out just one role within the relational system.

Thematic categories (Lin & Murphy, 2001) group a set of items that all take part in a common relational system or event, though they play different roles within that event. So, often people will categorize *milk* with *cereal* even though milk probably bears more similarity to other drinks. This relationship is different from role-governed categorization because thematic relations group things together that play different roles in the same event, as opposed to the same role across different events.

Feature-based categories have been given the most attention in the literature on artificial and natural categories. When researching natural categories, researchers positing feature-based representations primarily focus natural kinds and artifacts, e.g. *dog* and *chair*. These categories are the most commonly used stimuli in semantic priming tasks (e.g. Cree, McRae, & McNorgan, 1999 among many others). They are used for Rosch's classic property listing tasks. In addition, virtually every artificial category learning experiment has subjects learning categories easily represented by features. In our perspective, feature-based categories are critical because for every instance of a role-governed category, there is a member of a feature-based category fulfilling that role. A full treatment of the relations between feature-based and role-governed categories are discussed below

1.1 Relational Categories

I start with relational categories because their representation is primary; role-governed categories are induced from them (Markman & Stilwell, 2001). Relational categories pick out the relations amongst elements in the environment. These relations may only relate two elements, as in x *visits* y, or name more complex systems with many interrelated entities, e.g., *scandal*, which involves the perpetrator, a misdeed, a potential cover up, and a set of people who care enough about the perpetrator that the misdeed is considered a scandal. That is, the misdeed is drastic violation of the expectations for the perpetrator, given his public role. For example, a bum buying underage minors alcohol is not a scandal, while the president doing so is.

There are at least five kinds of relational categories: verb categories, preposition categories, event scripts, taxonomies, and comparatives. The first two receive their name from grammatical categories, as they are reliably lexicalized as such.

Verbs are the primary linguistic medium for encoding complex relations. Their representation is one of the most well studied topics in linguistic theory (e.g., Jackendoff, 1990). Verbs point outward to other concepts; specifically to the entities they bind to their argument slots. They bind their arguments by specifying how their arguments relate. Most frequently verb arguments relate through causal action. By binding their arguments, they assign each a "thematic role, " i.e. their role in the relation named by the verb. For example, if a verb names a causal action, then the verb would assign one argument as the *agent*, or doer of the action, and another as the *patient*, or recipient of the action.

Analogous to how feature-based object categories decompose into their features, verb representations decompose into more primitive relations (Gentner ,1975; Jackendoff, 1990, 2002; Levin & Rapaport-Hovav, 2006; McKoon & MacFarland, 2002; Wechlser, 1995). We give examples of verb relational structures for *give* and *exchange* to illustrate this point (for a more detailed analysis of the semantics of verbs of possession see Gentner, 1975). Give involves a transfer of possession of some object Z from person X to person Y.

Give: X CAUSE [Y POSSESS Z]

Exchange is a reciprocal giving event, so when one person (X) gives some object (Z) to a second person (Y), Y in turn gives a second object (W) to X. That leads to the rather complex relational structure

Exchange: [X CAUSE [Y POSSESS Z]] CAUSE [Y CAUSE [X POSSESS W]]

In these representations, as mentioned above, the place-holders for potential arguments are frequently represented by variables, but we believe that these variables are not just empty place-holders. Instead, they include rich conceptual information about what is typically bound to the variables, i.e. the stuff of role-governed categories (see more below; Ferreti, et al., 2001).

Preposition categories are similar to verbs, as they point outward to other concepts and bind a small number of arguments. Unlike verbs, which tend to describe actions, prepositions tend to encode spatial and temporal relations, e.g., *above*, and *before.* Prepositions tend to play a supporting role to verbs in the composition of the relational structure of events described by sentences. For example, when verbs describe events that entail a change of location, often the resulting spatial relation amongst the verb's arguments will need to be specified by a preposition, e.g., "Bob put the book *on* the shelf." While verbs represent rich conceptual information about their typical arguments (Ferreti et al., 2001), prepositions seem to specify more skeletal perceptual information at most. For example, *along* specifies something linear, but *on* only requires that its arguments have spatial extension, i.e, they are objects (Landau & Jackendoff, 1993).

While verbs refer to actions and events, they only bind a small number of arguments. "Event scripts," however are larger-scale relational knowledge structures that organize and sequence a number of component actions, binding an indefinitely large number of entities (e.g., Schank & Abelson, 1977). For example, *going on a date* entails a series of characteristic actions. Event scripts have a somewhat hierarchical organization (Shank, 1982), as we may have separate scripts for *going on a dinner and movie date*, or *going to a wedding with a date*. Clearly, each has their own set of subcomponents, and is highly interrelated with other event scripts, e.g., *attending weddings.* Support for scripts as representations in memory comes from a variety of sources, including the ability to make inferences about missing pieces of information in a discourse, or recalling more script- relevant information for a passage of text than irrelevant (e.g., Owens, Bower & Black, 1979). Event scripts contain rich relational information, thus role categories based in scripts reflect this richness (Ross & Murphy, 1999; see below).

Taxonomies are large knowledge structures that can bind many elements together for the purpose of classifying elements by their place in that system. For example, the military hierarchy, or any rank based system is a taxonomy. Rank systems exist for the purpose of classifying individuals to establish how they can relate to one another by assigning role-governed duties and privileges. *Family* is another taxonomy; *mother*, *uncle, etc* name roles within this. There has not been much psychological research into taxonomies as relational categories. Often research into "taxonomic categories" refers to research on natural kinds, because biologists have developed taxonomies organizing them hierarchically based on their biological similarities, creating classifications of increasing generality. There has been much less research into taxonomic systems as concepts in their own right.

The final kind of relational category is "comparatives." These state quantitative relations among sets of or properties of entities e.g., *identical, greater*, and *fewer.* Because these relations seem relatively simple, there has not been much research into these concepts in adults, but their learning and representation has been a topic of both developmental (Smith, 1989) and comparative psychology (Premack, 1983). For example, Thompson and Oden (2000) suggest that non-human primates do not display an understanding of the concept `identical ' (e.g. that two identical squares are in the same relation as two identical triangles) unless they have been given training with symbols.

While there has been a lot of research on relational categories, this work is not typically framed explicitly as the study of categorization (cf., Morris & Murphy, 1990). In fact, research on event scripts has sometimes had the goal of showing how different their representations are from categories (Barsalou & Sewell, 1985). One reason that relational categories are not be discussed as if they were categories, is that research into natural categories tend to rely on verbal report (e.g., in a property listing task), and (with the exception of event scripts) most information about relational categories is not verbally accessible (Sieck, et al. 1999). For example, if one were to ask subjects to list the properties of "giving" or "between," subjects would most likely give a definition, and then just seem confused. However, to have a complete theory of categorization, the study of relational categories must be unified with the study of feature-based, e.g., by having subjects learn artificial relational categories using similar paradigms as their learning of feature-based categories (c.f., in developmental psychology, the same paradigms have been used to examine noun learning and relational word learning, e.g. Casasola & Wilburn, 2004).

1.2 Role-Governed Categories

Role-governed categories are induced from relational categories. However, relational categories will differ from each other in their ability to support such inductions, and in theory, role-governed categories will differ depending from which kind of relational categories they are induced. Prepositional categories and comparatives generate the least support for role-governed categories, as they only constrain role-fulfillment to entities that can relate to others spatially, temporally, or quantitatively, which essentially all entities can. All roles are fulfilled by entities that could be cross-classified as members of feature-based categories, e.g. *guests* tend to be *humans*, but if the relational structure places no constraints on role-fulfillment, then there is no motivation to create a concept for that role because classifying something as a member of such a role-concept adds no information about that entity.

Unlike prepositions and comparatives, verb categories support role-governed category creation, i.e. verbs license role-governed categories. Many role-governed categories are clearly related to verbs. As an additional example to *guest* above, the concept of a *thief* relies crucially on verbs like *steal*, because the defining characteristic of a thief is that this individual is the first argument to the relation x *steals* y. Table 1 lists more corresponding verbs and role-governed categories. In addition to these example sets that do not bear a morphological relation, the "–er" morpheme in English allows us to freely derive terms for typical agents, a type of role-governed category, from verbs, e.g. *dance* and *dancer*.

In addition to this linguistic evidence, there is experimental evidence supporting the claim that verb representations include the conceptual information of their typical arguments. Ferretti et al.. (2001) show that in a lexical decision task, verbs prime their typical agents, e.g. *arrest* primes *cops,* their typical patients, e.g. *arrest* primes *criminal,* typical instruments, e.g. *stir* primes *spoon*, and features of patients, *manipulate* primes *naive*. Interestingly, computational models of semantic memory currently cannot account for these findings because they rely on featural overlap between concepts (Cree et al.., 1999); they cannot represent relational structures.

Ferretti, McRae and colleagues' results (e.g, McRae, Ferretti, & Amyote, 1997) are important support for role-governed concepts because if these concepts were parts of relational structures, they would be as specific as their root relational categories. Traditionally, thematic roles have been seen to be very general (Jackendoff, 1990); *agent* is just an entity that causes something. That thematic roles are verb-specific, or as specific as classes of just a few verbs (e.g., Filmore, 1982), is predicted from the current framework.

In addition to verb categories, event scripts support category formation because they are rich knowledge structures. For example, we have lots of cultural knowledge based on scripts for eating. We know how eating lunch, dinner, and breakfast differ, and how eating at an upscale restaurant differs from eating at a fast-food restaurant. One way these scripts differ is simply by which foods are eaten in each situation, creating categories such as *breakfast foods*. Ross and Murphy (1999) demonstrate that foods are categorized by their script-based roles, e.g., *waffles* with *eggs*, in addition to categorizing foods based on featural overlap, e.g., *waffles* with *cookies*.

Critically, Ross and Murphy were first to examine the function of role-governed categories. Categories do not exist just to support classification. The most highly studied function of categories is induction i.e., given the knowledge that an entity is a member of a category, what can be predicted about that entity (e.g., Osherson, et al., 1990)? In the feature-based view of categories, there has been a focus on what features could be induced from shared category membership, e.g., if waffles have a particular nutrient in them, then cookies may be predicted to have that same nutrient. Script-based categories support inductions like these to a lesser degree than feature-based, but Ross and Murphy (1999) demonstrate that they do support inductions about situations. Knowing that waffle*s* were served at an unfamiliar ceremony, subjects were more likely to think that eggs would be served than cookies. Showing the role-governed categories support different kinds of inductions than feature-based shows why humans would be motivated to form such categories.

Taxonomies support role-governed categorization, e.g., *private* and *general* are induced from *military hierarchy*. As discussed above, taxonomies can be used to organize pre-existing concepts, as in the natural kind taxonomy. Presumably, concepts for individual animals predate taxonomic classification, thus, unlike role-governed categories like *colonel, supervisor,* and *cousin,* natural kinds are not representationally defined by their place in the taxonomy. However, even in this case, the existence of the taxonomies licenses categories that are the placeholders for natural kinds. That is, *phlyum, genus*, etc are role-governed concepts derived from a taxonomy used to organize already existing feature-based categories.

Two other proposals about category representations bear some similarity to the idea of a role-governed category. First, Rehder and Ross' (2001) abstract coherent categories come closest to bridging the gap between categories used in artificial category learning experiments and role-governed representations. In their studies, the exemplars of their abstract coherent categories are all represented by distinct feature sets, but what makes the exemplars part of the same category is that each distinct set is made coherent because all aim to satisfy a common function (e.g. to get rid of pollution). Many functionally-based categories are role-governed categories, because they play a particular role within a causal relational structure.

Barsalou (1983, 1985) discusses goal-derived categories that people are able to create in a given situation. For example, the category *of things to take out of the house in the event of a fire* are all objects that satisfy a particular role in a (somewhat strange) relational structure. The proposed framework suggests that goal-derived categories in general may be one type of role-governed category. Barsalou suggested that one difference between goal-derived categories and feature-based categories is that prototypical category members are important for feature-based categories, but ideal category members are important for goal-derived categories. For example, diet foods are a goal-derived category, and the ideal diet food is low in calories and tastes great. It is possible to extend this work to role-governed categories and predict like goal-derived categories, they have more prominent ideals than feature-based categories (see below).

1.3 Thematic-Relation Categories

Thematic relation categories group entities together that play corresponding roles in a relational structure, e.g. *king* and *crown*, *milk* and *cereal*, *dog* and *bone*, *fly* and *dumpster*. Up until recently, most of the research on these categories was conducted with children, because these categories were seen as a developmental precursor to more "mature" feature-based categories (e.g., E. Markman, 1989). However, recent research by Lin and Murphy (2001) shows that adults often exhibit a preference for thematic relation categories over feature-based categories. Lin and Murphy (2001) argue that adults represent both kinds of categories, and they serve different functions. Members of thematic relation categories co-occur in situations, thus, properties acquired by contagion could be spread from one member to another. Lin and Murphy show that knowing one member of a thematic relation category has a certain bacterium, e.g., *fly,* leads to a stronger induction that another member of a thematic relation category, e.g., *dumpster*, carries that bacterium than another member of a feature-based category, e.g., *mosquito*.

Thematic relation categories group entities playing complementary roles in the same situation, while role-governed categories group entities together playing the same role across situations. Each type of category is rooted in learning about situations, i.e., relational structures. However, it is unknown how the two kinds of categories relate in the learning process. Does learning about one entail learning about the other? For example, can one learn about *guests* without learning about *hosts*, and that *guests* and *hosts* form a thematic relation category? Or does each type of categorization involve an attention shift at the consequence of the other kind? No research has yet examined learning novel role-governed and thematic relation categories from novel situations together.

1.4 Feature-Based Categories and Role-Governed Categories

Because feature-based categories are the primary area of study in categorization research, a summary of this work would be beyond the scope of this dissertation. Instead, I focus on the way that feature-based categories fit into the current framework, and particularly how they relate to role-governed categories.

Roles are filled by entities that also have descriptive features. Often, these features become crucial parts of the representation of role-governed categories. In particular, these features are in the service of recognition processes. For example, military officers of different ranks have different numbers of stripes on their uniform. These signs of rank allow superior officers to be recognized without parsing complex social relations defined by respective ranks. Indeed, processing unstructured feature-based representations demands fewer resources than fully bound relational structures (Markman, 1999). To ease such processing demands, we set up our environments to enable the identification of relational structures via featural correlates.

The intentional replacement of relational processing with featural can be clearly seen with the design of common artifacts for home-use. Many artifacts seem to be role-governed categories because they are defined by their functional roles. However, in practice artifact representations may be primarily feature-based because their functional roles are consistently fulfilled with the same features. The consistent use of features, for example, eases recognition of and induction about products at Walmart. Members of the class of blenders or vacuum cleaners, for example, all tend to look the same. To be a blender, all something has to do is blend stuff, but in reality, blenders all blend stuff in essentially the same way, and even when the function is not performed in the same way, many design features are kept consistent with other products. In addition to these reasons for hypothesizing artifacts are represented featurally, they are often the objects of study in research conducted in the feature-based framework, so I will also operate with that assumption. In the next chapter, I present experimental evidence that artifacts are represented as feature-based and differently from role-governed natural categories.

In addition to artifacts, there are other situations where the line between role-governed categories and feature-based categories appears fuzzy because roles are filled by the same objects very consistently, and those objects play those roles very consistently. For example, *homes* are often *houses*, *pets* are often *dogs*, and vice versa. As a result, their representations are likely to share content. However, this overlap does not deny that their representational distributions are separate. And, the difference is not just one of generality. *Pet* is not more general than *dog*, as *animal* is, because even though *pet*'s extension includes some cats and birds, there is not strict nesting. There are many dogs that are not pets, but all dogs are animals. It is an important avenue of future research to investigate the consequences of the aggregation of relational and featural knowledge for both role-governed and feature-based categories.

1.4.1 Representational Format and Category Definitions

In addition to predicting novel phenomena, the current framework offers new perspectives on classic puzzles posed with feature-based views of categorization. The first is whether categories have definitions. From Aristotle to the mid twentieth century, western thought has generally considered categories as definable, leading to clear binary decisions about whether something is a member of a given category. Wittgenstein (1968) countered this long-standing assumption, using *game* as his primary example of a category with no definition. He argued there were no defining features common to all games, but all games share some features with at least one other, forming a *family resemblance* structure. Rosch and Mervis (1975) showed that people learn categories with such structures. However, this argument conflates categories having definitions and having feature-based representations. Markman and Stilwell (2001) consider *game* to be a role-governed category, i.e. the y in x *plays* y. The reason there are no common defining features are because roles can be fulfilled in an indefinite number of ways. By considering categories as collections of features subordinate to the category label, many have missed the potential defining power of superordinate relational structures.

A related example to the confusion about *game* is that at first glance, *bachelor* seems to have a straightforward definition: an unmarried man. Priests fit this definition, but are not considered *bachelors*. However, as discussed in Filmore's "frame semantics" (e.g., Filmore, 1982) *bachelor* is not just a collection of features, but is a point in a stereotypical timeline of male life. This timeline binds phases of life together into a temporal relational structure. This timeline starts with childhood, and eventually gets to married life with kids, and then grandkids and retirement, etc. The reason priests are not bachelors is because they are not going through this typical life trajectory; bachelorhood is a point that assumes marriage is coming next.

1.5 Conclusion

This chapter's goal was to lay out a working model of category representation. In the next section, I lay out my recent work supporting these views. The goal of the first three studies I review was to empirically distinguish natural role-governed categories from natural feature-based. The following two examine the proposed connection between role-governed and relational categories, that novel relational categories will license novel role-governed categories.

2. Recent Work

2.1 Empirically distinguishing role-governed from feature-based categories

We know a lot about our natural categories, and role-governed categories are no exception. A productive way to investigate our knowledge is through a property-listing task (Rosch, 1973; McRae et al., 1997; Bar & Caplan, 1987), thus, this task should reveal differences between role-governed and feature-based categories. Markman and Stilwell's (2001) proposal is that role-governed categories are pieces of, and are defined by, larger knowledge structures. Our knowledge about them should contain a lot about the larger knowledge structures of which they are a part and are defined. In contrast, our knowledge of feature-based categories should consist primarily of descriptive properties of category members.

Barr and Caplan (1987) pointed out that there are two kinds of properties listed in these tasks. *Intrinsic* properties describe category members. For example, being furry is an intrinsic property of dogs. *Extrinsic* properties point outward to other objects. For example, being owned by people is an extrinsic property of dogs. To clarify, I used the term "property" to refer to items listed in a property listing task, while the term "feature" is reserved for a mental representation.

Barr and Caplan's research assumed that all categories are represented as sets of features, and so they did not have predictions about which categories were likely to have relatively more extrinsic than intrinsic properties. They predicted only that categories would differ along this dimension, and that other aspects of their representation, like exemplar typicality gradients, would differ correspondingly. Because my working model makes a claim about different formats of representation, it generates novel predictions about which categories will elicit more intrinsic vs. extrinsic properties. Of course, what is listed is not identical to mental representations, however, what is listed is a window into mental representations.

The analysis I have given of role-governed categories suggests that extrinsic properties should be produced more for role-governed categories than for feature-based categories. I propose that role-governed categories are part of larger knowledge structures, and thus point outward to other aspects of these structures. Descriptions of such relations among entities, by definition, are coded as *extrinsic.* Feature-based categories should primarily elicit intrinsic properties; knowledge about them is relatively self-contained as their features are subordinate to their labels.

Goldwater et al. (2009) conducted a property-listing task with role-governed and feature-based categories, and coded properties listed using Barr and Caplan's coding criterion. The feature-based categories used were artifacts. Artifacts exist because of their functions, and functional information is considered extrinsic by Barr and Caplan, so they provide a strict control condition. When listing properties for "typical examples" of the categories, subjects listed approximately 70% intrinsic properties for artifacts, and 75% extrinsic properties for the role-governed categories, confirming predictions generated by the proposed framework.

To relate these findings to natural use of these categories, I conducted a study of how labels are used to describe images on the world-wide-web using the photo-sharing website flickr.com (see Figure 3). On flickr.com, people upload their pictures and give them titles and descriptions. Any user then can "tag" any of the uploaded photos with a label. For example, there is a photo of a dragonfly on a flower titled "My beautiful guest" that has been tagged with *dragonfly*, *nature*, and *flower* among other things. Essentially, flickr.com can be seen as a corpus of natural categorization free of artificial laboratory conditions. People tag photos because they feel like it, not because they need to fulfill a course requirement. Tagging is relevant, because tagging is hypothesized to be based more on (representations that elicit) intrinsic properties than titling and describing a photo. The person who titles the photo took the picture, experienced the moment it captured, and has a representation of the situation of which the picture is depicting an aspect. However, the taggers just come along see a picture and name it, without all the relational situation knowledge. They probably just name what they see, i.e. the intrinsic properties of the objects pictured.

"My beautiful guest" exemplifies this proposal. The photographer experienced this dragonfly as her guest, i.e. the dragonfly was seen as visiting her. Looking at the picture, not being there to be visited, the tagger just sees a dragonfly and a flower, two feature-based categories.

Flickr.com provides the user with an easy way to verify this proposal. There are two ways to search for photos on flickr.com. One type of search uses all the text associated with the photo, which includes the titles, descriptions and the tags, and the other just uses the tags. Definitionally, the first type of search will get more hits because it is a superset of the second. Categories that primarily elicit extrinsic properties should have a smaller proportion of the superset "all text" search hit by the "tag only" search (the "flickr tagged proportion"). Goldwater et al. searched for all the categories used in the property-listing experiment with both search commands. They found that the flickr tagged proportion was significantly lower for role-governed category labels (M = .2) than it is for the feature based category labels (M =.34). This is important real world support for the role-governed category proposal.

In a second analysis, Goldwater et al. ignored the distinction between role-governed and feature-based categories and created a continuum based on the proportion of extrinsic properties listed for every category used in the previous experiment. The proportion of extrinsic properties was then regressed against the flickr tagged proportion. Increases in extrinsic properties had corresponding decreases in flickr tagged proportion. This effect was quite large. The extrinsic proportion accounted for approximately one third of the variance in the flickr tagged proportion. That results taken from a laboratory investigation with 30 subjects could be such a powerful predictor of millions of real-world behaviors is quite encouraging, and compelling support for the view of categories laid out in this proposal.

Goldwater et al.'s third study distinguished role-governed from feature-based categories. It was inspired by a common finding in the categorization literature that category members differ in their goodness of membership to the category, e.g. a Toyota Camry is a better example of a car than a Ferrari (even if a Ferrari is a better car). Often, the average member of a category is judged to be the most typical or "best" exemplar of a category. Indeed, important data produced to evaluate early prototype and exemplar models was the finding that central category members are deemed more typical than peripheral category members (e.g., Medin & Schaffer, 1978; Nosofsky, 1986; Posner & Keele, 1968).

Central tendencies are clearly important for determining the typicality of exemplars of feature-based categories. However, research on ad hoc, and goal-derived categories (Barsalou 1983, 1985) and studies of categorization by experts (Lynch, Coley & Medin, 2000) demonstrate that ideal category members can influence people's beliefs about goodness of membership as well. Ideal members have extreme values on the relevant dimensions. For example, the goal-derived category "diet food" has an ideal value of 0 along the calorie dimension. If one were to average all the foods classified as diet foods, not only would the average be much larger than 0, there may not be a single exemplar that has that value, and yet the goodness of membership to the category is based on distance from the ideal, and not the central tendency. Lynch et al. (2000) found that ideal category members also influenced judgments of tree experts who have goals related to their interactions with trees (e.g., arborists).

Goldwater et al. suggested that role-governed categories are like goal-derived categories because they have prominent ideals in their representations. Like all categories, exemplars vary in their goodness of membership. In role-governed categories, category members differ in the degree to which they fulfill the role in the relational structure. Because the ability to serve the relational role is crucial, goodness of category membership should be related to this ideal rather than to the average values of the particular items that happen to be part of that category.

This study assessed the role of typical and ideal category properties as predictors of category goodness. The measure used for category goodness was the utility of category properties for explaining the category to someone else. Explanatory theories are crucial parts of the information people use to classify items (Ahn, Kim, Lassaline, & Dennis, 2000; Keil, 1989; Murphy & Medin, 1985; Rehder, 2003). If ideals are more important for role-governed categories than for feature-based categories, then properties that pick out ideal category values should have a more prominent role for people trying to explain the category to someone else than should properties that pick out typical category values.

To test this hypothesis, some participants were asked to list properties of ideal members of the same feature-based and role-governed categories for which the other participants had listed typical properties. Goldwater et al. constructed lists of the five most frequently listed ideal properties and the five most frequently listed typical properties. Properties that were listed frequently for both ideal and typical examples were not included. A new group of participants was shown the ideal and typical property lists and were asked which list they would use to explain the category to someone with no knowledge of it. Confirming their predictions, subjects chose ideal property lists more frequently for role-governed categories than for feature-based. This demonstrates another dimension upon which role-governed and feature-based categories differ, the prominence of ideals.

2.2 Inducing role-governed categories from relational categories

Goldwater et al.'s first three studies empirically distinguished role-governed from feature-based categories, their fourth investigated the hypothesis that role-governed categories emerge from relational categories. This view proposes that instantiating novel relational categories licenses novel role-governed categories. A side benefit of this study is that it suggests role-governed categories are straight-forward to acquire. Goldwater et al.'s fourth study investigated these two hypotheses by examining how role-governed categories might be learned in the context of learning new verbs.

Their participants instantiated relational structures by interpreting novel denominal verbs (i.e. novel verbs derived from nouns). Many familiar verbs are denominal, e.g. *dust*, *shelve*, *saddle*, *google*, etc (see Clark and Clark, 1979 for a comprehensive taxonomy).

Interpreting a novel denominal verb creates a relational representation from a feature-based category. While this is certainly a complex process, *novel* denominal verbs are readily understood in their first encounter in on-line sentence comprehension (Goldwater & Markman 2009, see Kaschak & Glenberg, 2000 for an off-line comprehension task) suggesting they are a reasonable object of study for the instantiation of novel semantic representations.

The use of novel denominal verbs also allows the instantiation of relational representations that are rooted in pre-existing knowledge, allowing concept learning to be rapid and to be embedded in the simple reading of short passages. For example, consider the sentence: "On the first night of Mardi Gras, Paul *whiskied* himself stupid." We can test whether this instantiation licenses a novel role-governed category by using the –er morpheme in English to create a novel agent term. So, later in a passage containing the sentence about Paul, we could refer to him as *the whiskier*. If people are able to form novel role-governed concepts when they instantiate a relational structure, then being exposed to a novel denominal verb should immediately allow them to understand the role-governed category.

The control passages were identical except for the sentence introducing the novel verb. The first control uses a paraphrase of the sentence with the novel denominal verb, but without the lexical innovation, e.g. "On the first night of Mardi Gras, Paul used *whiskey* to drink himself stupid." The second control condition used a novel adjective derived from the same root noun, e.g. "On the first night of Mardi Gras, Paul had a *whiskeylicious* time." This second control condition was included to rule out a "general novelty effect." One could argue that the novel agent term is read faster after the novel verb just because once one has read a novel word derived from a root word, it is then easier to understand any other novel word derived from that same root word. If the novel denominal verb condition has an advantage over the novel adjective condition, then this general novelty effect cannot explain the results in entirety, and instead there is support for the special link between verbs and agents.

In all conditions, because the event of Paul drinking a lot of whiskey is understood, the term *whiskier* is understandable, but because of the special link between verbs and agents, between relational and role-governed categories, understanding the same novel agent term after the novel denominal verb should be easiest. The final sentence of the passages for each condition were identical "The next day, the *whiskier* slept in until 2:00." Goldwater et al. found that processing time for the novel agent term was faster when it was preceded by the novel verb than by either control, confirming the hypothesized relation between relational and role-governed categories.

Goldwater et al.'s reading time effect is strong evidence for the connection between relational and role-governed categories, but it does not shed much light on the nature of the novel agent licensing process. Goldwater et al. (2009) hypothesized that the licensing process is run by the implicit semantic system because interpreting novel word uses in context can seem so effortless. However, it is possible that more explicit reasoning processes are involved. A third possibility is that the effect is simply one of morpho-syntax because agent terms marked by "-er" are derived from verbs. To extend the reading time study, Goldwater et al. replicated it (without the adjective condition, as the general novelty effect was already ruled out) using an Event-Related Potential paradigm because different ERP components are markers of these different processing systems. A relative difference in the N400 component is a sign of implicit semantic memory processes, while a difference in the Late Positive Component (LPC), aka the P600, is a marker of explicit decision processes and morphosyntactic processing (Heinze, Muente, & Kutas, 1998; Kim & Osterhout, 2005; Voss & Paller, 2005). The novel agent term showed a reduced N400 when preceded by the novel verb than when preceded by the paraphrase, and there were no reliable differences in the LPC. Thus, Goldwater et al. concluded that licensing novel role-governed category terms from interpreting novel relational terms is an implicit semantic process.

2.3 Conclusion

There were two primary goals of this reviewed work. One was to develop empirical markers distinguishing role-governed categories from feature-based categories. The two categories were distinguished by: 1) The larger proportion of extrinsic properties 2) Smaller flickr tagged proportion 3) Greater explanatory value of ideals for role-governed categories as compared to feature-based. The second goal was to show that novel relational structures license novel role-governed categories. These studies confirmed this hypothesis and demonstrated that the licensing process is run by the implicit semantic system. I now discuss the next set of questions and present the principal studies of this dissertation.

3. Dissertation Experiments

My framework claims that role-governed categories are formed by categorizing entities together by their common relational role. Much of my previous work examined familiar role-governed categories, and so did not reveal how these categories were initially formed. The goal of my dissertation experiments is to test this claim that entities are categorized by common role directly, i.e., by presenting participants with entities that share a relational role, and eliciting judgments of categorization.

Having participants make explicit categorization judgments allows for a number of research goals to be achieved. First, these experiments provide critical evidence for the framework laid out in this dissertation. Other models of categorization do not make predictions about whether entities will be categorized by common roles, whereas it may be my model's most straightforward prediction. Second, the experiments help to further specify the presented framework by revealing the processes involved in making such categorization judgments.

Categorizing two entities by common relational role requires first representing each entity by its relational role in its domains, and then representing the relational commonality across domains. Retrieving relational connections across domains is far from trivial (Gick & Holyoak, 1981), yet it is a prerequisite for role-governed categorization. Thus, it is quite possible that without something to make participants shift their focus to the across-domain common relational structure, they will only represent within-domain connections. The proposed experiments examine ways to induce such shifts in attention.

The experiments examine two processes that may lead to increased across- domain connections, and hence more role-governed categorization judgments. Experiment 1 investigates the extension of category labels. Experiments 2 -4 examines the structural alignment of relational representations (Gentner & Markman, 1997). I now discuss each in turn.

The first process I examine is the how category labels affect role-governed categorization. Labels have a very important role in category representation, as they have been shown to attribute essences (Gelman & Heyman, 1999) and generate category coherence (Yamauchi, 2005). Evaluating how labels direct attention towards common roles will connect role-governed categorization research with the field of categorization more broadly. In addition, it may potentially alter how the field views the role of labels in categorization because the effects of labels are thought to be privileged for feature-based categories (e.g., E. Markman, 1989. Role-governed categories were not considered, as the idea was yet to be conceived).

Labels have been shown to focus children on feature-based categories, and divert their attention from thematic relations. Markman (1989) showed that young children will show a preference for thematic relation categories when they are asked to make categorization judgments without a category label. For example, when children are queried in the form of “what goes best with this picture?” and they are shown a picture of a robin, they will choose a picture that forms a thematic relation, e.g. a picture of a nest, over a picture of another bird. However, if the robin is given a novel label e.g., a *goppin,* and then children are asked which other object is also a *goppin*, they will choose the object with greater feature-based similarity instead of the thematic relation match. Natural thematic-relation categories do not have labels, whereas natural feature-based categories do. This label advantage for feature-based categories is interpreted to show the importance of feature-based categories in mature conceptual systems.

Natural role-governed categories also have labels. If labels promote role-governed categorization over thematic relation-based, then at least for adults, role-governed categories should be seen to share this importance in the conceptual system.

However, the extension of novel labels across domains does not necessarily shed light on the mechanism behind finding the across domain relational correspondences. One potential mechanism is *structural alignment*, the mechanism underlying analogical comparison. When forming an analogy, people seek similarities in the relations that describe a pair of domains, even if the objects that take part in those relations are not identical. Thus, the atom was once seen to be similar to the solar system, because it was hypothesized that there is one element that revolves around another in each, even though electrons do not look like planets, and nuclei do not look like the sun. Structural alignment explains this perceived similarity through a process that aligns relational representations by putting entities into correspondence by role, ignoring featural dissimilarities.

Markman and Gentner (1993) showed that comparison induces structural alignment. The task was to indicate which character in one scene corresponded to a character in a second scene. For example, one scene would depict a woman giving food to a squirrel, whereas the second scene would depict a pizza delivery guy giving food to a woman. When prompted to select which character in the second scene corresponded with the woman from the first scene, participants would pick the second woman. However, if the subjects compared the two scenes (via a similarity rating) before making their selection, then they overwhelmingly dismissed the feature-based commonality between the two women, and put the first woman in correspondence with the delivery guy, as both were playing the *food giver* role.

While putting objects into correspondence seems like a good basis for categorizing them together, it may not be enough if people do not think sharing a relational role is a good basis for a category. However, if structural alignment is a mechanism of role-governed categorization, then rating two entities similarity, e.g., *bodyguard* and *forcefield*, should increase the likelihood they are judged to form a category.

Experiments 2 -4 do not just investigate whether similarity ratings increase role-governed categorization, but how making those comparisons support it. They contrast two possibilities. 1. The similarity rating induces structural alignment of the particular elements in the comparison, affecting the representation of each comparison independently. 2. Similarity ratings prime a general mode of role-governed responding, such that the effect of making similarity ratings on one set of entities transfer to sets of entities that were not directly compared.

In sum, the dissertation experiments will primarily examine two potential processes of categorizing entities by their common relational role. The use of novel category labels will potentially extend the "label privilege" of feature-based categories to role-governed categories. Having participants rate entities' similarity before making category judgments will reveal whether structural alignment is a mechanism of role-governed categorization. Now, I turn to describing the experimental methods, pilot studies, and stimuli norming.

3.1 Methods and Pilot Study 1

The primary method these studies use to elicit categorization judgments is the classic triad categorization task, where there is a target word or picture, and the subject is asked which of two alternatives goes better with the target to form a category. A recent relevant example is Lin and Murphy (2001) who used this task to contrast thematic relation with feature-based categorization. Participants were presented with a target word, e.g., *cat*, and asked to indicate whether an entity that formed a thematic relation with the target, e.g., *litter box*, or a feature-based match, e.g., *lion*, better made a category with the target.

Two pilot studies use this method to establish baseline levels of role-governed categorization. After describing these pilot studies, I describe a series of experiments that move the rates of such categorizations above baseline levels, and in doing so reveal underlying processes.

The first question to ask is: will people choose to categorize objects together based on common relational role at all? Thus, in this first pilot study, the two alternates with which to form a category with the target was a concept that shared the target's relational role (the role match) and a relatively unrelated term. For example, *bodyguard*'s role match is *force-field* and its unrelated alternate was *scissors* (see Table 2 for complete set).

Methods

Participants

Thirty University of Texas undergrads participated in this experiment for course-credit.

Materials

The stimuli are listed in Table 2 (only targets, role-matches, and unrelated alternates used for this experiment).

Procedure

Thirty participants were tested individually, at computers using the experimental software *E-Prime*. The instructions were modified from Lin and Murphy (2001). Participants were told:

In this experiment we are interested in how you think about everyday things in the world.

Specifically, we are interested in what things you think go together to make categories.
A category is a set of things or people that share some commonalities—be it genetic makeup, functions, physical or perceptual characteristics, purposes, or behavioral predispositions.
Every trial in this study will present 3 words. One will be at the top of the screen. The other two will be below, and marked by an "A or B". Your job is to choose which one of the two lower words goes best with the higher word to form a category by pressing A or B.

Every trial presented the three words as described, with the target word put in the sentence "Which better goes with X to form a category?" On every trial participants selected A or B. On half of the trials the role match was choice A, and half choice B. Their selections and response times were recorded. There were 20 trials. This took approximately 5 minutes.

Results and Discussion

There was a strong preference for the role matches, as they were chosen at an average proportion of .91, while the unrelated alternate was chosen at a rate of .09.

Participants were clearly able to make affirmative categorization judgments based on common relational role. This was crucial to establish. However, the unrelated alternates were not strong foils. Because role selections were essentially at ceiling, the results do not provide a window into the mechanism underlying such judgments. To reveal such a mechanism, one would need to perform a manipulation that changes the rate of role selections. Thus, the goal of this second study is to demonstrate a scenario where role selections are much lower, so they will then have room to be increased by further experimental manipulations.

3.2 Pilot Study 2

The second pilot study uses the same methods to contrast role-governed categorization with thematic relation categorization. The former entails categorizing objects by their common role across situations, while the latter entails categorizing objects together that play corresponding roles in the same situation. For example, *bodyguard*'s thematic match is *celebrity*. Thematic relation categorization should be preferred because the connection is more salient; it entails activating one situation schema where all the elements are directly bound to each other, while role-governed categorization entails activating two relational representations and finding connections across them. Indeed, thematic relation categories were shown to even be preferable to feature-based matches in many contexts (Lin & Murphy, 2001).

Methods

Participants

Thirty University of Texas undergrads participated in this experiment for course-credit.

Materials

The stimuli are listed in Table 2. Here, thematic matches are used in place of the unrelated alternate.

Procedure

This procedure is identical to that from the first study.

Results and Discussion

In contrast to the 1st study, participants showed a strong bias away from the role matches. They selected the role matches, at an average rate of .2, while selecting the thematic match at a rate of .8.

In addition, response time patterns were suggestive. I claimed that thematic relation connections are more easily accessed as only a single situation schema needs to be activated, while multiple relational representations needs to be activated and compared in role-governed categorization. While these response times need to be interpreted with caution, they are consistent with that claim. Role selections (M = 5.163 sec) were reliably slower than thematic relation selections (M = 3.335 sec), $t\,(24) = 3.55$, $p < .01$ by subjects (not including subjects who did not select a single role match), and by items, $t\,(19) = 2.69$, $p < .05$. In addition, the number of role selections made was predictive of an increase in average response time for all selections, showing a trend by subjects, $r\,(29) = .31$, $p < .1$ and $r\,(19) = .47$, $p < .05$ by items.

Participants preferred the obvious associations rooted in thematic relations to form categories in this second study. This was the goal. Now, a baseline is established. From here we can examine what will focus subjects on across situation connections, and de-emphasize thematic dyads.

3.3 Stimulus Norming

In addition to the pilot studies, these stimuli were normed, ensuring empirically that the stimuli are implementing what I am claiming, that is, that the role matches are really role matches. Some of the role matches could also be cross-classified as members of a common feature-based category e.g., patients and customers are both people. In these cases, I chose an unrelated alternate also cross-classifiable in this same category, e.g., *blacksmith*. However, these intuitions need to be confirmed with data. Subjects listed things that the target words have in common with all three kinds of alternative choices from the pilot studies. These were coded as intrinsic or extrinsic. Two findings are critical. 1. Role and thematic matches elicit primarily extrinsic commonalities, confirming that judgments about shared category membership with the target is not rooted in common intrinsic features. 2. Role matches must not have more intrinsic commonalities with the targets than either of the other two. The results are shown in Table 3.

For all three kinds of matches, extrinsic commonalities were listed more frequently than intrinsic commonalities. In addition, all three alternate types had virtually identical rates of listing intrinsic properties. However, they differed in the number of extrinsic properties, as revealed by a one-way matched samples ANOVA $f(2, 38) = 7.89$, $p < .01$, by items. The Tukey HSD test showed that the role matches elicited more extrinsic commonalties than the unrelated matches, but no other comparison was significant. Importantly, these results confirm the suitability of these materials. The role match has no reliable advantage in terms of *number* of properties in common with the thematic match, and no advantage in number of intrinsic properties over the other alternates. Finally, in a qualitative analysis, suggesting our intuitions were correct about role matches being role matches, many of subjects explicitly listed their relational commonality. For example, for *coffee* and *Red Bull*, one subject reported "they both give you energy."

3.4 Experiment 1

Category labels are more than just another property of the category: category labels are thought to attribute essences (Gelman & Heyman, 1999), and generate coherence (Yamauchi, 2005). Proposed experiment 2 examines whether novel category labels focus subjects on role-governed categorization.

Labels have been shown to focus children on feature-based categories, and divert their attention for thematic relations in the triad categorization task (E. Markman, 1989, and see above). Natural thematic-relation categories do not have labels, whereas natural feature-based categories do. This label advantage for feature-based categories is interpreted to show the importance of feature-based categories in mature conceptual systems. Natural role-governed categories also have labels. If labels promote role-governed categorization over thematic relation-based, then at least for adults, role-governed categories should be seen to share this importance in the conceptual system. Experiment 1 examines whether presenting the categorization judgment in the form of a novel word extension task will promote role choices.

This experiment has two conditions using the materials from pilot study 2. The first condition examines the extension of category labels. The target is labeled a *goppin* and then the question is asked which alternate is better also called *goppin*. The second condition examines the extension of properties across items, in the form of a novel adjective. This will function as a control, ensuring that any advantage for role choices is not just about novel word extension, but particularly about category labels. For example, in the category label condition, the query reads: "That bodyguard is a *goppin.* Which of these other two is better called *goppin*." In the novel property condition: "That's a *goppin* bodyguard. Which of these other two is better called *goppin*." The first condition should boost role choices, as role-governed categories, and not thematic relation categories have labels. In contrast, the same advantage may not be present in the property condition because people often extend properties across members of thematic relation categories (Lin & Murphy, 2001).

Methods

Participants

Twenty-seven University of Texas undergrads participated for course credit.

Materials

The materials are those presented in Table 2, using the role and thematic matches. Instead of the question asked being about category membership as in the pilot study, the query concerned novel word extension. There were two novel words used, *goppin* and *blicket*. One condition used one nonce word, and the other condition used the other nonce word (counterbalanced across subjects which word went with which condition). For example, in the category label condition, the query reads: "That bodyguard is a *goppin.* Which of these other two is better called *goppin*." In the novel property condition: "That's a *blicket* bodyguard. Which of these other two is better called *blicket*."

Procedure

The procedure was the same as from the pilot studies, except the instructions were altered to explain about the use of the nonce words. They read:

In this experiment we are interested in how you think about everyday things in the world. Specifically, we are interested in what things should be described or named with the same words. In this experiment, you will read statements about everyday things that are described or named with a new word. Your job is to indicate what else should be named or described with this same new word.

There is one new word used for all the new names, as well as a single new word for the descriptions. But treat each presentation of the new word separately. Meaning, when thinking about the new word, and what else should be named or described with it, do not consider what this new word has already been used to name or describe for your previous choice.

Every trial presented the three critical words, the target, and the role and thematic matches, with the target word put in the sentence "That X is a goppin?" or "That X is a blicket" for the novel noun condition, or "That's a goppin X" or "That's a blicket X" for the novel adjective condition. Below the target sentence, a sentence reads "Which of these other two is better called goppin (or blicket)?" Then, below are the two alternates, marked by an A or B. On half of the trials the role match was choice A, and half choice B. Their selections and response times were recorded. There were 20 trials. This took approximately 5 minutes.

Results and Discussion

If people think role-governed categories are more worthy of labels than thematic relation categories, then the label condition should produce more role-choices than the property condition. This prediction was confirmed, with the label condition eliciting a rate of role choices of .66, while the property condition elicited a .38 proportion, t (26) = 4.46, $p < .001$ by subjects, t (19) = 9.53, $p < .00001$, by items. For clarification, unlike the pilot studies where the statistical tests compare thematic to role choices without a separate experimental manipulation, these statistical tests are comparing the proportion of role choices across two conditions.

These results suggest that the category label privilege first conceived in regards to feature-based categories should be extended to role-governed categories. Also suggestive is the proportion of role choices in the property condition. The baseline established from Pilot study 2 was 20% role choices. It was almost twice this in the property condition. This suggests that people make property inductions across members of role-governed categories. However, this does not tell us anything specific about these inductions. This is an important topic of future research (see below).

3.5 Experiment 2

The second experiment examines a potential mechanism for role-governed categorization. The candidate mechanism is "structural alignment," the mechanism underlying analogical comparison (Gentner, 1983; Gentner & Markman, 1997). As explained above, when forming an analogy, people seek similarities in the relations that describe a pair of domains, even if the objects that take part in those relations are not identical. Structural alignment explains this perceived similarity through a process that aligns relational representations by putting entities into correspondence by role, ignoring featural dissimilarities.

Intuitively, placing entities into correspondence, and increasing their perceived similarity seems like a good basis for categorization. However, this hypothesis has not been directly tested, because research on analogy and categorization has been somewhat disconnected. This experiment promotes structural alignment in subjects by having them rate the alternative choices' similarity with the target. Similarity ratings have been shown to promote role-based correspondences over feature-based (Markman & Gentner, 1993). This experiment tests whether similarity ratings promotes role-governed categorization.

Experiment 2 has two conditions. In the first, subjects rate the similarity of the target with both the role and thematic relation match. Then the triad categorization judgment will be queried. In the second condition, subjects compare the relative imageability of the pairs of words. This will ensure that any advantage for role-matches in the first condition is not simply due to making any comparative rating. If structural alignment is a mechanism of role-governed categorization, then there should be more role-choices in the similarity condition.

Methods

Participants

Thirty-nine University of Texas undergrads participated for course credit.

Materials

The materials used are identical to those from Pilot Study 2, with additional judgments made on the sets of words. For similarity ratings, two words were presented, the target and one of the alternates. The query read: how similar are [Target] and [Alternate]? The target and alternate were presented in large font on their own line. There was a 7 point Likert scale, that ranged from Not Similar to Very Similar. For relative imageability ratings, two words were presented, the target and one of the alternates. The query read: Which is easier to picture in your head: A [Target] or B [Alternate]? There was a 7 point Likert scale, that ranged from A is easier at 1 to Both are equal at 3 and B is easier at 7.

Half of the trials had similarity ratings, and half had imageability ratings. The order was randomized. Which items were rated for similarity and which for imageability was counterbalanced across subjects.

Procedure

The procedure was identical to Pilot study 2, with the addition of the other judgments and modified instructions, which read (the description for the category judgment trial was identical to the description from the pilot studies, so is not included):

In this experiment we are interested in how you think about everyday things in the world. Specifically, we are interested in:

1. What things you find similar.
2. What things are easy to picture in your head.
3. What things you think go together to make categories.

There are three kinds of trials. One kind for each of the three questions we are interested in.

In the first kind of trial, you will judge how similar you find two words to be. They will be presented in the middle of the screen. Below, there will be a scale from 1 - 7. 1 means that you see no similarity between the two words. 7 means that you see very high similarity between the two words. Please press the number that best matches how similar you find the two words, from 1-7.

In the second kind of trial, you will judge how easy it is to picture one word in your head relative to another. Each word will be presented in the middle of the screen, and will be marked by an 'A' or a 'B.' Below, there will be a scale from 1 - 7. 1 means that you find word A much easier to picture.7 means that you find word B much easier to picture. 4 means that you find them both equal. Please indicate which one you find easier to picture by pressing a number, 1 - 7.

Each trial consisted of three judgments. The first two judgments were of the same type. Both of the two judgments were either similarity ratings comparing the target with each of the two alternates or both were relative imageability ratings, also comparing the target with each of the two alternates. Half of the trials presented the judgment comparing the target to the role match first, and half the thematic match first. The third judgment was the triad categorization judgment identical to the pilot studies.

Results and Discussion

If structural alignment is a mechanism of role-governed categorization, then the category judgments preceded by the similarity comparisons should elicit a higher rate of role choices than when preceded by the imageability ratings. The predicted pattern was found, with the similarity condition eliciting a .3 proportion of role choices, as compared with a .24 proportion elicited by the imageability rating. However, this was merely a trend by subjects, t (38) = 1.73, $p < .1$, and reliable by items, t (19) = 2.18, $p < .05$.

While the results went in the predicted direction, the effect, in terms of raw scores, was not very large. The similarity condition only elicited 6% more role choices than the imageability condition (compare that with Experiment 1, where the category condition elicited an increase of 28% compared to the property condition). There are a few possibilities to explain the small effect. 1. Structural alignment is not a large part of role-governed categorization. 2. Structural alignment is an important process of role-governed categorization, but there was some influence of switching between judgment types on successive trials. Experiment 3 attempts to tease these two explanations apart.

3.6. Experiment 3

The goal of Experiment 3 is to demonstrate that structural alignment is a mechanism of role-governed categorization. Structural alignment is induced through similarity ratings, and to reduce the potential influence of making other kinds of comparison in Experiment 3, the two kinds of judgments, similarity and imageability, was manipulated between subjects. Otherwise, it is identical to Experiment 2.

Methods

Participants

Thirty University of Texas undergrads participated for course credit.

Materials

The stimuli were identical to those from Experiment 2.

Procedure

The procedure was identical to that from Experiment 2, except the comparison manipulation was between subjects. The instructions were modified accordingly.

Results and Discussion

If structural alignment is a mechanism of role-governed categorization, then the category judgments preceded by the similarity comparisons should elicit a higher rate of role choices than when preceded by the imageability ratings. The predicted pattern was found, with the similarity condition eliciting a .41 proportion of role choices, as compared with a .21 proportion elicited by the imageability rating. This was reliable both by subjects, t (29) = 2.24, $p < .05$, and by items, t (19) = 4.18, $p < .001$.

Experiment 2 showed that rating two items' similarity increased their likelihood of being categorized together when the two items play a common role across domains, but the effect was relatively small when similarity judgments were intermixed with relative imageability ratings. Experiment 3, with half of the subjects making similarity ratings on all trials and half making relative imageability ratings, the power of the manipulation increased. Here, the similarity condition elicited nearly twice as many role choices as the imageability condition. This suggests that the small effect seen in experiment 2 was due to the influence of switching judgment type from trial to trial.

There are two explanations of how similarity ratings may be affecting role-governed categorization consistent with the pattern across these two experiments. 1. The similarity rating induces structural alignment of the particular elements in the comparison, affecting the representation of each comparison independently. 2. Similarity ratings prime a general mode of role-governed responding. Because these two accounts cannot be distinguished on the basis of the current pattern, Experiment 4 attempts to tease the two apart.

3.7 Experiment 4

Similarity ratings may be inducing structural alignment on a per-trial basis, or it might be putting subjects into a general mode of role-governed responding. These explanations differ because the latter predicts that this general mode, once induced should be transferable to trials that do not directly induce structural alignment via a similarity rating. While it is clear from Experiment 2 that the role-governed mode of responding does not transfer from trial to trial when similarity ratings and imageability ratings are intermixed, this does not rule out the possibility that if subjects made an entire block of only similarity ratings, whether this would solidify the role-governed mode of responding enough to then allow transfer to a block of trials where imageability ratings were being given. Experiment 4 tests this hypothesis by having two blocks of trials each entirely composed of trials of one judgment type and manipulates their order.

Methods

Participants

Thirty-seven University of Texas undergraduates participated for course-credit.

Materials

The materials were identical to that from Experiment 2, except there were 4 experimental scripts to counterbalance which items were rated for similarity or imageability, and whether similarity or imageability ratings were given first.

Procedure

The procedure was identical to that from Experiment 2, except the similarity and imageability trials were blocked, instead of intermixed. Order within-block was randomized.

Results and Discussion

If similarity ratings induce a general role-governed responding mode, then the imageability trials should elicit a higher rate of role-governed responding when following the block of similarity trials than when preceding it. However, this order effect must be greater for the imageability condition than the similarity condition to claim the effect is about a transfer of role-governed responding, and not just a more general order effect. That is, block order and judgment type must interact.

2 (judgment type: similarity, imageability) X 2 (order: similarity 1^{st}, imageability 1^{st}) mixed-measures ANOVA (by subjects) and repeated-measures ANOVA (by items) revealed main effects of judgment type, $f(1,35) = 22.4$, $p < .001$, by subjects, and $f(1, 19) = 23.7$, $p < .001$, by items, as the similarity condition ($M = .30$) elicited a higher proportion of role-governed responses than the imageability rating ($M = .12$). There was no main of effect of block order, $f(1, 35) = .145$, by subjects, and $f(1, 19) = 2.12$, by items p's > .15. Importantly, there was no interaction between block order and judgment type, $f(1, 35) = 1.07$, by subjects, and $f(1, 19) = .79$, by items p's > .3 (see Figure 4).

Thus, there is only evidence that similarity ratings induced structural alignment on a per-comparison basis, and not for a general role-governed responding mode.

The large main effect of judgment type replicates Experiment 3. However, here, it is a within subjects manipulation as in Experiment 2. This suggests that the original interpretation that the small effect found in Experiment 2 was due to the costs of switching between judgments types throughout the experiment was correct, as the current experiment only has a single judgment-type switch halfway through the trials.

3.8 Summary of Results

These four experiments advance our understanding of role-governed categorization, and make crucial evaluations of the model of conceptual representation described throughout the dissertation. Traditionally, feature-based categories are seen as privileged in mature conceptual systems. One sign of this advantage is that they have category labels, and even very young children take advantage of this to guide their word learning (E. Markman, 1989). Experiment 1 argues that role-governed categories should be seen as equally prominent, as role-governed categories have labels, and novel label use can direct one's attention to role based commonalities across entities, just as they do for featural.

Experiments 2 – 4 examine an underlying mechanism of role-governed categorization by connecting it with analogical reasoning through the shared process of structural alignment. The pattern across three experiments shows that aligning relational representations, that is, putting entities into correspondence by role, leads to increased role-governed categorization judgments. Additionally, aligning relational representations works on a per-comparison basis, and does not induce a general role-governed response mode.

4. General Discussion, Implications and Future Directions

Traditionally in Cognitive Psychology (and Western Philosophy more broadly), categories are thought to be represented as sets of features. This dissertation, following Markman & Stilwell (2001), looks to establish a new way to think about category representation that seeks to unify work on categorization and concept learning with other areas in higher-level cognition such as analogical reasoning. The framework emphasizes structured relational knowledge representations, and creates a taxonomy of categories based on different components of relational structures. Relational categories are categories of entire relational systems while, role-governed categories, the focus of this dissertation, are represented as the roles in the relational systems. Lastly, thematic relation categories group entities together that play complementary roles within these systems.

Four experiments were presented that support this framework and further specifies it. They contrasted thematic relation categorization with role-governed categorization. When the two are put in direct conflict, that is, when one has to choose which two items go better together to make a category, people prefer two that form a thematic relation category over two that form a role category. This is predicted because within-domain connections are easier to find than across-domain connections. However, when one is to extend a label across a pair of items, then role-based commonalities are preferred. This is predicted as natural role-governed categories have labels while thematic relation categories do not. In addition, 3 experiments demonstrated that structural alignment boosts the likelihood of preferring the role-based pair.

The first finding demonstrates the prominence of role-governed categories in our conceptual system and lexicon, and the second specifies an underlying mechanism that is shared with analogical reasoning. However, there are many open questions, and many more predictions this framework generates. Here, I relate the current findings to a few avenues of future research.

4.1 Learning Role-Governed Categories from Events

The current experiments looked at processes of categorizing entities across domain by common role. While the connections across domains may have been novel to the participants, all the individual entities were familiar, and the relational correspondences were found through the subject's pre-existing knowledge of the world. However, it seems very likely that many role-governed categories are first learned from observation of and interaction with the world in novel situations. For example, when children first learn about guests and hosts, it is probably through witnessing and participating in novel visiting events. That said, the current findings are relevant in that the framework predicts that the use of labels and structural alignment will be critical in forming such concepts from direct observation (see Gentner & Loewenstein, 2002).

Thus, one important line of research would be to examine how people learn event role categories from observing events. Participants will view a series of events, able to be categorized together based on their relational structure. After viewing such events, subjects would then be queried into what categories were formed. What would cause people to form categories of participants across the events based on common role? Similar manipulations to the current studies will be done to test their generality, e.g., inducing the structural alignment of the event exemplars. Also, completely novel questions would be asked in such a paradigm. One interesting question is to what degree roles are learned as systems or as independent. For example, can one learn about guests without also learning about hosts? Goldwater et al. (2009) show that instantiating relations licenses role-governed categories. Would learning event role categories be dependent on first learning the event categories? Once the event categories are learned, are all event roles licensed as a system?

Research investigating how role-governed categories are learned from viewing events is already underway (in collaboration with Cathy Echols). This project is examining how children learn that a noun refers to an event role and not a feature-based category. One process that may be of critical importance to children in learning about roles is turn-taking (Tomasello, personal communication). That is, seeing the same entities taking turns playing different event roles may be critical to abstracting the role away from the features of the particular participants who are playing it.

4.2 Labels, Induction and Category Coherence

An important function of categories is supporting induction about objects and events. Labels increase inductive strength. Category labels are thought to assign category essences to individuals. That is, they increase the inductive strength of attributing category essential properties to individuals, e.g., labeling someone a “feminist” increases the likelihood that that person believes in equal rights for women (Gelman & Heyman, 1999; Yamauchi & Markman, 1998, 2000). Yamauchi (2005) extended these findings by showing labels increase inductive strength for even unlikely characteristics of a person, e.g., labeling someone a “feminist” increases the inductive strength that she is a bank teller. Yamauchi hypothesizes that category labels generate coherence for categories by connecting disparate pieces of information together.

Experiment 1 showed that people think role-governed categories should have labels, as they do for feature-based categories. However, many role-governed categories are defined by a situation, that is, they are contextually bound. People are guests, hosts, and passengers for limited amounts of time. A fundamental question about situational role-governed categories is the degree to which their labels generate coherence as other types of “context-free” categories do. It would seem somewhat strange if labeling someone a passenger leads to any inductions about permanent properties about that person. Yet, we have rich knowledge about such categories that seems to cohere.

Alternatively, they may generate coherence not by tying together disparate pieces of information concerning permanent properties of category members, but through properties of the situation in which the category members are playing a role. For example, the role category of car passenger may only increase attribution for behaviors or properties that exist for the duration of the event, e.g., listening to the radio. I am currently testing this hypothesis.

4.3 Analogy and Conceptual Change

Experiments 2-4 demonstrate that role-governed categorization and analogy share an underlying mechanism. What is unclear is how coupled the two processes are. In the most extreme case, every analogy results in new role concepts. Perhaps, all one needs to do after learning about the Rutherford model of the atom to have a role concept that includes atomic nuclei, the sun, and any other object in the center of an orbiting system, is just to attach a label to it. This perspective that analogy and role concept learning are tightly coupled will generate many new research ideas.

Analogies are central to conceptual change, both throughout history (Arabatziz, 2009) and in individuals (Vosniadou, 2009). Arabatziz (2009) examined the use of analogies in early 20th century physics to move away from the Rutherford model to the Bohr model. Vosniadou (2009) showed the importance of analogy in teaching young children the proper mental model of the day/night cycle. Many young children have a mental model of day and night that represents a motionless Earth, and a sun and moon that alternate between rising into the sky and hiding behind mountains. Explicitly stating to a group of Greek elementary school children that the Earth's rotation was a critical piece of the day/night cycle had no effect on their model. However, making an analogy to the rotating obelisk of gyro meat, a rich base domain for Greek children, dramatically changed their understanding of the domain.

Studies such as Vosniadou (2009) tend to look at understanding domains as wholes, that is, the form of the entire mental model. Something the current framework adds to such work is the suggestion to examine elements of these relational systems as concepts in their own right. That is, perhaps what it means to have a new mental model of a domain is to have a new system of role concepts. A main function of analogies is to generate inferences. In many cases, analogical inference may be akin to role-governed category based induction. Of course, this is rather speculative at this point, however, making further connections between analogy and concept learning should be a fruitful line of research. Establishing their connection in these simple experiments presented here is just a first step.

4.4 Expertise and Learning Relational and Role-Governed Categories

While all people have relational and role-governed categories, as evidence by our lexicon, expertise is often defined as a shift away from feature-based categorization towards relational categorization. Chi, Feltovitch & Glaser (1981) showed that physics novices categorized problems based on their literal features, while experts used the physical relations that define the structure of the problems. This kind of expertise is presumably not learned through observation alone, as perhaps learning simpler relational categories may be.

While gaining expertise in physical-relational systems is the topic of much science education research, using the current framework leads to interesting new questions. For example, after physics experts learn to categorize problems by their abstract relational principles, do they form explicit concepts of the roles in these physical systems? If so, what function do these role concepts serve? What inferences do they help generate? Because role-governed categories are pieces of larger relational structures, perhaps role concepts are used as cues to the larger structures. Identifying a single role could allow one to activate an entire schema to aid in problem solving.

During my post-doctoral fellowship with Dedre Gentner, I will investigate relational and role-governed category learning in a science education context. Particularly, how people learn complex systems principles that explain climate change (in collaboration with Duncan Sibley, professor of Geo-Sciences).

4.5 Conclusion

This dissertation laid out a framework for category representation that has a rather broad scope. Category research, on the whole, has had a too narrow a focus for far too long (Markman & Ross, 2003). The work presented here establishes a ground-work, but the goal of this final chapter is to give a flavor of the large research program taking this view of concepts can generate. Categorization is a central cognitive process, and broadening the scope of category research will keep it a central part of Cognitive Psychology.

Table 1

Verbs and corresponding role governed categories

Verbs	Role Governed Categories
x steals y	x= thief
x visits y	x = guest, y = host
x trains/advises y	x = mentor, y = protégé
x gives birth to y	x = mother, y = son or daughter
x defeats y	x = victor/winner, y = loser
x eats y	y = food
x plays y	y = game
x writes y	x = author
x gives y to z	y = gift
x shops in y	x = customer
x lives in y	y = home

Table 2

Materials for Pilot Studies 1-2 and Experiments 1 - 4

Target	Role Match	Thematic Match	Unrelated Alternate
Appetizer	Opening Act	Entrée	Bed
Bird's Nest	House	Tree	Court
Bodyguard	Force-field	Celebrity	Scissors
Bouncer	Tollbooth	People in line outside a club	Remote Control
Coach	Manager	Athlete	Cop
Coffee	Red Bull	Cream	Lemonade
Customer	Patient	Salesman	Blacksmith
Director	Author	Actor	Crossing Guard
Drawer	Wallet	Socks	Statue
General	CEO	Private	Accountant
Hard Drive	Filing Cabinet	Electronic Files	Medicine
Jar	Closet	Pickles	Pillow
Lock	Security Guard	Key	Doctor
Paralegal	Nurse	Lawyer	Janitor
Police Chief	Admiral	Beat Cop	Athlete
Quilt	Campfire	Bed	Truck
Real Estate Agent	Drug Dealer	House-Buyer	Spy
Teacher	Camp Counselor	Student	Linebacker
Tripod	Cast and Sling	Camera	Book

Vacuum	Soap	Carpet	Paint

Table 3

Stimulus Norming Results: Mean and SD of Commonalities listed for Each Target

	Role Match	Thematic Match	Unrelated Alternate
Commonality	Mean (SD)	Mean (SD)	Mean (SD)
Extrinsic	2.43 (.51)	2.18 (.40)	1.83 (.56)
Intrinsic	0.70 (.38)	0.76 (.37)	0.73 (.44)

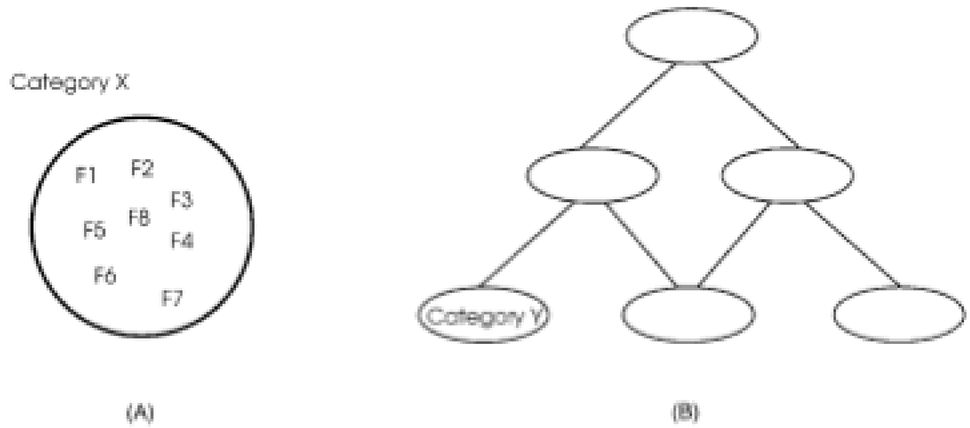

Figure 1: A) Schematic of feature-based category representation as subordinate to the category label B) Schematic of category structure superordinate to the role governed category label (taken from Markman & Stilwell, 2001).

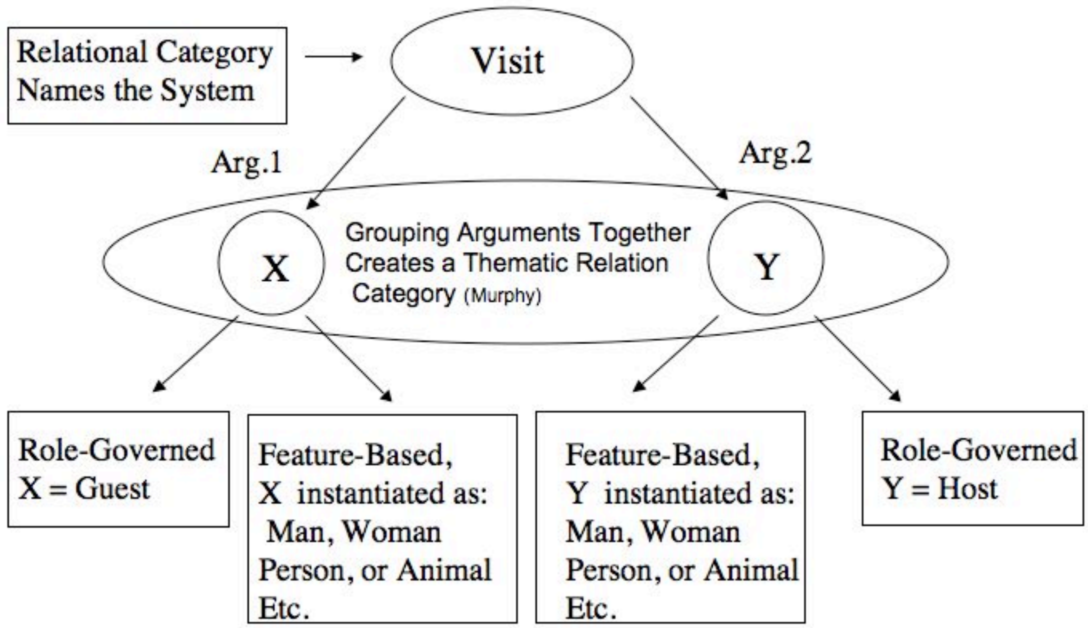

Figure 2: A relational system and its constituent categories. A relational category, e.g. *visit* names the system. Role-governed categories, *e.g. guest and host*, name individual arguments. Grouping the arguments together, creating a category of entities that includes guests and hosts, define a thematic relation category. In any given relational system, e.g. a visit, the roles are fulfilled with members of feature-based categories, e.g. men and women.

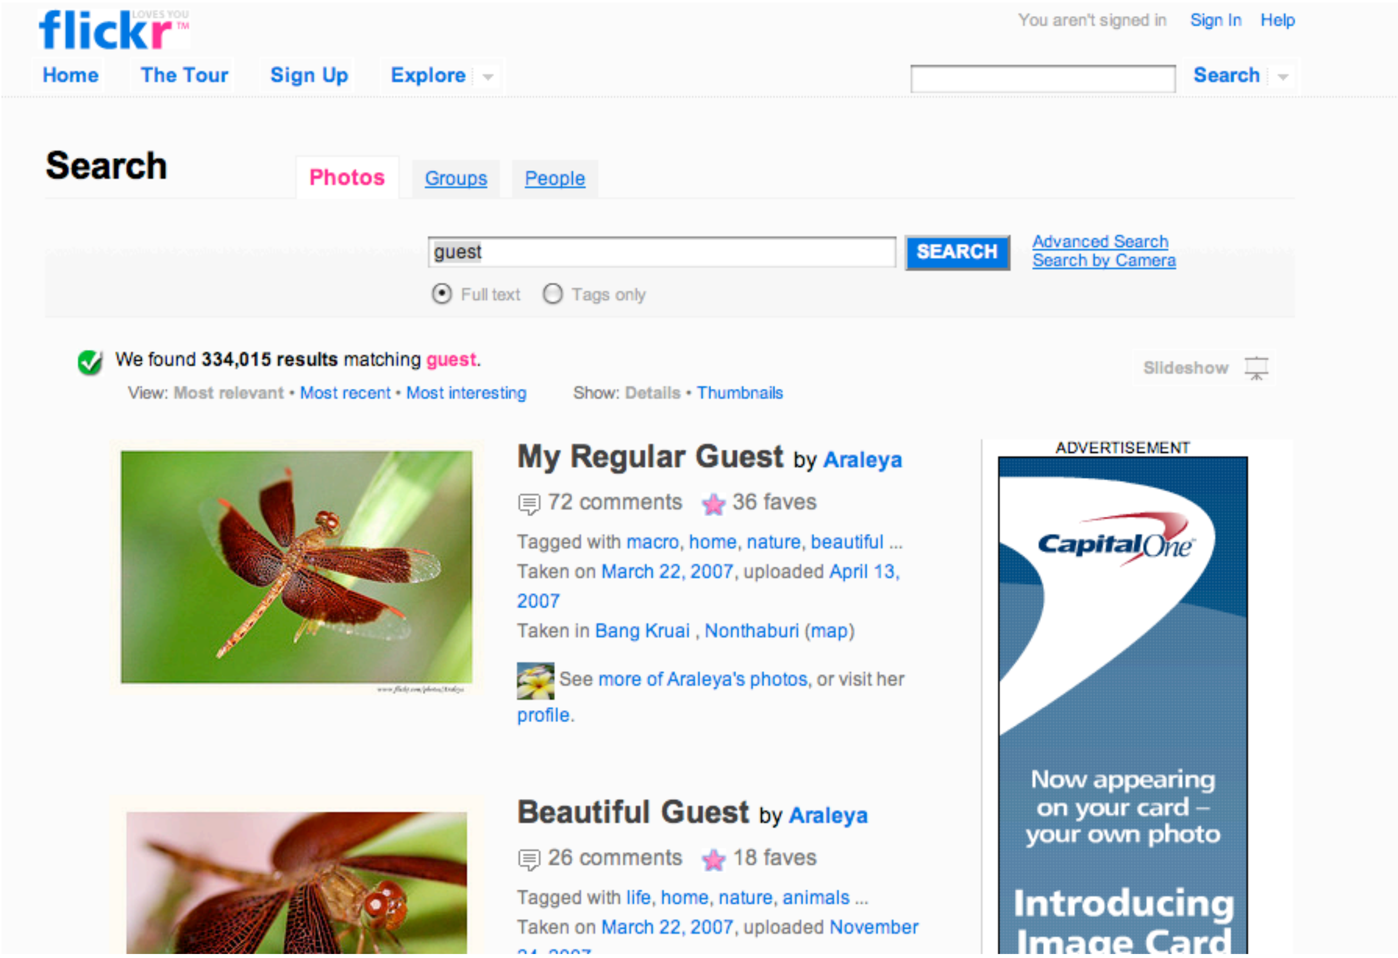

Figure 3: Screen shot of full text search for *guest*.

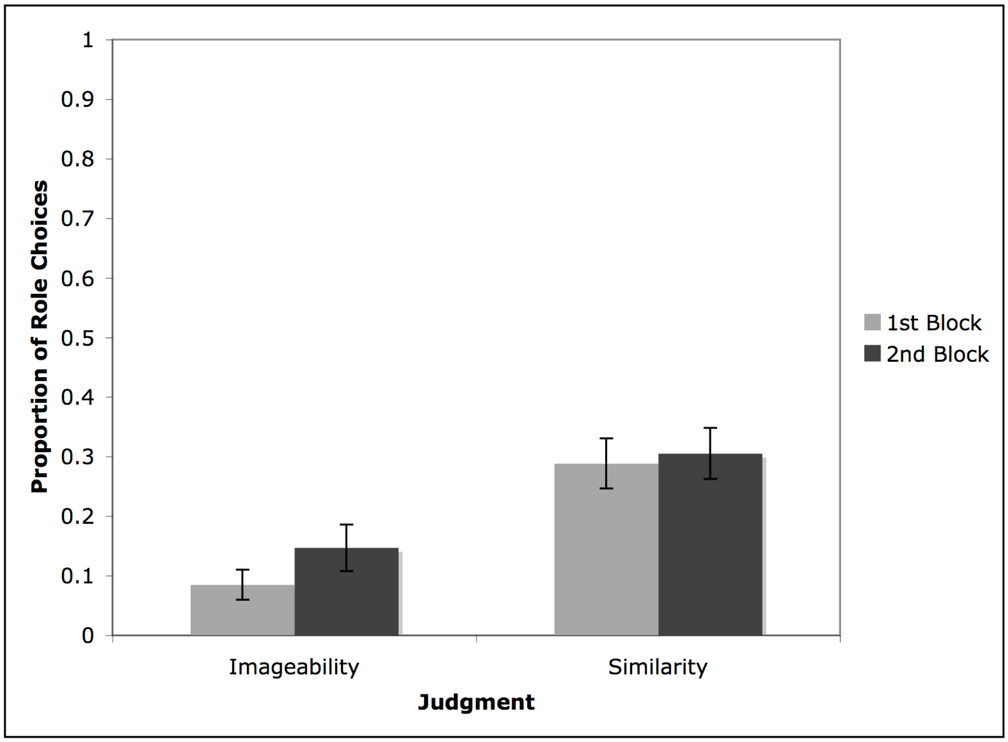

Figure 4. Items means and standard errors of proportion of role choices in Experiment 4

References

Ahn, W., Kim, N. S., Lassaline, M. E., & Dennis, M. (2000). Causal status as a determinant of feature centrality. *Cognitive Psychology*, *41*, 361-416.

Arabatziz, T. (2009) Models and Analogies in Conceptual Restructuring. *Symposium presented at the 31st Annual Meeting of the Cognitive Science Society.*

Barr, R. A., & Caplan, L. J. (1987). Category representations and their implications for category structure. *Memory and. Cognition, 15,* 397-418.

Barsalou, L. W (1983). Ad hoc categories. *Memory and Cognition, 11,* 211-227.

Barsalou, L. W (1985). Ideas, central tendency, and frequency of instantiation as determinants of graded structures in categories. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 11,* 629-654.

Barsalou, L. W. & Sewell, D. R, (1985) Contrasting the representation of scripts and categories. *Journal of Memory and Language, 24*, 646-665.

Casasola, M. and Wilbourn, M. P. (2004). 14-month-old infants form novel word-spatial relation associations. *Infancy, 6*(3), 385-386.

Clark, E. V., & Clark, H. H. (1979) When nouns surface as verbs. *Language, 55* (4), 767-811.

Chi, M. T. H., Feltovich, P., & Glaser, R. (1981). Categorization and representation of physics problems by experts and novices. *Cognitive Science, 5,* 121-152.

Cree, G.S., McNorgan, C., McRae, K. (2006). Distinctive features hold a privileged status in the computation of word meaning: Implications for theories of semantic memory. *Journal of Experimental Psychology: Learning, Memory & Cognition, 32*, 643-658.

Ferretti, T. R., McRae, K., & Hatherell, A. (2001). Integrating verbs, situation schemas, and thematic role concepts. *Journal of Memory and Language, 44,* 516-547

Fillmore, C. J. (1982). Frame semantics. In *Linguistics in the Morning Calm*. Seoul, Hanshin Publishing Co., 111-13

Gelman, S. A., and Heyman, G. D., (1999) Carrot-eaters and creature-believers: The effects of lexicalization on children's inferences about social categories. *Psychological Science, 10,* 489- 493.

Gentner, D. (1975). Evidence for the psychological reality of semantic components: The verbs of possession. In D. A. Norman, D. E. Rumelhart, & the LNR Research Group (Eds.), *Explorations in cognition* (pp. 211-246). San Francisco: Freeman.

Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. *Cognitive Science, 7*, 155-170.

Gentner, D., & Kurtz, K. (2005). Relational categories. In W. K. Ahn, R. L. Goldstone, B. C. Love, A. B. Markman & P. W. Wolff (Eds.), *Categorization inside and outside the lab.* Washington, DC: APA, p. 151-175.

Gentner, D., & Loewenstein, J. (2002). Relational language and relational thought. In J. Byrnes & E. Amsel (Eds.), *Language, Literacy, and Cognitive Development* (pp. 87-120). Mahwah, NJ: LEA.

Gentner, D., & Markman, A. B. (1997). Structure mapping in analogy and similarity. American Psychologist, 52(1), 45-56.

Gick, M. L., & Holyoak, K. J. (1981). Analogical problem solving. *Cognitive Psychology, 12,* 306-355.

Goldwater, M. B. & Markman, A. B, (*in press*). Constructional sources of implicit agents in sentence comprehension. *Cognitive Linguistics.*

Goldwater, M. B., Markman, A. B. & Stillwell, C. H. (2009) The empirical case for role-governed categories. *In preperation.*

Goldwater, M. B., Markman, A. B., Tujillo, L., & Schnyer, D. (2009). Licensing role-governed categories: An ERP analysis. *In preperation.*

Heinze, H. J., Muente, T. F., Kutas, M. (1998). Context effects in a category verification task as assessed by event-related brain potential (ERP) measure. *Biological Psychology, 47*, 121-135

Jackendoff, R. (1990). *Semantic Structures*. Cambridge, MA MIT Press.

Jackendoff, R. (2002). *Foundations of Language: Brain, Meaning, Grammar, Evolution.* Oxford, UK; Oxford University Press.

Kaschak, M. P., & Glenberg, A. M. (2000). Constructing meaning: The role of affordances and grammatical constructions in sentence comprehension. *Journal of Memory and Language, 43,* 508 –529.

Keil, F. C. (1989). Spiders in the web of belief: The tangled relations between concepts and theories. *Mind and language, 4*, 43-5

Kim, A., & Osterhout, L. (2005). The Independence of Combinatory Semantic Processing: Evidence from Event-Related Potentials. *Journal of Memory and Language, 52*(2), 205-225.

Landau, B., & Jackendoff, R. (1993). "What" and "where" in spatial language and spatial cognition. Behavioral and Brain Sciences, 16(2), 217-238, 255-265.

Leech, R., Mareschal, D. & Cooper, R. P. (2008): Analogy as relational priming: A developmental and computational perspective on the origins of a complex cognitive skill. *Behavioral and Brain Sciences, 31*, 357-378.

Levin, B. & Rappaport-Hovav, M. (2005) *Argument Realization*. Cambridge, UK; Cambridge University Press.

Lin, E. L., & Murphy, G. L. (2001). Thematic relations in adults' concepts. *Journal of Experimental Psychology: General, 130,* 3-28.

Lynch, E. B., Coley, J. D., & Medin, D. L. (2000). Tall is typical: Central tendency, ideal dimensions and graded category structure among tree experts and novices. *Memory and Cognition, 28*, 41-50.

Markman, A.B. (1999). *Knowledge Representation.* Mahwah, NJ: Lawrence Erlbaum Associates.

Markman, A.B., & Ross, B.H. (2003). Category use and category learning. *Psychological Bulletin, 129*, 592-615

Markman, A. B. & Stilwell, C. H. (2001). Role-governed categories. *Journal of Experimental & Theoretical Artificial Intelligence, 13*, 329 – 358

Markman, E.M. (1989). *Categorization and naming in children: Problems of induction.* Cambridge, MA: MIT Press.

McKoon, G. & MacFarland, T. (2002). Event templates in the lexical representations of verbs. *Cognitive Psychology, 45*, 1-44.

McRae, K., Ferretti, T. R., & Amyote, L. (1997). Thematic roles as verb-specific concepts. Language and Cognitive Processes, 12, 137-176.

Medin, D. L., & Schaffer, M. M. (1978). A context theory of classification learning. *Psychological Review*, 85, 207-238.

Morris, M. W., & Murphy, G. L. (1990). Converging operations on a basic level in event taxonomies. Memory & Cognition, 18, 407-418.

Murphy, G. L., & Medin, D. L. (1985). The role of theories in conceptual coherence. *Psychological Review, 92*, 289-316.

Nosofsky, R. M. (1986). Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology: General*, 115(1), 39-57.

Osherson, D. N., Smith, E. E., Wilkie, O., Lopez, A., & Shafir, E. (1990). Category-based induction. *Psychological Review,* 97**,** 185–200.

Owens, J., Bower, G. H., & Black, J. B. (1979). The "soap opera"effect in story recall. *Memory and Cognition, 7,* 185-191.

Posner, M.I., & Keele, S.W. (1968). On the genesis of abstract ideas. *Journal of Experimental Psychology, 77*, 353-363.

Premack, D., (1983) Animal cognition. *Annual Review of Psychology*, *34:* 351-362.

Rehder, B. & Ross, B.H. (2001). Abstract coherent concepts. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 27*, 1261-1275

Rehder, B. (2003). Categorization as causal reasoning. *Cognitive Science, 27*, 709-748

Rips, L. J., Shoben, E. J., and Smith, E. E. (1973) Semantic distance and the verification of semantic relations. *Journal of Verbal Learning and Verbal Behavior, 12*: 1-20.

Rosch, E.H. (1973) Natural categories. *Cognitive Psychology 4*, 328-50.

Rosch, E., and Mervis, C. B., (1975) Family resemblances: Studies in the internal structure of categories. *Cognitive Psychology, 7*: 573-605.

Ross, B. H., & Murphy, G. L. (1999). Food for thought: Cross-classification and category organization in a complex real-world domain. *Cognitive Psychology, 38,* 495-553.

Schank, R. C., (1982) Dynamic Memory (New York: Cambridge University Press).

Schank, R. C., and Abelson, R., (1977) Scripts, Plans, Goals and Understanding. Hillsdale, NJ: Lawrence Erlbaum Associates.

Schnyer, D.M., Maddox, W.T., Ell, S., Davis, S., Pacheco, J., & Verfaellie, M., (in press). Prefrontal Contributions to Rule-Based and Information-Integration Category Learning. *Neuropsychologia*

Sieck, W. R., Quinn, C. N., and Schooler, J. W., (1999) Justification effects on the judgment of analogy, *Memory and Cognition, 27*(5): 844-855.

Smith, L. B., (1989) From global similarities to kinds of similarities: The construction of dimensions in development. In S. Vosniadou and A. Ortony (eds) *Similarity and Analogical Reasoning* (New York: Cambridge University Press), pp. 146-178 .

Thibaut, J.-P., French, R., & Vezneva, M. (2009) Analogy Making in Children. *Proceedings of the 31st Annual Meeting of the Cognitive Science Society.*

Thompson, R. K. R., and Oden, D. L., (2000). Categorical perception and conceptual judgments by nonhuman primates: the paleological monkey and the analogical ape. *Cognitive Science, 24*(3): 363-398.

Tversky, A., (1977). Features of similarity. *Psychological Review, 84*(4), 327-352.

Wechsler, S. (1995) *The Semantic Basis of Argument Structure: A Study of the Relationship Between Word Meaning and Syntax*. Stanford University: CSLI Publications.

Wittgenstein, L., 1968, Philosophical investigations, trans. G.E.M. Anscombe (New York: MacMillan) .

Vosniadou, S. (2009) Models and Analogies in Conceptual Restructuring. *Symposium presented at the 31st Annual Meeting of the Cognitive Science Society.*

Voss, J. L., & Paller, K. A. (2006). Fluent Conceptual Processing and Explicit Memory for Faces Are Electrophysiologically Distinct. *Journal of Neuroscience, 26*(3), 926-933.

Yamauchi, T. (2005). Labeling bias and categorical induction: Generative aspects of category information. *Journal of Experimental Psychology: Learning, Memory and Cognition, 31*, 538-553

Yamauchi, T., & Markman, A.B. (1998). Category learning by inference and classification. *Journal of Memory and Language, 39*, 124-14

Yamauchi, T., & Markman, A.B. (2000). Inference using categories. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 26*, 776-795.

Zeithamova, D., Maddox, W.T. (2007) The Role of Visuo-Spatial and Verbal Working Memory in Perceptual Category Learning. *Memory & Cognition,35*, 1380-1398.

VITA

Micah Balser Goldwater attended Newton South High School, Newton, MA. In 1999 he entered the University of Rochester in Rochester, NY. He received his Bachelor of Arts from the University of Rochester in 2003. In September 2004, he entered the Graduate School at the University of Texas at Austin.

Permanent address: 1471 W. Carmen Ave. Apt. 3N, Chicago, IL, 60640.

This manuscript was typed by the author.