NORTHWESTERN UNIVERSITY

Comparison Driven Representational Change

A DISSERTATION

SUBMITTED TO THE GRADUATE SCHOOL

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

for the degree

DOCTOR OF PHILOSOPHY

Field of Computer Science

By

Balasubramanian ‘Subu’ Kandaswamy

EVANSTON, ILLINOIS

December 2016

ABSTRACT

How mental representations are constructed and how they evolve are central problems for cognitive science. Representation decisions help determine what computations are hard or easy. Structured, relational representations are a hallmark of human cognition. Developmental studies show that children do not perform as well as adults in tasks that require noticing relational similarity. What drives this development? Gentner and her colleagues have argued that comparison and language are two forces driving this change. This thesis explores these ideas further by presenting a computational model of forced choice tasks to illuminate the roles of comparison and language in driving representational change.

The model simulates the following roles of comparison. First, comparison can be used to make selections in forced-choice tasks. Second, comparisons from recent positive experiences are assimilated as *interim generalizations* which are retrieved for subsequent tasks and influence encoding by highlighting relevant structure. Third, comparisons suggest opportunities for re-representation. Finally, verifying candidate inferences resulting from a comparison provides a way to augment encodings with background knowledge, thus enriching representations. The model simulates the role of language in facilitating the creation and enrichment of generalizations as follows. When two objects are given the same label, the model compares them. This leads to an interim generalization associated with that label, enriched with commonalities from background knowledge.

We tested these hypotheses by extending the Companion cognitive architecture and simulating three developmental studies. To reduce tailorability, the visual stimuli were provided as sketches and the objects were labeled using simplified English. The model was evaluated by comparing its behavior

and learning trajectory to that of children in the developmental studies. The performance of the model in the simulations provide evidence for the claims of this thesis.

ACKNOWLEDGEMENTS

I thank Northwestern University and the Office of Naval Research for their generous financial support.

I am grateful to my advisor Ken Forbus. I thank him for giving me this wonderful opportunity. Without Ken's support and guidance, I could not have gotten this far. He taught me how to seek knowledge and stay excited about research. He was always patient with me and showed me how to turn every mistake into an opportunity for learning. Ken is an example of how a 'guru' (गुरु) should be.

There is a unique advantage in being a member of the QRG Cogmodeling group. You get to work with two great professors, Ken Forbus and Dedre Gentner. I am very proud to have had the opportunity to collaborate with Dedre in research. I thank her, Chris Riesbeck and Ian Horswill for being in my thesis committee and for providing constructive feedback.

I thank Tom Hinrichs and Madeline Usher for their guidance and support. Thanks for helping me solve problems that I thought was impossible to solve. No matter how many times I came running to you for help, you greeted me with a smile. Thanks for being such wonderful people.

There is an informal buddy system in QRG. Every new student works closely with a senior student. As a new student, I had the opportunity to work closely with Andrew Lovett. I thank him for showing me the ropes. I thank Matt Mclure for being a good friend and my office mate for the past 6 years. Matt is my comrade in building edge segmentation modules for CogSketch. I thank Maria Chang and David Barbella for teaching me how to write a thesis. I thank my fellow QRG researchers and friends, Joe Blass, Clifton McFate, Max Crouse and Irina Rabkina for all their help. I also thank Dedre's students, Christian Hoyos, Ruxue Shao, Linsey Smith and others, in our Cogmodeling group for helping me understand things from the perspective of a cognitive psychologist.

I thank Chandra Bhagavatula, Thanapon Noraset and Chen Liang for teaching me AI techniques outside my research area. I benefited a lot from all the discussions we had, while walking to tech-express to get coffee.

I thank my dear friends Neeta Boga, Christopher Francis, Spoorthi Sanjeevi, Arjun Kannan, Roopa Natarajan, Ishan Saxena and Aditi Bose. They made me appreciate my decision to come back to school for a PhD. I thank them for all the fun board-game nights and dinners. I thank my brother-in-law Vamsi Yalamanchili for helping me with R studio and data analysis.

Most of all, I thank my mom, Lalitha, and my wife, Sravanthi, for believing in me and for supporting me in everything I do.

This dissertation is dedicated to

My teachers,

My mom, Lalitha Kandaswamy

and

My wife, Sravanthi Yalamanchili

Contents

List of Tables

List of Figures

Chapter 1: Introduction

1.1 Motivation

One of the important problems in cognitive science is to understand how humans make sense of their surroundings. Humans are constantly bombarded with abundant information, yet they manage to make sense of things around them seemingly without much effort. Although a complete answer is yet to emerge, many agree that our extraordinary ability to pick out patterns is part of how we manage to weed through the information chaos.

Analogy is central to our ability to think about relational patterns and is fundamental to human cognition (Holyoak, Gentner, & Kokinov, 2001). The importance of analogy is well established in the field of cognitive science. One widely regarded theory of analogy is Gentner's *structure-mapping theory* (Gentner, 1983). According to the structure mapping theory, mental representations involve relational structure, and comparison is the process of aligning them. Representations are important for analogical matching. A good comparison, which highlights relevant similarities and differences, emerges from encoding the problem/scenario using uniform representations.

The *Structure mapping engine* (SME) is a computational model of structure-mapping theory (Falkenhainer, Forbus, & Gentner, 1989). SME takes as input two descriptions, a base and a target, and produces mappings between them. The descriptions can be any kind of predicate representation, including stories, problems, and descriptions of visual scenes. SME has been used successfully in many cognitive simulations and also as a component of large scale artificial intelligence systems (Forbus & Hinrichs, 2006)(Friedman, 2012).

Despite successes in modeling a wide range of phenomenon, structure mapping theory and SME is not without its critics. One criticism of structure mapping comes from Hofstadter and his colleagues (Chalmers, French, & Hofstadter, 1992; Hofstadter, 2008). Chalmers et al. (1992) propose a different approach for analogy called High-level perception (HLP). They argue that flexibility of human cognition arises out of a highly interwoven interplay between lower order perceptual processes and higher order cognition. They argue that these processes cannot be studied in isolation. They criticize computational models like SME for separating the processes and accuse SME of bypassing the process of perception, by starting with pre-derived representations.

Forbus, Gentner, Markman, & Ferguson (1998) addressed the criticisms by Chalmers et al. (1992). They agree with CFH with the importance of studying the processes for building representations. Nevertheless, they argued for the advantages of decomposing analogical processing into constituent subprocesses. In their view, there are three coarse-grained ways to think about how perception and cognition could interact (see Figure 1)

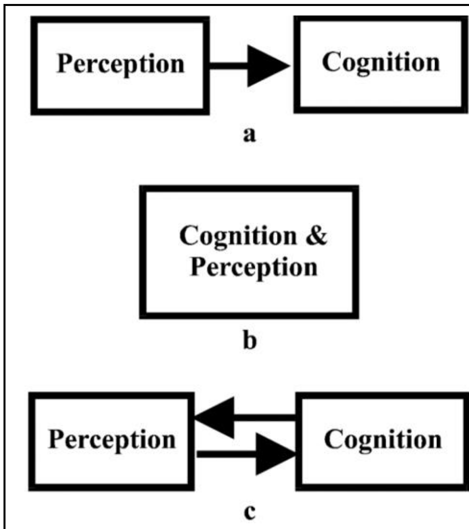

Figure 1: Three abstract views of perception and cognition (Forbus et al., 1998).

Forbus et al. (1998) points out that the classic stage model depicted in part (a) is the strawman that Hofstader et al. argue against and is not the same as what they propose, which is part (c). The model depicted in part(c) is the primary motivation for this thesis. Both the criticism by Chalmers et al. (1992) and the response by Forbus et al. (1998) centers around the importance of representations. Analogy is crucial to cognition and representations are crucial to analogy. Given the centrality of representations in analogy, it is imperative to understand the origin, acquisition and the evolution of our representational capabilities. Markman & Dietrich (2000) stress the importance of studying representational change. They say that cognitive science cannot proceed without studying representations and the computational processes defined over them. According to them, representational change provides a way to account for fluidity in cognition.

The mechanisms driving representational change can be understood better within the framework of the Structure Mapping theory. Gentner and her colleagues propose that structure mapping comparisons, language learning and progressive alignment, which states that experiencing concrete similarities enables appreciation of more abstract similarities, are among the main factors driving representational change (Gentner & Namy, 2006; Kotovsky & Gentner, 1996). One useful set of distinctions in understanding the role of comparison in representation change concerns the kinds of information aligned. A mapping is considered an *appearance match* if most of what matched are attributes of objects (and first-order perceptual relations, e.g. above). It is an analogy if most of the match involve relational information, with few object attributes. It is *literal similarity* if the match contains a high number of relations and object attributes in correspondence.

These distinctions are useful in understanding human processing of comparisons. To explore how children, improve their representations, Gentner and colleagues used *forced choice tasks*, which we simulate here. In a forced-choice task, children are presented with a standard and two or more alternatives. Children are required to pick the choice that is more similar to the standard(s) (Christie & Gentner, 2010)(Namy & Gentner, 2002)(Kotovsky & Gentner, 1996). They controlled the type of similarity between the standard(s) and the choices, as well as the details of the instructions, to shed light on how comparison can be used to promote learning.

They found that younger children exhibit a reliance on holistic or object-level similarities. Children interpret comparisons like “The prison guard’s heart was hard as a stone” to mean that the person’s heart was literally as hard as a stone (Rattermann & Gentner, 1998). In subsequent studies, they also showed how comparison and progressive alignment can be used to change the way children represent, enabling the children to appreciate relational similarity (Kotovsky & Gentner, 1996)(Rattermann & Gentner, 1998).

The studies clearly show evidence for developmental change in the recognition of relational similarity. According to the Relational Shift hypothesis (Gentner, 1988), this is because as children learn, they understand more relationships, and hence have more relationships that they can encode, and thus the relational content of their mental representations increases. The shift is considered epistemological rather than maturational, and hence it can (and does) happen at different times in different domains (Gentner, 1988). The appreciation of relational similarity increases with experience and with the accretion of relational knowledge via language. Gentner and her colleagues demonstrated this via several developmental studies. These studies are prime candidates for studying representational change via computational simulations. The knowledge of how to represent a situation or a problem is learned over experience and is a long term process. However, as the studies show, the change can occur even within a single experimental session. The objective of the thesis is to provide a computational model sufficient to explain the phenomenon of representational change observed within such short durations. We demonstrate the model's effectiveness by simulating three of the developmental studies by Gentner and her colleagues.

1.2 Claims and Contributions

Claim 1: Recent experiences affect how new problems or tasks are encoded. This can be modeled using interim generalizations.

The Sequential Analogical Generalization Engine (SAGE) model of analogical generalization focuses on learning generalizations that are stored in long-term memory (McLure, Friedman, & Forbus, 2010). In SAGE, *generalization pools* provide models of concepts. We propose that there are also generalizations built up in a similar way in working memory, which we call interim generalizations (Kandaswamy, Forbus, & Gentner, 2014). Interim generalizations are proposed to be involved in within-task comparisons and learning. Only a small number of descriptions can be stored in interim generalization pools, and retrieval is based on similarity, biased by recency. In our model, comparisons of a current problem with interim generalizations provides a filter, removing elements that may be irrelevant for the task.

Claim 2: Forced choice tasks can be modeled using structure mapping comparisons.

1) The difference in structural evaluation scores between a standard and the alternatives is used for selecting a winning choice.
2) Verification of candidate inferences produced by comparisons is used to improve mappings.
3) When comparisons result in scores that are too close to discriminate, re-representation is triggered to attempt to differentiate between the alternatives.

Candidate inferences specify what additional knowledge in the base can potentially be transferred to the target. Consequently, the candidate inferences which were verified to be true becomes part of the target and thereby increases the score of the resultant mapping. Verification involves comparing the candidate inferences to prior knowledge about the target. If scores are close, a decision cannot be made. We propose that re-representation (Yan, Forbus, & Gentner, 2003) is performed in order to improve discrimination. In our model, we use the difference in average-self matching score to determine if re-representation is worth the effort.

Claim 3: Labeling two examples the same triggers a comparison for the purpose of understanding the meaning of the label. This can be simulated using structure mapping comparisons and generalizations.

1) The examples that are labeled the same are compared and assimilated into a generalization.
2) The generalization highlights commonalities and deemphasizes dissimilarities.
3) Examples are augmented with conceptual commonalities if possible, resulting in enhanced generalizations.

For each novel word/label, a generalization pool is created to assimilate examples introduced with that label. Generalization pools have a threshold that specifies how similar the examples should be for being assimilated into a generalization. We model the tendency for a common label to lead to assimilation via lowering the assimilation threshold when a label is provided in a comparison. The augmentation of prior knowledge about concepts provides a means of connecting new information with existing knowledge.

Contribution 1: A model of interim generalization pool (SAGE-WM)

We built SAGE-WM, a model of interim generalization pools, based on SAGE (McLure, Friedman, & Forbus, 2010). We limit the number of elements that can be present to emulate working memory constraints. Psychological evidence suggests that the use of interim generalization pools is governed by similarity and recency. We model these by using SME for retrieving the most similar generalization or example, while biasing the retrieval based on recency.

Contribution 2: A model of forced choice tasks

The primary contribution of the thesis is a model of forced choice tasks built on the Companions Cognitive architecture (Forbus & Hinrichs, 2006). In addition to supporting analogical reasoning and learning, the Companion architecture provides natural language understanding capabilities and sketch understanding capabilities, which we use to automatically encode the stimuli in modeling the psychological experiments.

We simulated three studies to capture the role of comparison in the phenomenon of representational change within forced-choice tasks:

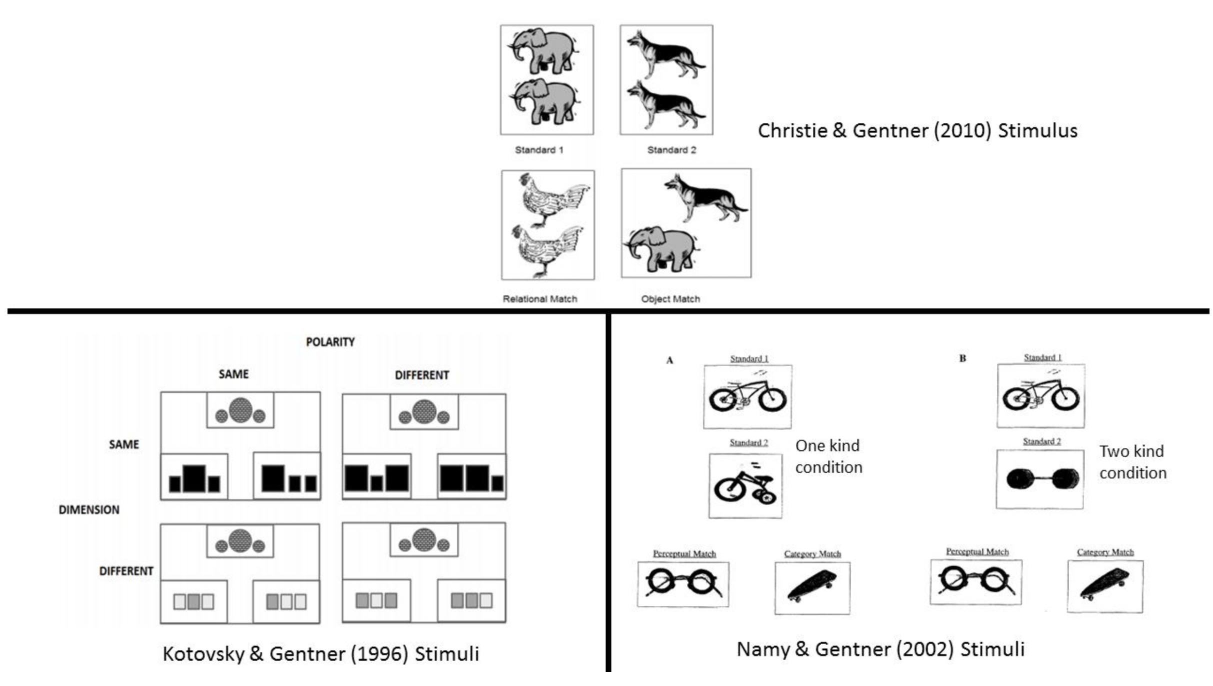

Figure 2: Examples of stimuli from Kotovsky & Gentner (1996), Christie & Gentner (2010), and Namy & Gentner (2002).

1) Christie & Gentner (2010): Where hypotheses come from: Learning new relations by structural alignment.
2) Kotovsky & Gentner (1996): Comparison and categorization in the development of relational similarity.
3) Namy & Gentner (2002): Making a Silk Purse Out of Two Sow's Ears: Young Children's Use of Comparison in Category Learning.

Each study highlights various roles played by structure mapping comparisons in representational change. Figure 2 provides examples from each study. Christie & Gentner (2010) showed the importance of common label and explicit comparison. They introduced the children in the comparison condition to two standards with the same label ("this is a jiggy") and gave an explicit instruction to compare them (example: "can you see why these both are jiggies?"). The children in comparison condition made the relational choice more often than the children in other conditions who did not get an opportunity to compare. Their results show the power of comparison in bringing to focus the relational similarity that were once hidden, perhaps because children may encode more object attributes by default.

Namy & Gentner (2002) studied whether comparison during word learning not only highlights commonalities but could also enrich the abstraction by utilizing prior knowledge about the objects. They used stimuli consisting of objects that were familiar to the children. The experiment had two conditions. In both conditions, they introduced the objects with the same label and invited them to compare. In the one-kind condition, the standards are from the same category, while in the two-kind condition they belong to different taxonomic categories. In both conditions, the standards share perceptual commonalities. One of the choices is perceptually similar to the standards but belongs to a different category. The other, called the taxonomic choice, is from the same category of at least one of the standards but does not share perceptual commonalities with them. (see Figure 2). As predicted, the children in the one-kind condition made the taxonomic choice more frequently than the children in the two-kind condition.

Christie & Gentner (2010) and Namy & Gentner (2002) show how within-task comparison leads to representational change. But comparison has a much larger role to play in shaping representations. Between-task comparisons of progressively alignable stimuli can provide valuable experience through which children can learn to discriminate between what is relevant or irrelevant to the given task. They addressed this in Kotovsky & Gentner (1996). They explored children's performance on tasks involving simple higher-order patterns, specifically symmetry and monotonic increase. They increased the stimuli complexity by changing the polarity and the dimension of change (see Figure 2). The 4 year olds performed poorly on all but the same-dimension-same-polarity triads. In a subsequent experiment, they prepared an ordered set of stimuli which promotes progressive alignment. Children in the progressive alignment (PA) condition improved in their performance in cross dimension tasks, while the children in the control condition did not.

The studies show how comparison, language and progressive alignment can enable children to appreciate relational similarity, perhaps by influencing the way they represent the stimuli. We evaluated our claims about representational change by building a model of forced choice tasks and comparing its response with that of the children in the studies. We discuss the studies and the model in more detail in later chapters.

1.3 Organization

Chapter 2 reviews the theoretical background and systems used in this thesis. It includes an overview of the Companion cognitive architecture plus the enhancements made to it as part of the thesis. We also introduce SAGE-WM, a model of interim generalizations, based on SAGE.

In Chapter 3, we present our model of forced choice task and analyze the roles played by structure mapping comparisons in representational change. First, we explore from how analogs from recent experience affect encoding. Second, we describe methods for conceptual augmentation, via enriching generalization and via candidate inference validation. Finally, we describe the role of re-representation in resolving impasses in forced choice tasks.

In chapter 4, we present our simulations and results. First, we explore the role of comparison and language in learning relational abstractions via modeling Christie & Gentner (2010) experiments. The simulation highlights the representational change brought by the effects of common labels and comparison. Second, via simulating Namy & Gentner (2002) experiments we show how word learning comparisons not only highlight commonalities, but also enhance the resultant abstraction by bringing in additional knowledge when possible, as a form of representational change. Finally, we show how recent comparisons assimilated as interim generalizations helps in the transfer of representational knowledge via simulating experiments from Kotovsky & Gentner (1996).

We discuss related work in chapter 5. In Chapter 6, We revisit the claims and finish with general discussion and suggestions for future work. The final sections of the thesis include references and appendices.

Chapter 2: Analogy and the Companion Cognitive Architecture

This chapter provides the background needed for the rest of the thesis, by summarizing structure-mapping theory, the simulations we are using, and the Companion cognitive architecture. A primary contribution of the thesis, SAGE-WM is introduced here, along with extensions to the Companion architecture and CogSketch.

2.1 Structure Mapping Theory

Structure mapping theory proposes that analogy is the process of mapping knowledge from one domain (the base) into another (the target). Human mental representations are structured and comparison is the process of aligning them. Analogy and similarity emerge out of the same process of structural alignment. The comparison is categorized into analogy, literal-similarity, or appearance match based on the nature of the alignment.

Structure mapping theory postulates psychological constraints on the alignment. First, the *parallel connectivity* constraint states that statements in alignment must have their arguments also in alignment. Second, the *one-to-one correspondence* constraint states that one element of the base can match to only one other element of the target. Together, these constraints imply *structural consistency* of a mapping. The *tiered identicality* constraint states that local matches are only proposed when predicates are identical, or when aligning non-identical functions can lead to larger matching structures, i.e. by ensuring that parallel connectivity is not violated. Finally, the *systematicity preference* is that mappings which involve systems of interconnected relations are preferred. Systematicity biases the alignment process towards interpretations that provide explanations. Since explanations involve statements of causal, inferential, and evidential relationships involving other statement, they are preferred over isolated relational structures.

Structure-mapping makes the following representational assumptions:

1) A structured description consists of objects, their attributes, and relations between objects. Objects (often referred to as entities) can be animate or inanimate, physical or conceptual. An instance of a triangle is an object, as is the number 3, in this sense. Attributes are statements using one-place predicates, e.g. the category of an object, its shape, color, and so on. Relations are statements using binary or higher arity predicates. Spatial relations and causality are two common types of relationships.
2) The *order* of a statement is one plus the order of its arguments. Objects have order zero. Thus attributes (e.g. `(Person John)`) are first-order statements, as are relations involving objects (e.g. `(inside John Cave)`). Statements that connect other statements are higher-order statements, e.g. `(cause (and (MovingEvent Run5) (performedBy Run5 John)) (inside John Cave))` is a second-order statement.
3) Functions are used to represent dimensions or components of an object or situation, e.g. `DarknessFn` to indicate how light something is. There is a tradeoff between using domain-specific relations and more general relations with domain-specific dimensions. As described in Chapter X, it appears that children start with more domain-specific relations and later re-represent using dimension independent relations. For example: `(darkerThan A B)` can be re-represented as `(greaterThan (DarknessFn A) (DarknessFn B))`.

2.2 Structure-Mapping Engine

The Structure-Mapping Engine (SME) (Falkenhainer, Forbus, & Gentner, 1986) provides a computational model of the structure-mapping process. It takes as input two propositional descriptions, the *base* and *target*, and produces one or more *mappings* as its output. The mapping consists of three components:

1) A set of *correspondences* between the structural elements (attributes, relations and objects).
2) A *score* that provides a numerical estimate of similarity, which can be normalized as explained in Chapter 3.
3) A set of *candidate inferences* projected based on non-overlapping structure according to the correspondences of that mapping.

SME typically produces only one mapping, but can produce up to three if there are close alternatives. It uses a greedy algorithm to provide good answers in polynomial time.

2.3 Sequential Analogical Generalization Engine (SAGE)

SAGE is a model of analogical generalization (McLure, Friedman, & Forbus, 2010). SAGE takes as input a sequence of positive examples of a concept, and produces in long-term memory a set of generalizations and unassimilated examples in a *generalization pool* representing that concept. The generalizations are probabilistic, in that they include the likelihood of each statement (attributes and relations). They can be disjunctive, if there is more than one generalization. Outliers are handled by maintaining unassimilated examples.

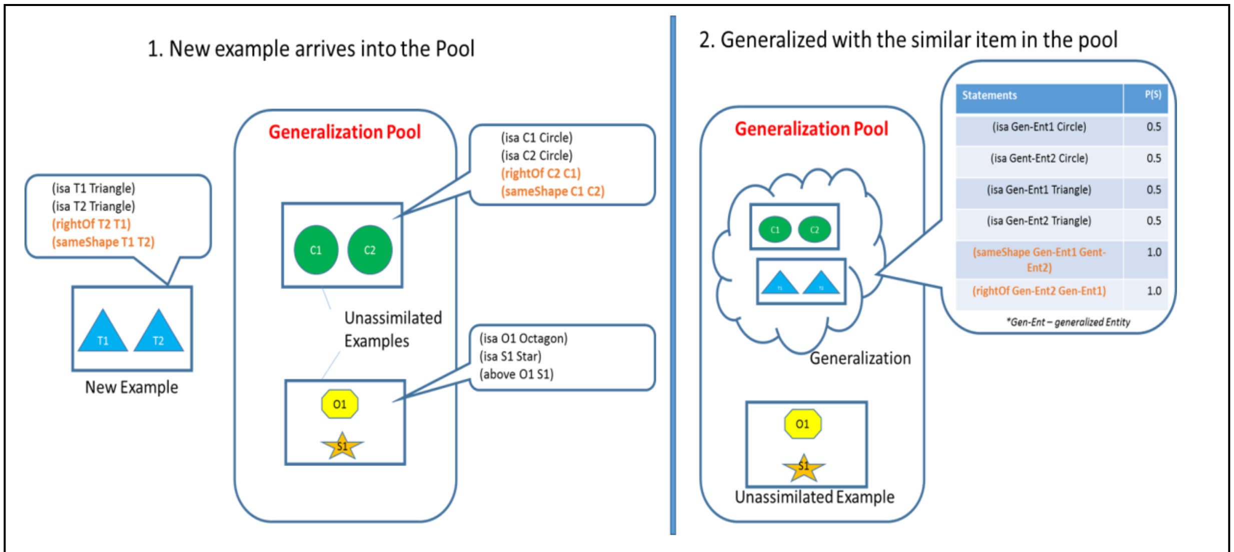

Figure 3: Illustration of SAGE operation.

SAGE's processing of an example works like this. Given a new example E, SAGE uses MAC/FAC (Forbus, Gentner, & Law, 1995), a model of similarity based retrieval, to find the most similar item in the pool to E. This item can be either a generalization or an unassimilated example. If it is a generalization and the similarity score is over an *assimilation threshold*, SAGE merges the new example into the generalization. If it is an example and the similarity score is over the threshold, SAGE merges the two examples into a new generalization. The merge process combines the overlap between its two inputs, and updates the frequency of occurrence of each statement in the alignment. That is, for merging two examples, a statement in both will have a probability of 1.0 while a statement occurring in one but not the other will have a probability of 0.5. When merging an example into a generalization, the frequency of occurrence for each statement is updated similarly. This can cause infrequent, variable, or noisy information to "wear away" with experience. An example of SAGE generalization is depicted in Figure 3.

SAGE has been used to model a variety of psychological phenomena, including learning words (Lockwood, Lovett, & Forbus, 2008), learning grammar (Taylor, Friedman, Forbus, Goldwater, & Gentner, 2011), and conceptual change (Friedman, 2012)) .

2.4 A model of interim generalizations (SAGE WM)

In analogical processing, new information about an object or event can be inferred based on its similarity to an existing generalization. There is psychological evidence that recent experience provides a context to interpret and understand our current situation (Day & Gentner, 2007). We propose that a SAGE-like mechanism is also used in working memory, to construct interim generalizations. The model for this processing is SAGE-WM.

Most of the operations of SAGE-WM are the same as SAGE. We assume that there is a default generalization pool in working memory that is used to help accumulate experiences with tasks. We assume that the interim generalization pool can only maintain a small number of descriptions. Given a new example, SAGE-WM retrieves the most similar item from this pool, based on both similarity and recency. Moreover, when an item has a label associated with it, the label is used as a filter for retrieval, to model people's preference for retrieving examples or generalizations that share the same label.

2.5 CogSketch

CogSketch is a domain-independent sketch understanding system (Forbus, et al., 2008). Users interact with CogSketch by drawing using a pointing device (e.g. a mouse or stylus). CogSketch interprets the ink by computing spatial, positional and shape relations between objects, and optionally within objects as well. Furthermore, CogSketch supports multiple subsketches within a single sketch. We use this feature to represent the stimuli in a forced choice task. Recall that a typical forced choice stimulus consists of one or more standards and multiple choices. The standards and choices are represented as subsketches within a sketch.

Figure 4 depicts an example from Christie & Gentner (2010) experiments and the corresponding subsketch. Cogsketch automatically computes qualitative visuo-spatial relations for objects and creates

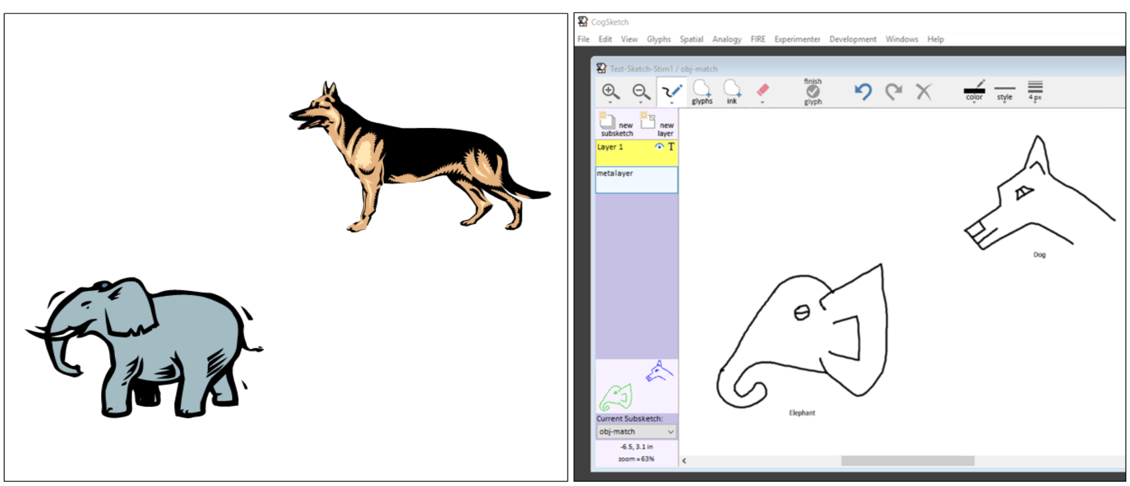

Figure 4: Example from Christie & Gentner (2010) stimuli and the corresponding sketch.

descriptions for each subsketch. For example, cogsketch computes that the dog in Figure 4 is *above* and to the *right-of* the elephant.

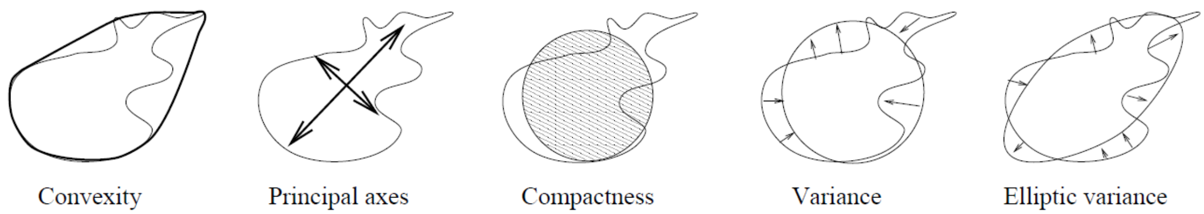

Figure 5: Simple shape descriptors (Peura & Iivarinen 1997).

Computing shape relations and shape attributes require segmenting the objects to its constituent parts such as edges and edge-cycles (McLure, Friedman, Lovett, & Forbus 2011). As part of the thesis work, we increased the accuracy and reliability of CogSketch corner detection and segmentation algorithms to handle organic-shaped ink representing real world objects, such as the ones used in the stimuli from psychological studies. Furthermore, we augmented the representations created by CogSketch to include shape-relations such as *same-shapes*, *reflected-among-X-Axis* (a mirror reflection among x-axis), *same-color*, etc. We also augmented shape-attributes created by CogSketch to include qualitative shape descriptions based on simple shape properties such as shape-convexity, compactness, orientation, circularity and elliptic variance, etc. (Peura & Iivarinen, 1997).

Additionally, CogSketch allows users to provide conceptual labels for objects in sketches using collections (i.e. concepts) and relations from the Cyc knowledge base. By labeling an item with a collection from Cyc, the user is indicating that the item is an instance of that collection. We use CogSketch's conceptual labeling as an alternative to object recognition. For example, instead of CogSketch having to recognize the ink in Figure 4 to represent a Dog and an Elephant, we label the objects with the Cyc Collections 'Dog' and 'Elephant' respectively.

2.6 Companion Cognitive Architecture

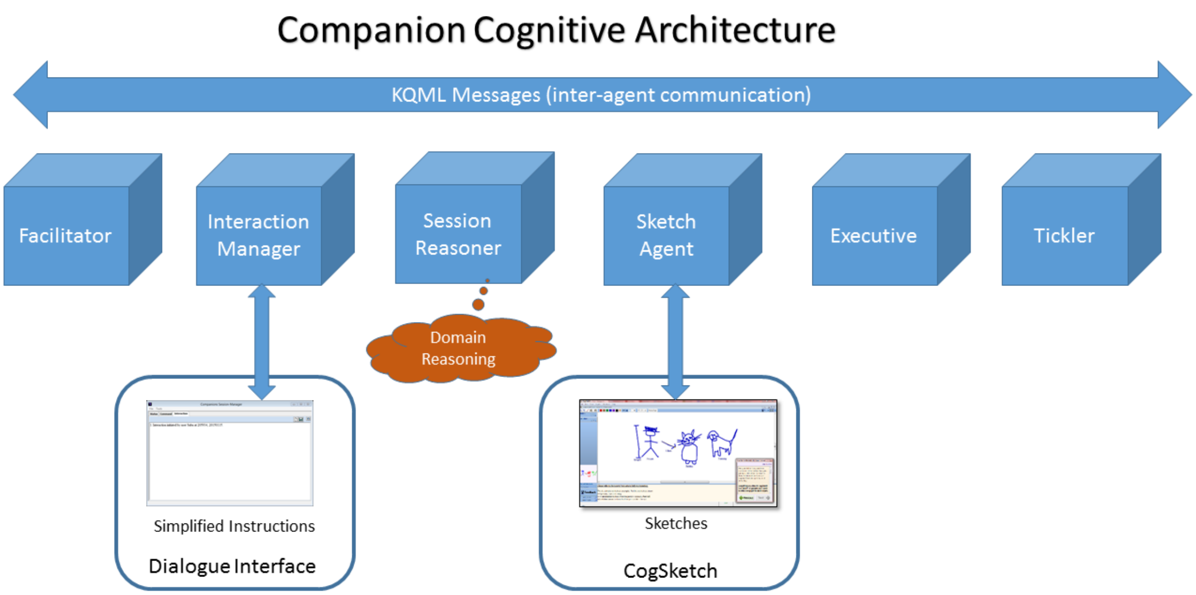

Figure 6: Companion Cognitive Architecture.

Every cognitive architecture is based on a fundamental idea. For example, the Soar cognitive architecture is based on the problem-space computational model (Laird, 2012). The Companion cognitive architecture is based on the idea that analogy lies at the heart of human cognition (Forbus, Klenk, & Hinrichs, 2009). It is a multi-agent architecture and inter-agent communication is achieved using KQML messages. Companion agents use computational models of analogical processes: Structure Mapping Engine (SME) for similarity comparison, SAGE for persistent generalizations, MAC/FAC for retrieval and SAGE-WM for interim generalizations. Some Companion agents are specialized for a specific service such as sketch processing (CogSketch agent) and processing user interaction (Interaction Manager). In summary, the Companion architecture acts as the glue that ties together different essential tools and provides us with a unified apparatus to model cognition.

2.7 Interaction Manager

The *Explanation Agent Natural Understanding System* (EA NLU) (Tomai, 2009) is used as the dialogue understanding system that processes the language portion of user interaction. EA uses head-driven phrase structure grammar and a lexical source to parse and interpret language. The interpretation process uses Discourse Representation Theory (Kamp & Reyle, 1993) and produces a semantic interpretation of the text. When producing the interpretation, EA maintains parse, sense and coreference ambiguities as choice sets. Once the ambiguities are resolved via selection of choices, the final interpretation is produced.

The Interaction Manager is the Companion agent that runs EA NLU. The parse, choice-set and interpretation information are stored in its working memory. To resolve reference ambiguities involving deictic references in multimodal dialogues, the Interaction Manager communicates with one or more CogSketch Agents, as described below.

2.8 CogSketch Agent

Each sketch being used by a Companion is linked into the system via a CogSketch Agent. This agent controls the CogSketch interface and provides reasoning services for other agents that need information from the sketch. References such as “the elephant”, for example, when one appears in a sketch, can be interpreted as referring to the entity in the sketch. Other agents can subscribe to events in the sketch, to detect user interactions such as changing what sketch is active or the selection of a glyph within a sketch. This helps maintain a shared focus for communication between a Companion and people. For example, to refer to sketched objects when using language to name them, they are selected first, so that the Companion knows which object(s) we are talking about.

2.9 Session reasoner

The Session Reasoner is used for domain reasoning. In this thesis, it is responsible for performing forced choice tasks, executing the instructions interpreted by the Interaction Manager and using data from it and from the CogSketch Agent. This is best illustrated by an example. When the user labels an object, e.g. "This is a Toma.", this statement is interpreted (by the Interaction Manager) as meaning that the label "Toma" should be applied to the currently selected entity (as found in the CogSketch Agent). The Session Reasoner gathers the relevant information from the CogSketch agent. Instructions to compare and choose are also handled by the Session Reasoner, which directs the CogSketch Agent to highlight the appropriate choice to indicate to the user what it has selected.

2.10 Multi-modal interaction

The forced choice task involves multi-modal interaction. The children performed the tasks according to the instructions provided by the experimenter, which involved the experimenter pointing to objects in

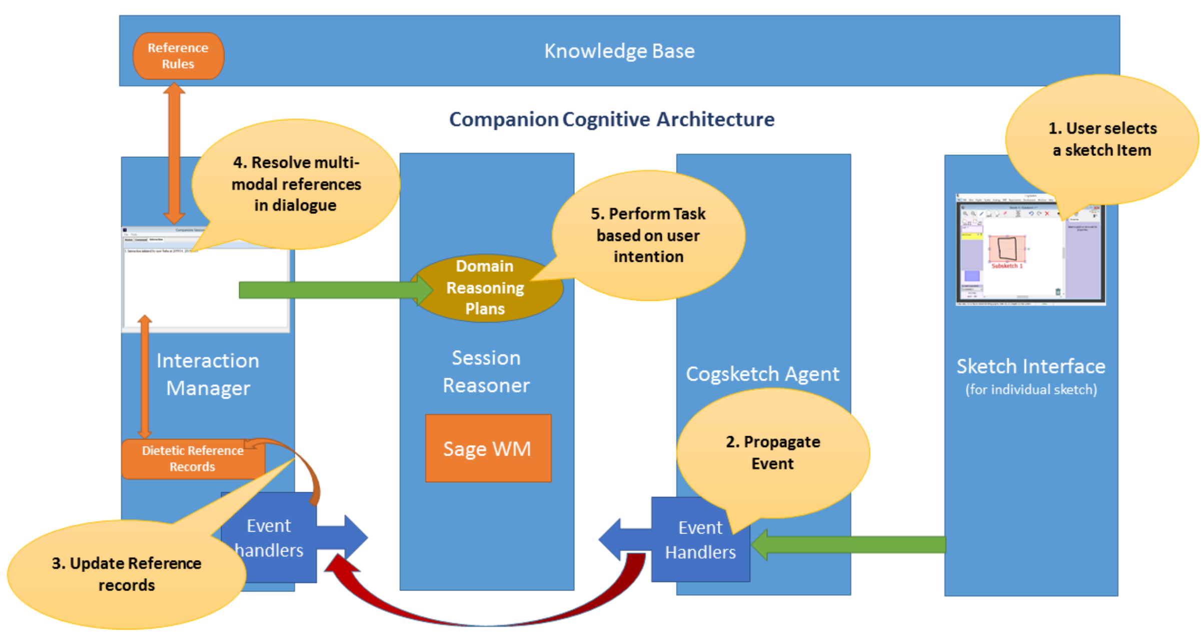

Figure 7: Inter-agent communication for multi-modal interaction.

pictures in front of the children and talking about them. Similar to the children, Companions receive input as CogSketch sketches and simple English language instructions. The equivalent of a pointing gesture for Companions is the selection of sketch items. As part of the thesis work, we added to Companion architecture the ability to resolve deictic references (pointing/selecting an object in the sketch and saying "this").

Figure 7 illustrates the deictic reference resolution process The Interaction Manager subscribes to the CogSketch Agent for sketch-related events. Based on the incoming event notifications, the Interaction Manager constantly updates the *DiecticReferenceContext* in its working memory. At any given time, the DiecticReferenceContext contains information about the currently active sketch/subsketch and the selected items in the sketch. The interaction manager resolves the demonstratives ('this' (or) 'these') in the dialogue to the selected sketch item (s) based on the DiecticReferenceContext. Thus, we emulate the interaction between the experimenter and the children (e.g. "This is a Toma").

When the experimenter says to the children "Which one of these is a toma?", they are asking the children to point to the choice that they think is the Toma. Such utterances in the Companion are marked as interrogatives, which involve demonstratives and are handled accordingly. The system resolves 'these' to the selected items in the sketch and uses the plans in the Session Reasoner corresponding to the forced choice model (chapter 3) to pick one of the choices as the answer.

Chapter 3: A Model of Forced Choice Tasks

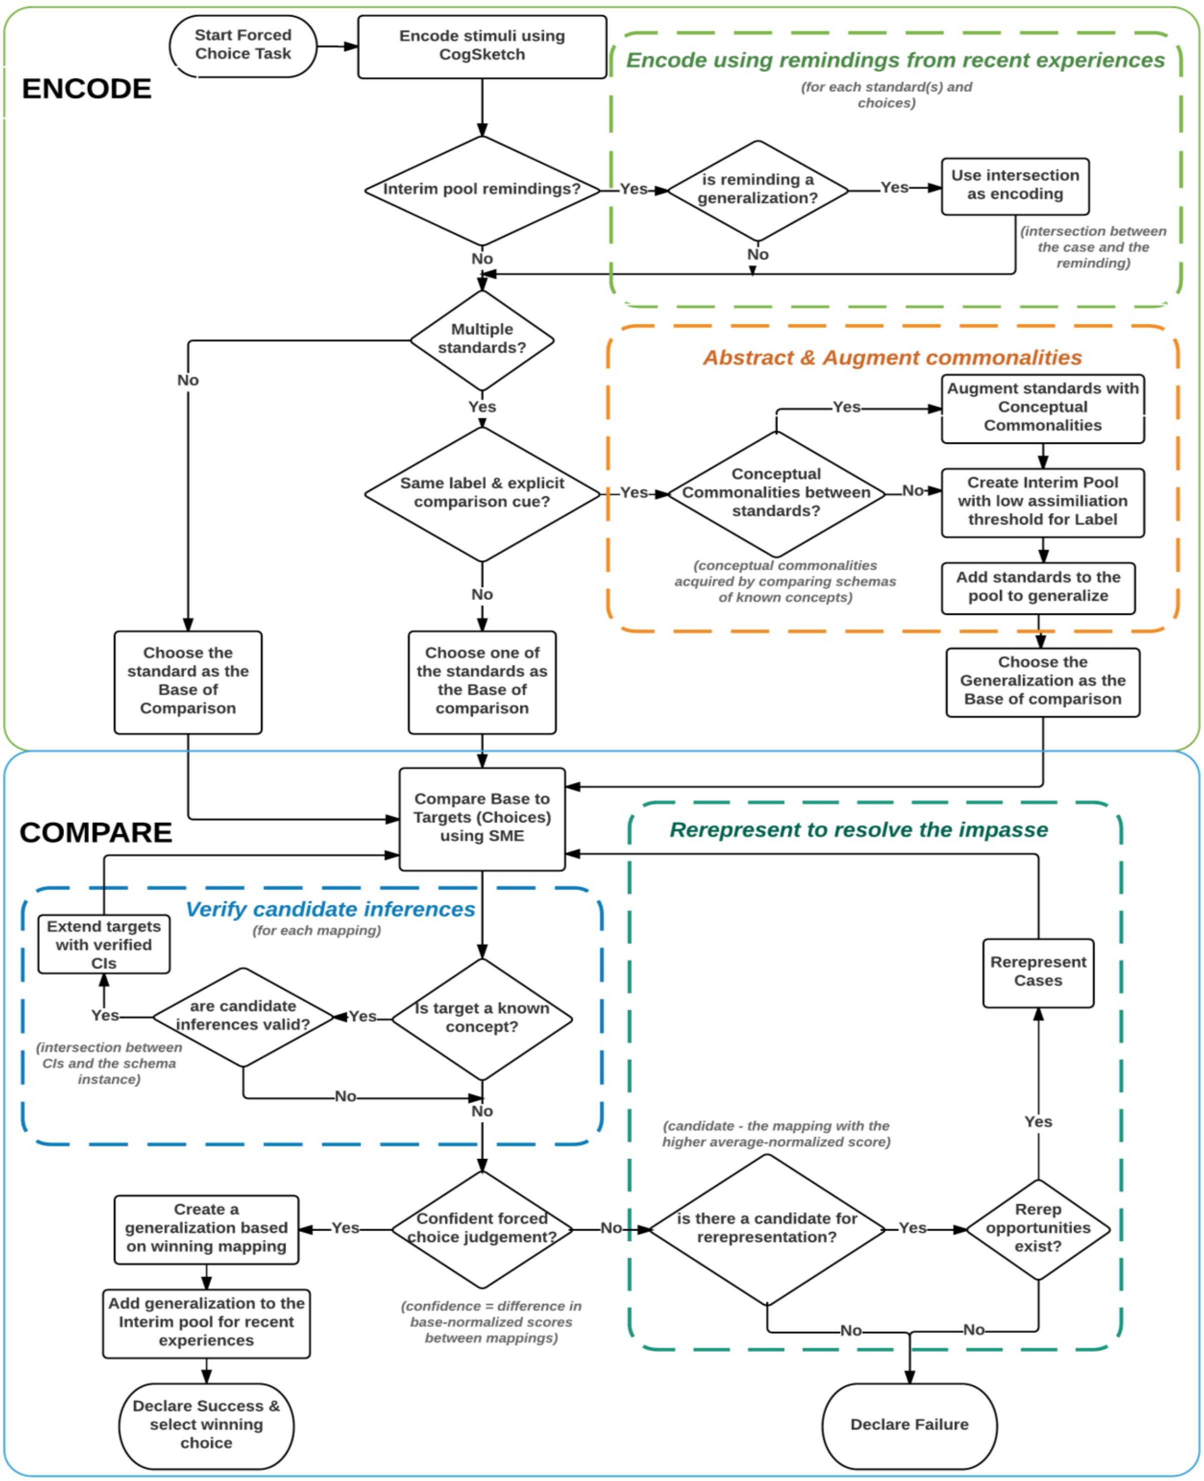

Figure 8: A model of forced choice task.

Forced choice tasks require choosing between two or more alternatives based on how similar they are to one or more standards. This provides a simplification of many real-world problems, which is why it is so heavily used in psychological experiments. They are easy to administer and provide a useful way of exploring many phenomena, especially involving similarity. There are many kinds of forced-choice tasks. For the purpose of the thesis, we are primarily interested in relational match forced choice tasks, where one of the choices shares relational structure with the standard.

A model of forced choice tasks should account for the following:

1) The processes determining the forced choice judgement and the ability to identify a confident decision, even in the absence of feedback.
2) The influence of recent experiences and prior knowledge in forced choice decision.
3) The processes for triggering and controlling representational improvements. (e.g. conceptual augmentation and re-representation).

The entire process can be divided into two phases: encoding and comparison. The encoding phase builds representations at the appropriate level of abstraction for solving the forced choice task. The comparison phase determines relative similarity and makes the choice. This may require re-representation if the choice is unclear. Even though we separate the operations into encoding and comparison, there is interaction between them via interim generalizations. We discuss each phases in turn.

3.1.1. Encode stimuli using CogSketch

This stage models the visual encoding processes that operate immediately upon the presentation of the stimuli. The visual stimuli used in forced choice experiments are presented as sketches to the CogSketch, which, as noted in Chapter 2, generates human-like qualitative visuospatial representations of the stimuli. Furthermore, we use CogSketch for conceptual labeling, which allows us to pick a concept from Cyc, to indicate the concept corresponding to the objects in the sketch, thereby sidestepping the need to model object recognition. The representations generated by CogSketch includes visuo-spatial, shape and conceptual information.

3.1.2. Encode using remindings from recent experiences

As noted earlier, psychological evidence suggests that people are applying recent experience to help encode new stimuli. We model this using remindings from SAGE-WM interim generalizations. When a choice is made, the winning mapping is added to the interim pool for recent experience. We encode the winning mapping by creating a generalization based on the mapping and filtering out facts that have probability less than 1.0. Thereby, eliminating statements that did not participate in the alignment. A case representing the winning mapping is created based on the remaining facts and added to the interim pool. Repeated exposure to similar tasks results in the assimilation of winning mappings as interim generalizations.

Kotovsky & Gentner (2010) studies show that children require multiple exposure to concrete similarity tasks to successfully transfer representational knowledge to more complex tasks that require abstract similarity. Based on their findings, we posit that humans, especially children, require more than one exposure to a task type in order to gain and transfer representational knowledge reliably. It is reasonable to assume that generalizations formed by repeated exposure (more than one) has a larger impact than an isolated experience. We model this by considering only remindings that are generalizations and ignoring remindings based on isolated experiences i.e. a single winning mapping. Given a subsequent task, the standards and choices are used as probes for remindings from interim pool. If an interim generalization is retrieved as the reminding for a standard (or) a choice, it is applied by filtering out all of the statements in the standard or the choice that do not align with the contents of the interim generalization. Thus, the interim generalization provides a means of highlighting what is likely to be important, by filtering out statements, that from recent experience were known not to have contributed for success in the task.

For word extension forced choice tasks, the labels are maintained as part of the generalizations in the interim pool and is used as a filter for retrieval. For example, if a new stimulus is presented with the label 'Toma', only interim generalizations corresponding to 'Toma' are retrieved.

3.1.3. Abstract & Augment Commonalities

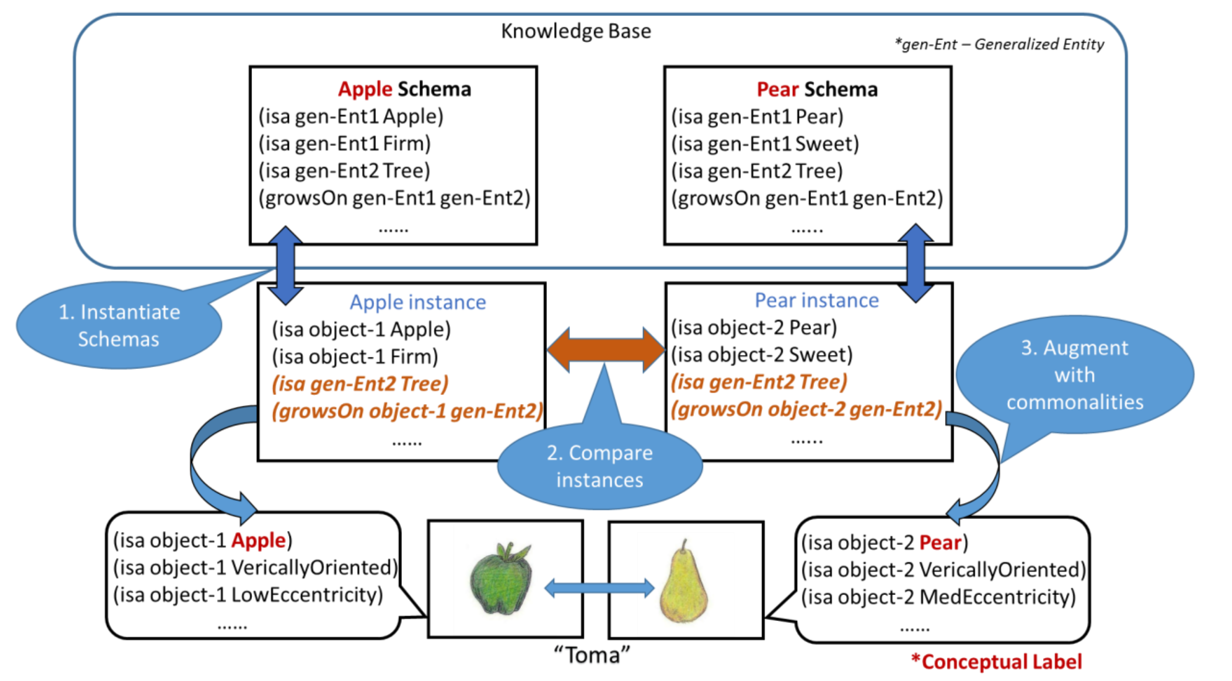

Figure 9: An example of conceptual augmentation through schema comparison.

When two things are labeled the same, it increases the tendency for people to compare them. Young children might require an explicit invitation to compare (e.g. "can you see why these both are tomas?"). This comparison process increases the odds of an interim generalization being created. We model this by creating an interim pool for the label and adding descriptions of items that share the label into the pool. The assimilation threshold of the pool specifies how similar two things should be, judged via structure mapping, for being generalized. We model the increased tendency to generalize by lowering the assimilation threshold of the interim pool.

Additionally, when the descriptions have a Cyc conceptual label, the generalization is enriched as follows. The background knowledge about concepts are approximated by formally encoding material from a children's dictionary (Wordsmyth Children's English Dictionary & Thesaurus) and stored in the knowledge base. The relation `synonymousExternalConcept` is used to connect the Cyc concept to the generalization corresponding to the concept. `LearnedSchemaSource` indicates that the generalization is created based on a text source (dictionary). For example: a schema is created for the Cyc concept Bicycle based on the dictionary definition and is connected to the concept using the proposition below.

```
(synonymousExternalConcept Bicycle LearnedSchemaSource (ConceptSchemaFn
Bicycle)))
```

The schemas are encoded as SAGE persistent generalizations (Appendix A), so as to model knowledge assimilated in long term memory. We compare the descriptions to acquire entities that correspond to each other in the mapping. The schemas associated with the conceptual labels of the entities are instantiated for each description and compared using SME. The expression correspondences in the mapping represents the conceptual commonalities. Before generalization, each description is augmented with the commonalities, as illustrated in Figure 9. Hence the generalization is enriched using background conceptual knowledge.

3.1.4. Forced choice comparison

The similarity score computed by SME is used to ascertain which choice should be selected, i.e. the higher similarity choice should be selected. The default score computed by SME is unnormalized, and thus it also depends on the relative size of the compared items. To remove this dependence on size, we normalize the SME score. We use base-normalized scores, i.e. we divide the similarity score by the score that would be computed by comparing the base to itself. This provides a measure of relative structural overlap between the choices and the standard.

Given: structural evaluation score between a base b and target t = M (b, t).

$$\textit{base normalized score} = \frac{M(b,t)}{M(b,b)}$$

The forced choice tasks modeled here did not involve feedback i.e. The experimenter in the psychological studies did not indicate to the children if the choice they selected is right or wrong. Likewise, the model does not receive any feedback about the choice it made. To utilize recent experiences for representational change, it is imperative to be able to discriminate between positive and negative experiences. If there is a clear difference in the base normalized scores, the experience is considered as positive as the model is able to select the choice with the higher score. Otherwise, if the scores are equal, the model arrives at an impasse and attempts to resolve the impasse using re-representation. When the model is able to make a clear decision, it creates a generalization to encode the winning mapping and adds it to the interim pool.

3.1.5. Verify candidate inferences

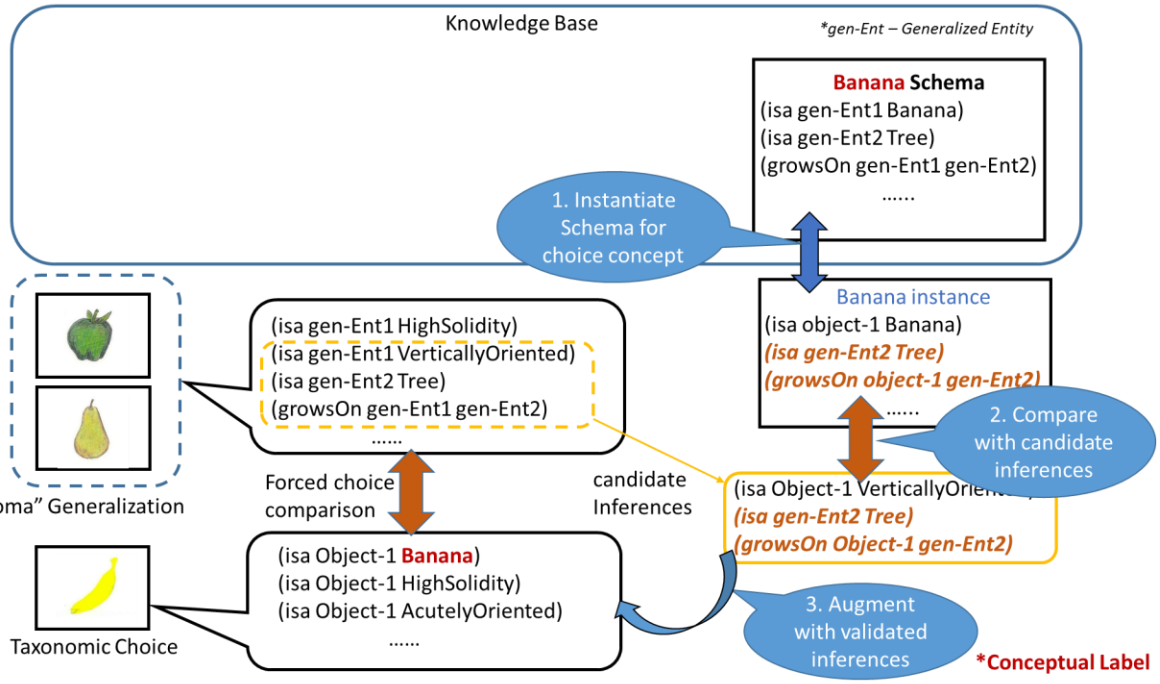

Figure 10: An example of candidate inference validation.

Candidate inferences can help provide more information about the stimuli, potentially increasing the differences between similarity scores to make a decision clear. Candidate inferences are conjectures, whose correctness must be verified before they are accepted. For comparisons involving familiar concepts, their schemas are utilized for verifying candidate inferences.

In a forced choice task, the standard is compared to each of the choices. The comparisons may result in the projection of candidate inferences. When the entities in the choice have Cyc conceptual labels, the corresponding schemas are instantiated and compared to the candidate inferences using SME. The expression correspondences in the resultant mapping represent validated inferences and are used to augment the target of the comparison i.e. the corresponding choice, as illustrated in Figure 10.

3.1.6. Re-represent to resolve the impasse

Rerepresentation re-construes parts of compared descriptions in order to improve a match (Yan et al., 2003). It is an important process in adding fluidity to analogical reasoning. Rerepresentation is necessary, given the variability of encoding processes. When the model arrives at an impasse, it uses re-representation to attempt to increase the similarity of one of the mappings, between the standard and the choices. While rerepresentation is important, it should be constrained. Unconstrained rerepresentation could make any description match to any other and is also computationally expensive. Hence, when there are multiple mappings, the model should be diligent about which mapping is chosen as a candidate for rerepresentation.

The forced choice model selects the mapping to re-represent based on differences in average normalized similarity scores. The average-normalization score captures the overall similarity between two descriptions based on both shared and unshared structures. The equation for average normalized score is given below:

Given: structural evaluation score between a base b and target t = M (b, t).

$$\textit{Average- normalize- score} = \frac{2M(b,t)}{M(b,b)+M(t,t)}$$

The primary objective of re-representation is to attempt to bring into alignment the non-aligned statements in the mapping. The structure-mapping theory of rerepresentation (Yan et al., 2003) outlines ways to improve the alignment in a comparison by reconciling statements that did not match because of violating a constraint of structure-mapping theory. Detecting such violations will enable identifying potential candidates (non-matching relations) for re-representation. These are called opportunities, and based on the type of violation the opportunities are categorized into Holes, gulches, rivals and left-overs.

Constraint	Violates parallel connectivity?	Opportunity
Identicality	Yes	Holes
	No	Gulches
One to One	Yes	Rivals
	No	Leftovers

Table 1: Re-representation opportunities (Yan et. Al, 2003)

SME identifies statements in the base that fail to align with statements in the target and, based on structural overlap, projects them as candidate inferences (Cis). When the *reverse-candidate-inference?* flag is set, SME projects non-aligned statements in target as reverse-candidate inferences (RCis). Cis and RCis provides a way to identify potential opportunities. We detect opportunities for re-representation using the candidate and the reverse candidate inferences, as illustrated via the example below.

Base	Target
`(cause (walk John Cave)` `(inside John Cave))`	`(cause (run Jill Chamber)` `(inside Jill Chamber))`
Expression correspondence: `(inside John Cave)` **<->** `(inside Jill Chamber)` **Entity correspondences:** `Jill` **<->** `John, Cave` **<->** `Chamber` **Candidate Inference (Ci):** `(cause (walk Jill Chamber) (inside Jill Chamber))` **Reverse Candidate Inference (RCi):** `(cause (run John Cave) (inside John Cave))`	

Table 2: An example mapping and inferences

The entities in the `cause` statements in both Ci and RCi are in correspondence. Also, both statements have a similar structure. However, they did not align due to violation of the identicality constraint i.e. the arguments `walk` and `run` are non-identical relations. Statements in Cis and RCis, whose entity arguments are in alignment and have similar structure, are paired together as potential opportunities. If the root predicate fails to match, it is identified as a Gulch. Otherwise, if one of the argument predicates did not match, as in the example above, then it is a Hole. Second, we filter the candidates based on the availability of *transformations* or *decompositions* in the knowledge base. Transformations are truth preserving rewrite rules. For example, `(hotterThan Coffee IceCube)` can be rewritten as `(greaterThan (Temperature Coffee) (Temperature IceCube))` without loss of information. we use '*equiv*' (bidirectional implication) statements in KB, for searching for transformations. Decompositions are rewrite rules which might lead to some loss of generality. For example, `(run John Cave)` can be rewritten as `(moveTo John Cave)`, but we might lose the information about the manner in which John moved into the Cave. We search for decompositions using 'implies' statements and 'genlPreds' statements, meta-predicates for stating that one predicate is a generalization of the other. Using genlPreds to find a common ancestor in the predicate hierarchy is similar to the minimal-ascension approach discussed in Falkenhainer (1988).

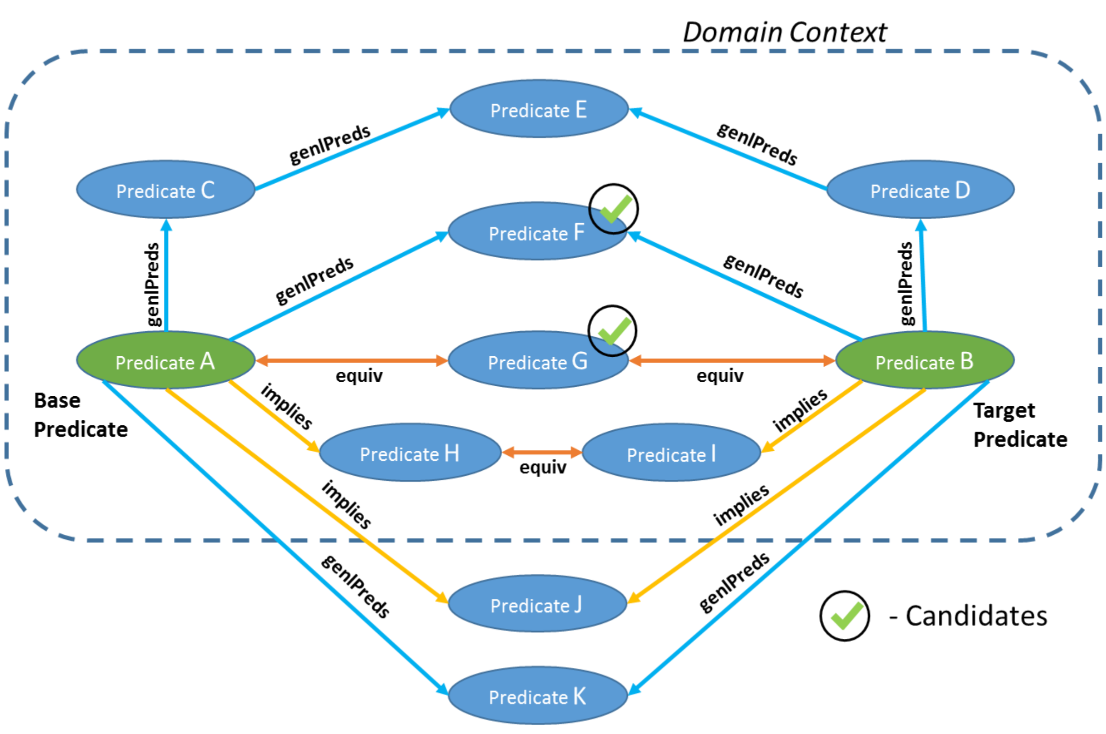

Figure 11: An example of search space for rerep-suggestions for Predicate A and Predicate B

The search space can be visualized as a graph, where the predicates are the nodes and the rewrite-rules are the edges. We constrain the search using context and depth. The search for rewrites are restricted to the current domain context, which is a KB Microtheory and is set when the companion is initialized for solving problems of a particular domain. The rewrite rules correspond to knowledge gained over experience in that domain. For our model, we restricted the search to a depth of 1 for tractability and to limit losing semantic information due to over generalization. We illustrate the search space with an example in Figure 11. The example shows the candidate rewrites for re-representing predicates `predicate A` and `predicate B`. The candidates are chosen based on whether they are within the depth of 1 and whether they are inside the domain context. For example, `predicate E`, `predicate H` and `predicate I`, are within context but are not considered as they exceeded the depth. Likewise, `predicate J` and `predicate K` are within depth, but are filtered out as they are outside the domain context.

The transformations and decompositions that are chosen as candidates are used to derive suggestions for re-representations. For the example in Table 2, assume that there is a general predicate `moveTo` for the predicates `run` and `walk`. The resultant rerep-suggestion is given below.

```
(baseRewrite (walk John Cave) moveTo GenlPredRewrite)

(targetRewrite (run Jill Chamber) moveTo GenlPredRewrite)
```

Finally, the set of <Ci, RCi, rerep-suggestions> for the mapping is returned. Potentially, there could be multiple suggestions for a non-aligned statement pair. We use the simplest strategy for the forced choice model and apply the first rewrite suggestion that could bring the statements into alignment.

After re-representation, the mapping is performed again i.e. between the re-represented base (standard) and the re-represented target (one of the choices). If the base-normalized score of the mapping is higher than the alternative, the model selects the choice as the winning choice, as shown in Figure 8. Otherwise, if there are no re-representation suggestions or if the score did not improve, the model declares failure, unlike the children, who always had to make a choice. The intent is to capture situations where impasses cannot be resolved using re-representations.

The re-representation system as described here is not intended to be complete, for example it does not support detecting rivals and leftovers. It also does not support re-representation mechanisms like entity-splitting or entity-merging. Efficiency and completeness of re-representation is not under primary evaluation in any of the contributions. Caching representational knowledge or employing a more sophisticated search might allow for a better re-representation system, but that is outside the scope of this thesis.

3.2 General analysis of the Relational Match Forced Choice Task

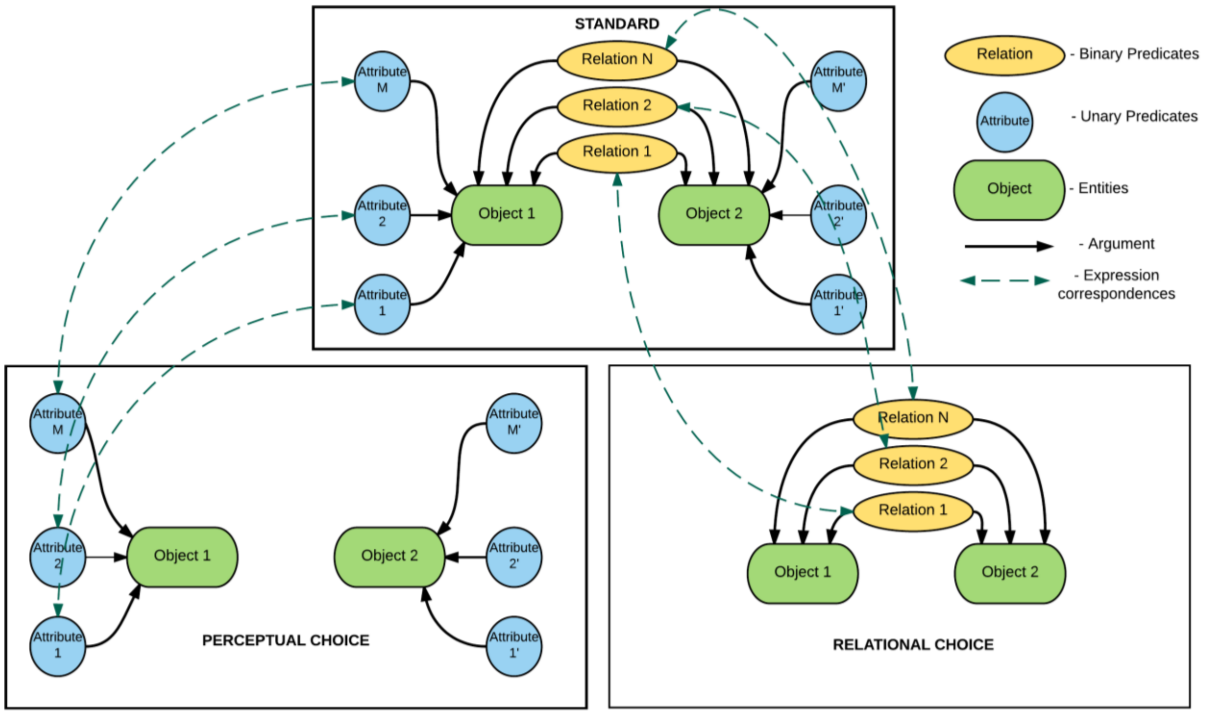

Figure 12: A simplified template of relational match forced choice task.

A general template of the relational match forced choice task is shown in Figure 12. The task has one standard and two choices. Both choices can share attributes with the standard, and even some relations, but there is at least one additional relationship shared between the standard and one of the two choices which will make the match different enough to make the correct choice discernable. How might the outcomes vary depending on the relative sizes of the relational versus attribute overlap?

To analyze this, we use synthetic representations, i.e. arbitrary predicates. For simplicity, we assume that the perceptual choice shares no relations with the standard and the relational choice shares no attributes with the standard. We vary the number of shared relations for the relational choice from zero to 50, and the number of shared attributes with the perceptual choice likewise from zero to 50.

As the plot in Figure 13 shows, shared relations have more influence on the decision than shared

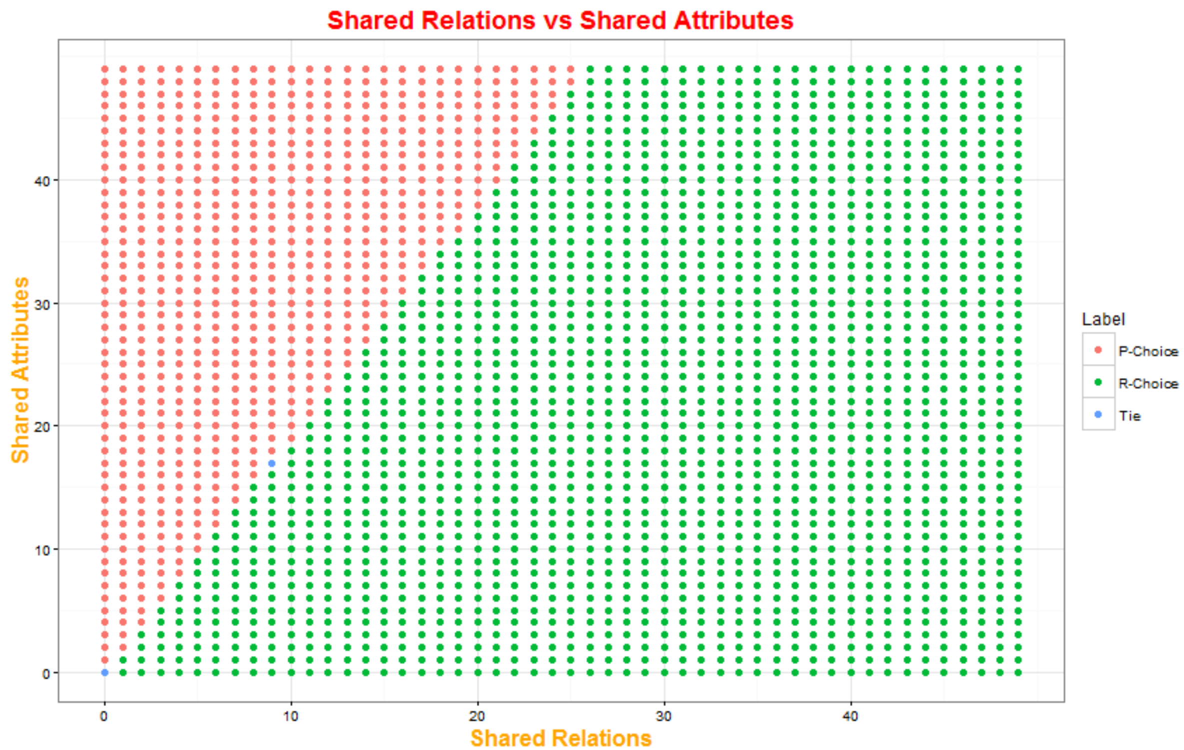

Figure 13: Response pattern of forced choice model.

attributes, leading to a higher relational response. This is not surprising, since SME's trickle-down method of implementing systematicity increases the score of each item below it, and each attribute contributes additional score to only one object, whereas first-order relations (which were all that were used here) contribute score to two or more objects. We assume that human representations typically contain more attribute information than relational information (the Specificity Conjecture, Forbus & Gentner, 1989), although the balance changes with expertise (and with age, which we consider to be gaining expertise across a wide range of domains) to include more relations. Thus for children, we would expect them to be more on the left-hand side of the plot. Thus, independent of the specific representational choices that are made, this model plus the Specificity Conjecture can explain the prevalence of perceptual choices with children (or the inexperienced in a domain) over relational choices in such tasks.

How then adults (and experts) respond more relationally in the forced choice tasks? The representational change mechanisms explained in the model could be part of a broader repertoire of representational learning which is responsible for adult/expert relational responding.

The next chapter describes how this model of forced choice tasks has been used to simulate three psychology experiments. The model initially responds more like the children (i.e. the left-hand side of the plot), but can learn via experience and language to respond relationally and move towards the right-hand (relational) side of the plot.

Chapter 4: Simulations

Young children are prone to focus on object matches rather than relational matches, while adults tend to respond more relationally. The Relational Shift hypothesis (Gentner & Rattermann, 1991) suggests that this difference is due to a lack of knowledge about relational structures in younger children, and that as they learn more, they gain the ability to perform more relational matches. Gentner and her colleagues explored two forces that could be driving relational shift: relational language and progressive alignment.

Language has a mutually facilitating partnership with relational representation and reasoning. Learning symbols and terms for relations, i.e. relational language, substantially augments our relational ability (Gentner & Christie, 2010). As children, we acquire a variety of relations, including spatial relations such as above, and on, and functional relations like edible and dangerous. Common labels help bring focus on to the common relational structure and thus fine-tune representations. Gentner (2003) has argued that common labels trigger comparison and thus promotes learning new relational abstractions via structure mapping. Another way in which representations are improved is via progressive alignment, where experience with sequences of highly similar examples lead to rapid learning by abstraction of common relational structure and thereby enabling recognition of more abstract relational structures.

There are three mechanisms of representational change facilitated by relational language and progressive alignment.

1. Abstraction: The common system resulting from the alignment becomes more salient and more available for future use.
2. Inference-projection: When the comparison involves one description that is richer than another, the spontaneous projection of candidate inferences can enrich the less-complete description. If there is background knowledge available, the inferences can be confirmed and be used to boost the alignment.
3. Re-representation: Re-representation is performed if there is a reason to believe that improving the overall match, could help in resolving the current situation/problem.

In this chapter, we present simulation studies that show how these mechanisms, in the context of our other assumptions about processes and representations, suffice to explain the results of three papers describing human-subjects experiments. The first two psychology studies (Christie & Gentner (2010) and Namy & Gentner (2002)) highlight representational learning facilitated by language. The simulations show how common labels help in the creation of abstractions, including filling in missing pieces by projecting inferences from prior knowledge. The third psychology study (Kotovsky & Gentner (1996)) shows the effects of progressive alignment and re-representation on representational change.

4.1 Christie & Gentner (2010) Simulation

Christie & Gentner (2010) showed that children (ages 3-4) can learn new relational abstractions via shared labels and comparison. They used novel spatial relational categories in a word extension task, as illustrated in Figure 14. Here the relationship might be characterized as “An animal above another identical animal”.

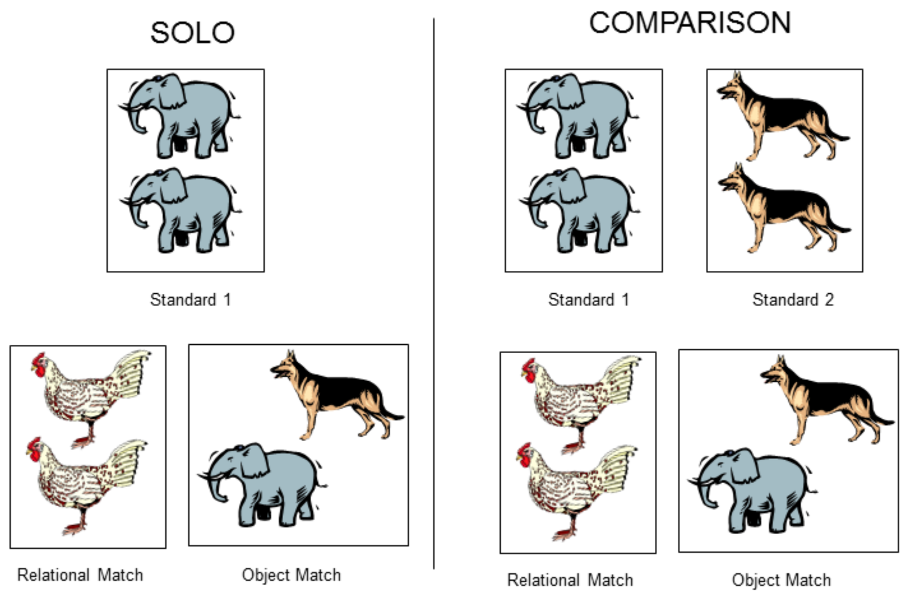

Figure 14: Example of stimulus from Christie & Gentner (2010).

In the Solo condition, children were shown a single standard (here, Standard 1) and told it was a novel noun (e.g. “Look, this is a jiggy! Can you say jiggy?”). In the Comparison condition, children were invited to compare two examples (e.g. “Can you see why these are both jiggies?” when presenting Standard 1 and Standard 2 simultaneously). In both conditions, children were then presented with a forced-choice task, where they had to choose which one of the alternatives is a jiggy (e.g. “Which one of these is a jiggy?” when presented with the relational match and object match cards). Children in the Solo condition preferred the object match, while those in the Comparison condition chose relational matches twice as often as object matches. This provides evidence that comparison can lead to learning new relational abstractions.

In a second experiment, a third condition was added. In the new Sequential condition, children saw two standards one at a time, to test whether or not simple exposure to more examples was sufficient to promote learning. They found significant differences between the Sequential and Comparison conditions, and between the Solo and Comparison conditions, but the difference between the Sequential and Solo conditions were not significant. This provides additional evidence that it is comparison, not just more examples, that is promoting learning.

By simulating the Christie & Gentner (2010) experiments we show that our computational model of forced-choice tasks is sufficient to explain this phenomenon. We describe below two simulation experiments, including sensitivity analyses to shed light on why it does so.

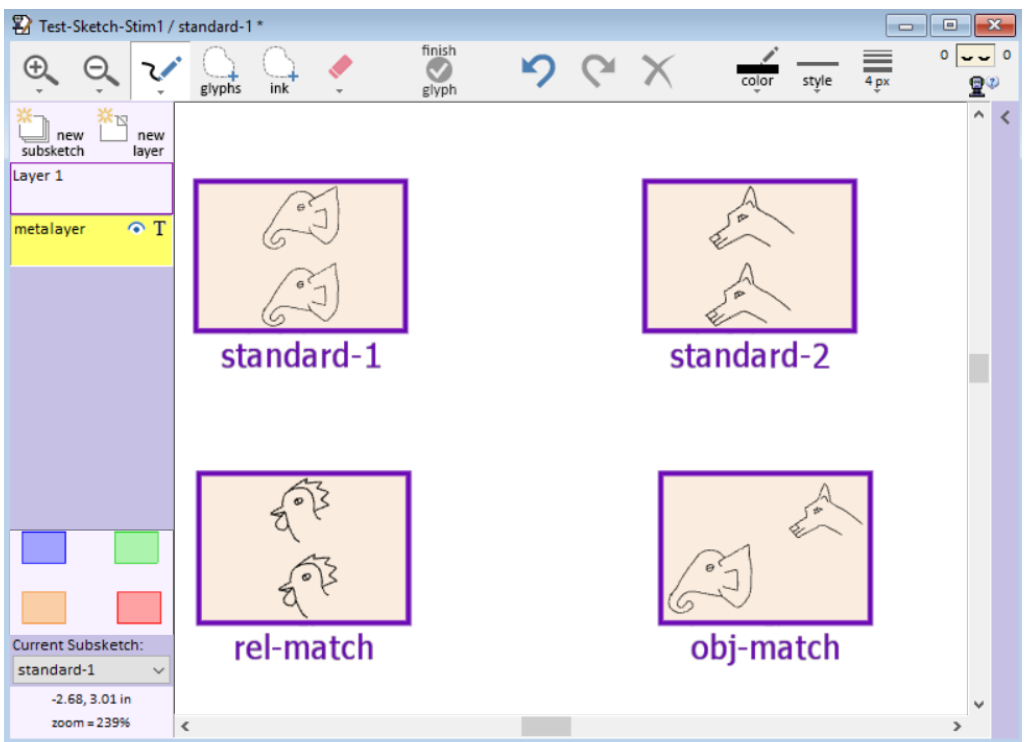

Figure 15: The sketched version of the stimulus in Figure 14, shown here as four subsketches.

4.1.1. Word Learning via Analogical Generalization

The forced choice task model can simulate word learning, as explained in Chapter 3. To recap, for each word, there is a generalization pool. Every time the word is used, an appropriate subset of the world, in our case the subsketch (Figure 15), is encoded to capture information about what that word denotes, and is added to the pool. The generalizations constructed can be considered as the meanings for the words. The ability to track multiple generalizations provides a mechanism for handling multiple senses of a word. The ability to store unassimilated examples provides a means of handling edge cases, and helps provide noise immunity in the face of changes in the underlying distribution of examples of a concept. This generalization-based account has been used to successfully model spatial propositions of contact in English and in Dutch (Lockwood et al., 2008). Recall that SAGE and SAGE-WM create generalizations in essentially the same way. In addition, SAGE-WM has a size limit and the retrieval is biased by recency, neither of which is necessary for Christie & Gentner (2010) simulation. For Christie & Gentner (2010) simulations, the forced choice model uses SAGE for creating relational abstractions instead of SAGE-WM, because these experiments were done before SAGE-WM was implemented.

4.1.2. Simulation Experiment 1

Experiment 1 used two conditions to show children the new concept, followed by a forced-choice task. We model these as follows:

a. **Solo Condition:** The single example is added to the generalization pool for the word and is chosen as the base of the comparison.
b. **Comparison Condition:** The two examples are added to the generalization pool, but since the experimenter has asserted that they are both examples of the labeled concept, we assume that the child is more likely to compare and assimilate them into a

generalization, which, as explained in chapter 3, is modeled by lowering the assimilation threshold from its default of 0.8 to 0.1. We also assume that the probability cutoff is 0.6, so that facts which do not appear in the shared structure will be eliminated from the generalization.

In both conditions, the model compares the base of comparison, the single standard or the generalization, to each of the choices. The choice with the highest based normalized score is chosen as the winning choice.

The original experiment used 8 stimulus sets. We encoded 8 sketches of animals, using CogSketch. Each element of the stimulus set (e.g. Standard 1, Standard 2, etc.) was drawn as a separate subsketch. Filters were used to automatically remove three types of information: Redundant information (e.g. given (rightOf B A), (leftOf A B) is redundant), irrelevant information (e.g., global estimates of glyph size like MediumSizeGlyph), and bookkeeping information (e.g. relationships describing timestamps of glyph creation). Table 3 shows the final encoding for the sketched stimulus set (Figure 15) and the resultant generalization.

<table>
<tr><th>Standard-1</th><th>Standard-2</th></tr>
<tr><td>(sameShapes Object-99 Object-420)
(above Object-99 Object-420)
(isa Object-420 Elephant)
(isa Object-99 Elephant)</td><td>(sameShapes Object-104 Object-425)
(above Object-104 Object-425)
(isa Object-425 Dog)
(isa Object-104 Dog)</td></tr>
<tr><th colspan="2">Generalization for “jiggy”</th></tr>
<tr><td colspan="2">(above (GenEntFn 1 0 jiggy) (GenEntFn 0 0 jiggy))
(sameShapes (GenEntFn 1 0 jiggy) (GenEntFn 0 0 jiggy)</td></tr>
</table>

Table 3: Encoding for sample sketch.

For this simulation, we treated the amount of conceptual and perceptual attributes that the children might encode as an interesting open parameter as we do not know of data that provided specific estimates. Consequently, we performed a sensitivity analysis by running the simulation while varying the number of conceptual attributes to ascertain their impact on the results. Specifically, we varied the number of attributes from zero to nine. We assumed that encoding is reasonably uniform, i.e. that the same attributes would always be computed for identical objects. For simplicity, we further assumed that the set of attributes computed for one entity had no overlap with the set of attributes computed for another entity whose shape is different. Given these assumptions, we used synthetic attributes (e.g. Uniquestandard-1MtAttribute8) for convenience.

Figure 16 shows the results. From the data, we can see that the model chose the relational match 100% of the time for the Comparison condition. This is qualitatively consistent with the behavior of participants in the Comparison condition, where participants chose the relational match around 60% of

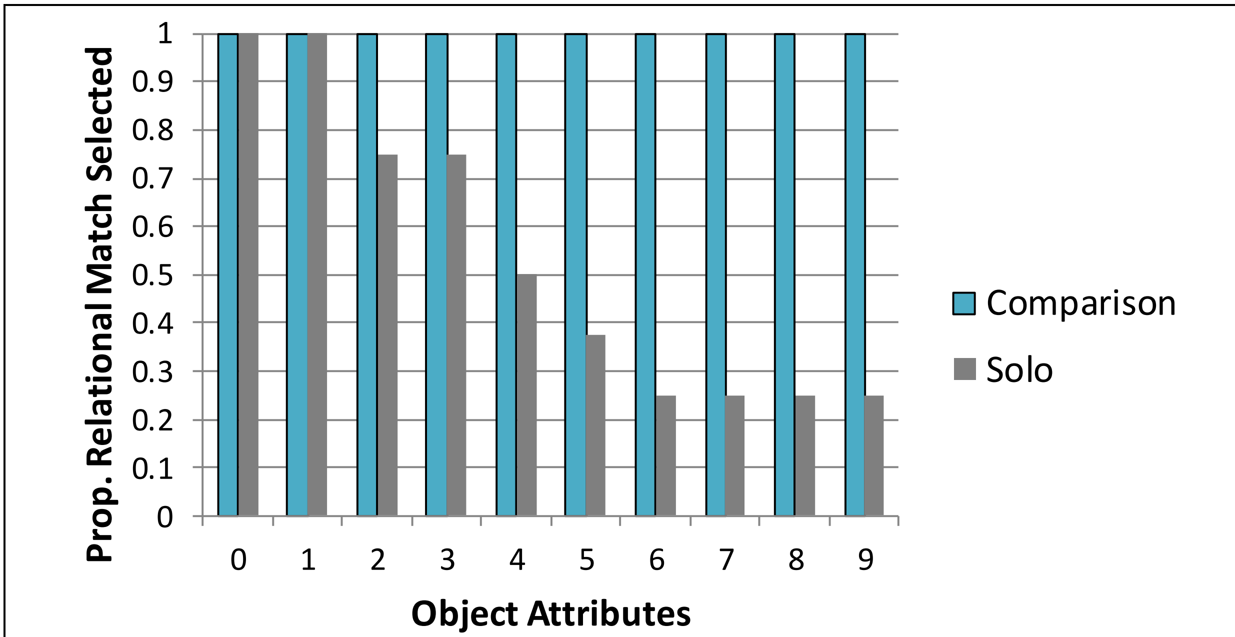

Figure 16: Results of Simulation experiment 1.

the time. We believe that the lack of object matches in this simulation condition are due to the use of completely independent attributes for each entity type in the stimuli sets. Since they are independent, no attributes are left in the generalization after assimilation. The more overlapping attributes there are, the more likely an object match is to become possible.

As depicted in Figure 16, in the Solo condition, as the number of attributes rises, the proportion of object matches rises (i.e., the proportion of relational matches falls). Again, this provides a good qualitative fit for the results of (Christie & Gentner 2010) Experiment 1. Since attributes are more salient to children, due to lack of relevant domain knowledge (Rattermann & Gentner, 1998), it is reasonable to assume that they would encode more attributes than relations, which is compatible with the simulation results.

Recall that we assume that the probability cutoff is set high enough that non-overlapping information is immediately filtered out. (Since these are novel concepts, there can be at most two examples in any generalization, and hence the probability of any fact not in the overlap would be 0.5, which is less than the 0.6 threshold.) Would adding in probabilistic information improve the fit of the model to human data? To determine this, we tried changing the probability cutoff to its usual default of 0.2. This leads to all attributes remaining in the generalization, which results in the score for the object match being boosted so high that it always wins over the relational match, regardless of the experimental condition used. This suggests that when children are invited to compare, they do indeed restrict themselves to keeping exactly the overlapping structure.

4.1.3. Simulation Experiment 2

Experiment 2 in (Christie & Gentner 2010) actually consists of two experiments. Both involved a new condition, the Sequential condition, designed to rule out non-comparison explanations. In Experiment

2a, fillers, in the form of pictures of familiar objects, were interposed between the serial presentation of the standards. No invitation to compare was issued. In Experiment 2b, no fillers were used, and the Solo and Comparison conditions from Experiment 1 were added, by way of replication. In our model, fillers would be added to some other generalization pool, thus 2a and 2b look identical from the perspective of our model.

For the Sequential Condition, the two examples are added to the generalization pool, but with the default assimilation threshold 0.8. Again we varied the number of conceptual attributes, in the same way as in Simulation Experiment 1.

Figure 17 illustrates the results. As anticipated, the results for the Sequential condition are similar to the results the model generates for the Solo condition. This is because of the model does not generalize the

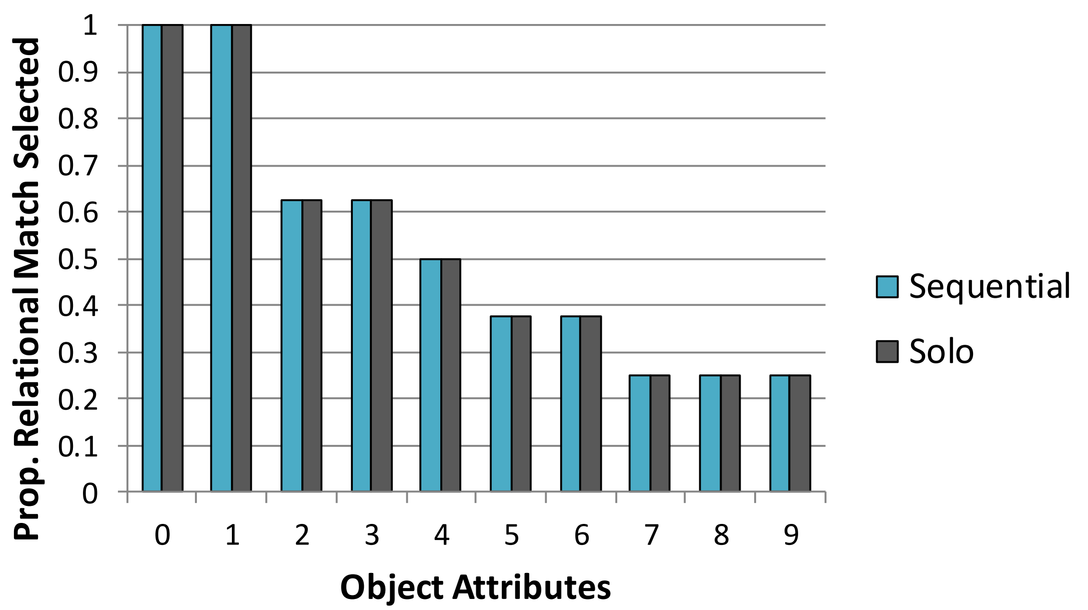

Figure 17: Results of Simulation Experiment 2.

two standards, and hence the choices will be compared to the examples in the generalization pool. This makes the results of the Sequential Comparison condition be the same as the Solo condition.

We know of no direct psychological evidence that would provide constraints on the value of the assimilation threshold. Consequently, we performed a sensitivity analysis by varying the assimilation threshold between 0.1 and 0.9, while varying the number of attributes from zero to nine. Figure 18 illustrates the results. The region marked as black indicates a high proportion of relational match choices and then the contour fades down gradually.

The slope of the contour indicates that the model readily generalizes the standards when both the assimilation threshold and the number of object attributes are low. This can be interpreted as follows. A low assimilation threshold corresponds to a higher willingness to accept the standards as belonging to the same category, which fits the assumptions of our model. A low number of object attributes indicates a leaner encoding i.e. not enough attention was paid to the object, or it may be unfamiliar. This is a second possible explanation for why some children chose the relational match for the Sequential condition.

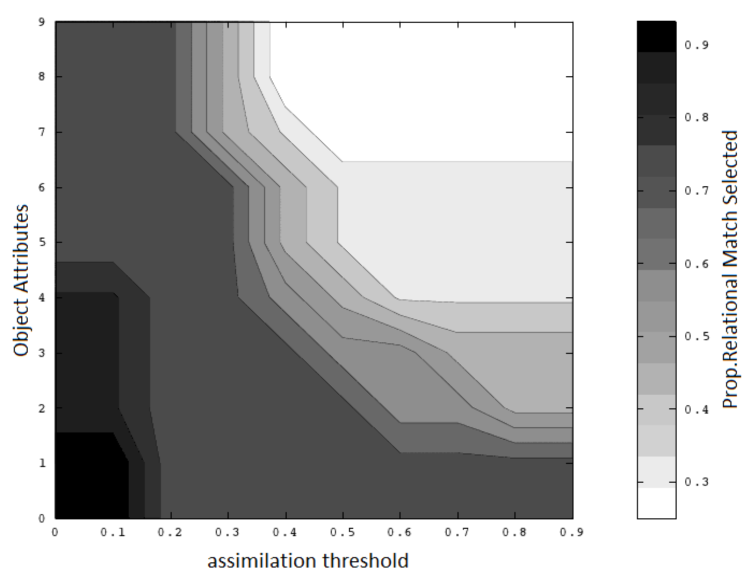

Figure 18: Sensitivity Analysis.

We have shown that a model of forced choice task can simulate the behavior found in (Christie & Gentner 2010). The invitation to compare, we argue, leads the child to aggressively attempt to form a generalization between the relational learning, as measured by responses in the forced choice task.

4.2 Namy & Gentner (2002) Simulation

Research on how children learn novel categories have resulted in a conflicting set of findings. On one hand, studies show that children's categorization reflect understanding of deep conceptual properties. On the other, there are results that suggest that children tend to rely primarily on perceptual similarity for categorization (Namy & Gentner, 2002).

Categorization requires grouping and generalizing objects that are perceived to be similar. An object/scenario will be categorized differently according to how it is represented. Subsequently, the conflicting findings might be explained by differences in the way children are encoding, perhaps due to the context and requirements of the task. The Namy & Gentner (2002) studies reconcile the conflict by means of structure mapping theory and provides insight into how representations can change based on the purpose of the comparison.

As shown earlier, common labels invite structural alignment, in order to help the learner understand what makes them the same. In addition to emphasizing common structure, inducing comparison via common labels enables accessing conceptual commonalities via access to long-term memory that may not have been evident prior to structural alignment. In other words, the process of aligning perceptual and lower-order relational information can give rise to the augmentation of representations with higher-order relational commonalities.

Namy & Gentner (2002) tested their hypotheses in two experiments. In both experiments, they had 4 year olds perform forced choice tasks involving simple daily objects (Figure 19). The first experiment was designed to provide evidence for structural alignment as the driving mechanism behind elicitation of conceptual commonalities.

The first experiment had two conditions. In the One-Kind condition, children were shown two standards from the same taxonomic category, e.g. two pieces of fruit. In the Two-Kind condition, the standards were from different categories. In both conditions, the standards were introduced with the same label

Figure 19: An example of stimulus from Namy & Gentner (2002).

and the children were invited to compare them (e.g. “This is a blicket and this is also a blicket. See how these are both blickets?”). Children were then presented with a forced choice task, where they had to choose from the alternatives (e.g. “can you tell me which one of these is a blicket?”). One of the choices was perceptually similar to both the standards. The other choice was not perceptually similar, but belonged to the same taxonomic category of at least one of the standards.

Children from One-Kind condition preferred the category/taxonomic choice more often than the children in Two-Kind condition. This provides evidence that structure-mapping comparison can play a role perceiving conceptual commonality, thereby enriching children's perceptual-centric representation with more conceptual information.

In Experiment 2, Namy & Gentner (2002) investigated the role of common labels vs different labels in structural alignment and in seeking out conceptual commonalities. The stimuli from the One-Kind condition were used. Experiment 2 had two conditions. For the Unifying-Word condition, the standards were labeled the same. For the Conflicting-Word condition the standards were assigned different labels. The results showed that the children who received standards with the same labels favored the category choice more often than the children who received conflicting labels.

In summary, the Namy & Gentner (2002) results show that common labels invite structural alignment, which facilitates the perception of higher order relational commonality. We simulate the forced-choice task experiments from Namy & Gentner (2002) using our model and demonstrate that the model is capable of exhibiting behavior consistent with their results.

4.2.1. Conceptual Augmentation and Candidate Inference Validation

Here we recap two important aspects of the forced choice model for simulating Namy & Gentner (2002) experiments. First, when familiar objects are introduced with the same label and are compared, we instantiate their schemas, compare the instances and augment the compared cases with commonalities found from the mapping. Second, when the choices are familiar objects, we use their schemas to validate candidate inferences of the forced choice comparisons. That is, the schema is instantiated and the instance is compared to the candidate inferences. The statements aligned in the mapping represents valid inferences and are used to augment the target (choice) case. Chapter 3 explains the processes in more detail. The schemas are represented as SAGE persistent generalizations derived from contents of a children’s dictionary and from Wiktionary, for two of the items not found in the dictionary. See Appendix A for examples.

4.2.2. Simulation Experiment 1

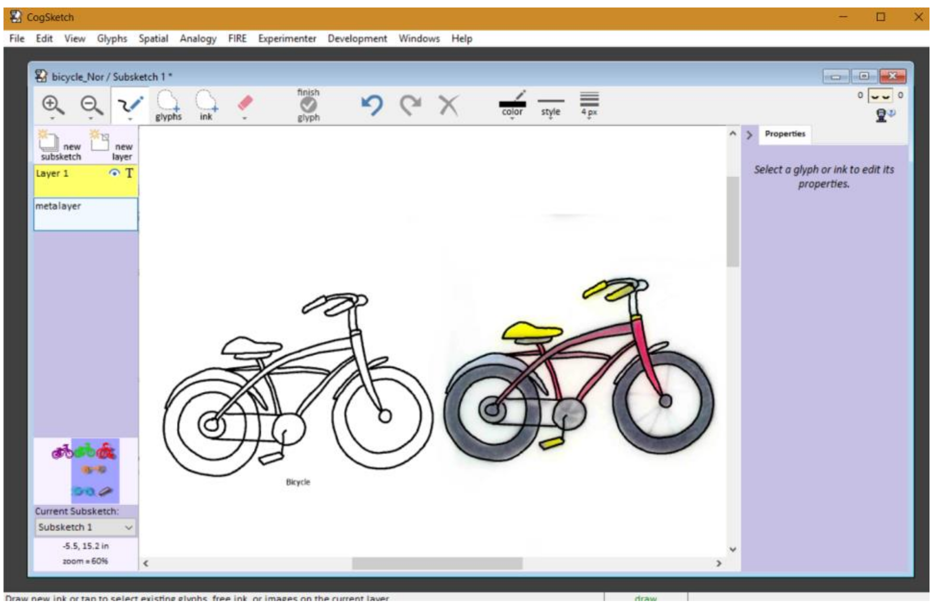

Figure 20: An example of a traced version of an object from Namy & Gentner (2002) stimuli.

The experiment used 10 sets of stimuli. Each of the original stimuli consisted of 5 color drawings of real-world objects. We encoded the stimulus sets using CogSketch. Each stimulus was a single sketch with five subsketches, one per drawn object. CogSketch allows importing images as backgrounds. We used this feature to import the scanned version of the original stimuli. Six Northwestern graduate students contributed by tracing over the images to create digital ink (Figure 20 illustrates). As per the above, this simulation experiment contains two conditions: One-Kind condition and Two-kind condition. We model these as follows.

One-Kind Condition: Descriptions of the two standards are compared and the entity correspondences are used to identify seeds, i.e. conceptually labeled objects, for augmentation. The schemas corresponding to the conceptual labels are instantiated i.e. appropriate substitutions are made in order to create an instance, as shown in the Table 4.

Schema: `(ConceptSchemaFn Belt-Clothing)`

`(isa (GenEntFn 0 0 (ConceptSchemaFn Belt-Clothing)) Belt-Clothing)`
`(isa (GenEntFn 1 0 (ConceptSchemaFn Belt-Clothing)) Waist)`
`(isa (GenEntFn 2 0 (ConceptSchemaFn Belt-Clothing)) Leather)`
`(wornOn (GenEntFn 0 0 (ConceptSchemaFn Belt-Clothing))`
`        (GenEntFn 1 0 (ConceptSchemaFn Belt-Clothing)))`
`(mainConstituent (GenEntFn 0 0 (ConceptSchemaFn Belt-Clothing))`
`                 (GenEntFn 2 0 (ConceptSchemaFn Belt-Clothing)))`

Instance created for seed: `(isa Object-1 Belt-Clothing)`

`(isa` **`Object-1`** `Belt-Clothing)`
`(isa (GenEntFn 1 0 (ConceptSchemaFn Belt-Clothing)) Waist)`
`(isa (GenEntFn 2 0 (ConceptSchemaFn Belt-Clothing)) Leather)`
`(wornOn` **`Object-1`** `(GenEntFn 1 0 (ConceptSchemaFn Belt-Clothing)))`
`(mainConstituent` **`Object-1`**
`                 (GenEntFn 2 0 (ConceptSchemaFn Belt-Clothing)))`

Table 4: An example of a schema and its instance, given seed.

The instances are compared and the descriptions of the standards are augmented with the commonalities. An interim pool for the label is created and the augmented descriptions are added. As the comparison involves a common label, the assimilation-threshold is lowered to 0.1. As a result, a generalization is created, that captures both perceptual and conceptual commonalities of the standards. As in Christie & Gentner (2010) simulation, the probability cutoff is set to 0.6 for eliminating facts that do not appear in the shared structure from the generalization.

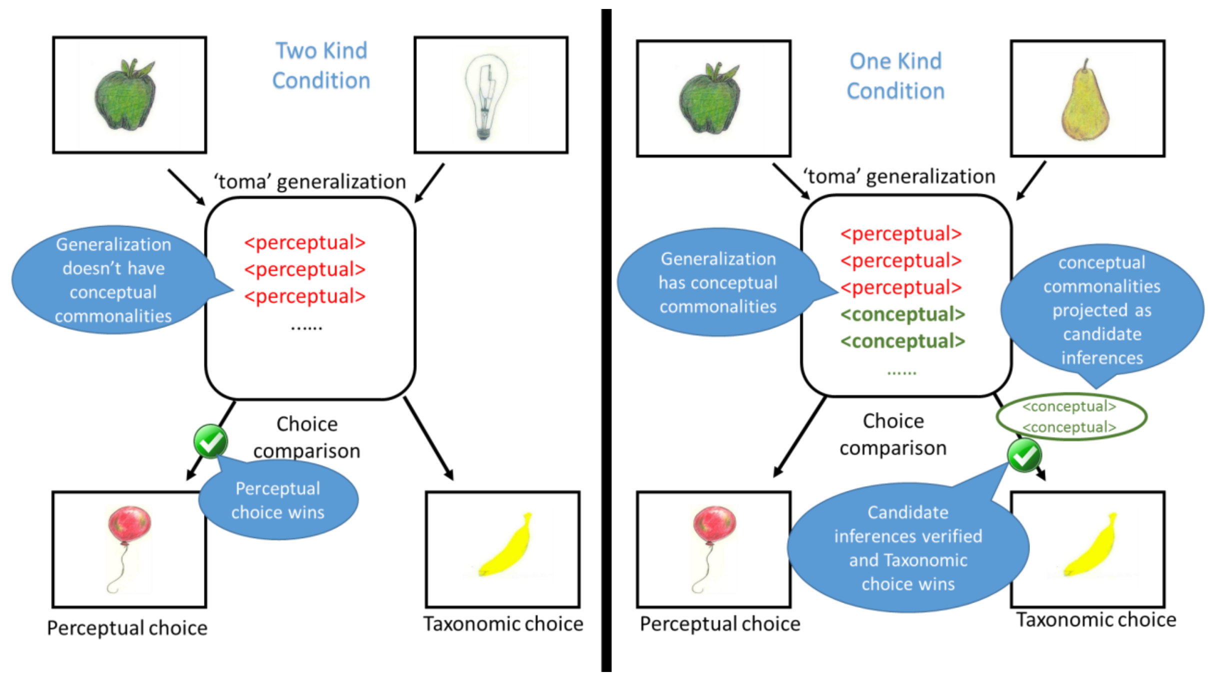

Figure 21: Simulation of One-Kind and Two-Kind condition.

The model compares the generalization to the choices. The augmented conceptual information in the generalization is projected as candidate inferences, as illustrated in Figure 21. Recall that the taxonomic choice belongs to the same category as the standards, which is reflected by its schema. Owing to that, most of the candidate inferences can be verified for the taxonomic choice, but not for the perceptual choice. This results in boosting the base-normalized score of the taxonomic-choice. Accordingly, the taxonomic choice is chosen as the winning choice 90% of the time for the One-Kind condition.

Two-Kind condition: The process is same as in the One-Kind condition. The schemas corresponding to the concepts are instantiated and compared. However, the standards are from different taxonomic categories. Consequently, the mapping between the schema instances does not have many aligned statements, indicating that the standards have low conceptual commonalities. As a result, a generalization is created that captures the perceptual commonalities but with an insignificant amount of conceptual content, as there were not many conceptual commonalities.

	Perceptual Choice win	Taxonomic Choice Win	Tie
One-Kind condition	10%	90%	0%
Two-Kind condition	80%	0%	20%

Table 5: Namy & Gentner (2002) Experiment 1 Simulation Results.

The model compares the generalization to the choices. But unlike the One-Kind condition, most of the candidate inferences are cannot be verified via schema instantiation. This results in the perceptual commonalities dominating the match and hence the perceptual choice wins 80% of the time. Table 5 shows the results.

4.2.3. Simulation Experiment 2

Experiment 2 uses the stimuli from Experiment 1 One-Kind condition. We simulated the conflicting label condition of Experiment 2 as follows. When same label invites alignment, differing labels could have an opposite effect. This might result in the children focusing on one standard while ignoring the other, or the children might engage in alignment but terminate it at an early stage (Namy & Gentner 2010). Consequently, they would not have enhanced the generalization with conceptual commonalities. We decided to test both possibilities via our simulation.

First, we simulated the possibility where the children might be focusing on only one of the standards. We expect this to be the default behavior of the model i.e. in absence of a common label, there will be different interim pools for different labels. The simulation was run with using only one of the standards as the base of the comparison. We tested both standards, i.e. for each stimulus we chose the first standard as the base for one experiment run and the second standard as the base for another run. The results are shown in Table 6. The model selected the taxonomic choice less than 10% of the time.

Second, we simulated the possibility where the children might engage in alignment and may even generalize, but terminate early. We modeled this by lowering the assimilation-threshold, but switching off the conceptual augmentation capability of the model. So the standards were generalized, but not conceptually augmented using schema comparisons. Hence the resultant generalization will only capture the perceptual commonalities between the standards. As expected, the results are similar to the Two-Kind condition of experiment 1. The model chose the perceptual choice 80% of the time. (Table 6)

	Perceptual Choice Win	Taxonomic Choice Win	Tie
Standard-1 as Base	100%	0%	0%
Standard-2 as Base	60%	10%	30%
Generalization (without conceptual augmentation) as Base	80%	0%	20%

Table 6: Namy & Gentner (2002) Experiment 2 Simulation Results.

In both cases, the results are as expected i.e. in absence of common label, the model favors perceptual choice more than the taxonomic choice. The results do not help us in figuring out which one of the possibilities may have happened in absence of a common label. Nevertheless, we have shown that our model is consistent for either possibility and responds similarly to the children i.e. selects perceptual choice most of the time for the Experiment 2 conflicting label condition.

The last two simulations showed us how common labels help in acquisition and augmentation of relational abstractions, and in turn improve representations. In the next simulation, we show that even in the absence of labels, representational learning can happen within a short duration via progressive alignment.

4.3 Kotovsky & Gentner (1996) Simulation

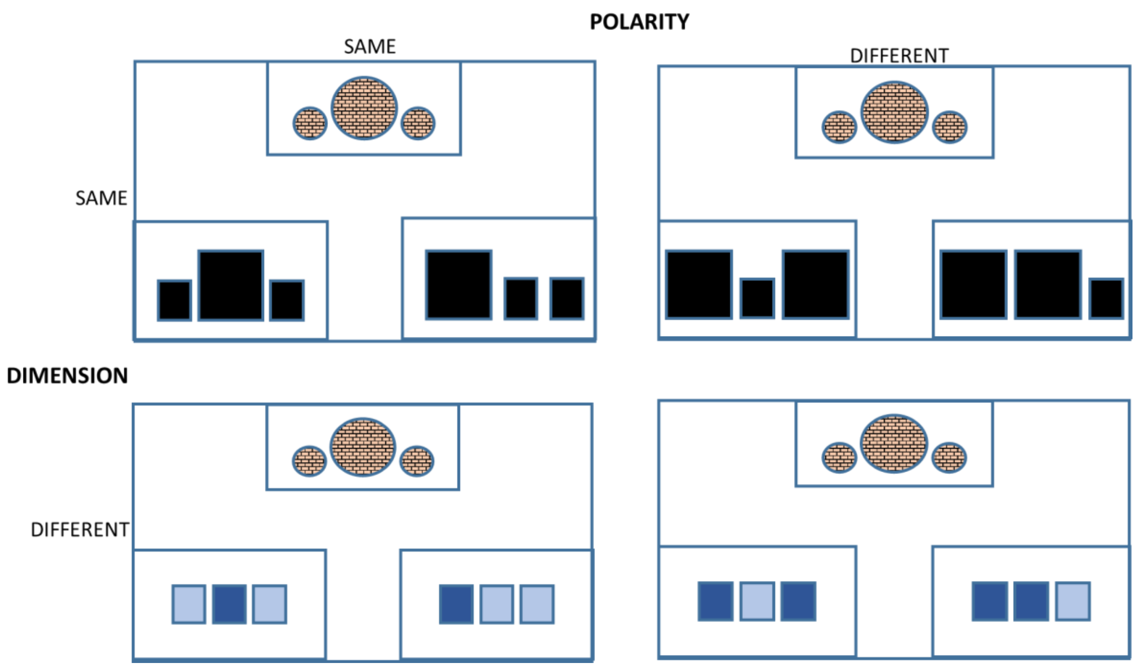

Figure 22: An example of four types of triads for a size symmetry standard.

Kotovsky & Gentner (1996) explored children's performance on comparison tasks involving simple higher-order patterns, such as symmetry and monotonic increase (Figure 22). In each triad, the top figure is the standard, and the bottom two figures are the choices from which a participant must pick. One choice always has the same higher-order relationship between its entities as does the standard, while the other has the same entities as the relational choice, but permuted so that the relationship does not apply. The triads in Figure 22 illustrate the 2x2 manipulation, namely the polarity (same or opposite) of the higher-order relation and the dimension (size or brightness) over which the relationship holds Children were asked to choose which one of bottom choices was most like the top one. No feedback was given at any time. However, some easy high similarity triads were provided as check trials.

The Relational Shift hypothesis predicts that older children will do better than younger children, and that all children will do better when there are lower-order commonalities supporting the higher-order commonalities. The results were consistent with these predictions: 4 year olds performed below chance on all but the same dimension/same polarity stimuli, where they were above chance. By contrast, 6-year-old and 8-year-old children were able to see the relational pattern to some degree without the support of first-order relational overlap, but better with it. The cross dimension/opposite polarity case was the hardest condition, even for eight year olds. Yet some children discovered this match over the course of the study. As Kotovsky & Gentner (1996) remark:

"The emerging appreciation of relational commonality can be seen in this comment by an eight year old, who after struggling with her first several cross-dimension matches, then excitedly articulated a startlingly apt description of relational similarity: "It's exactly the same, but different!" She proceeded to choose relationally for all the remaining triads"

Dimension	Dimension of Standard	High Order Relation
same	size	monotonic-increase
same	size	symmetry
same	color	symmetry
same	color	monotonic-increase
cross	size	symmetry
cross	color	symmetry
cross	size	monotonic-increase
cross	color	monotonic-increase

Table 7: Order of triad pairs in progressive alignment condition in Kotovsky & Gentner 1996.

How can we explain such learning within less than 20 trials, without feedback? It requires that a child be able to detect that they do not know a good answer. There is informal evidence for this in that children in the study often puzzled over the cross-dimensional triads, saying things like “A dark one and a big one make daddies. The other one has two twins and a daddy on the side.” Children further need to figure out ways to re-represent the stimuli so that the choice becomes clear. This re-representation process is aided by the experience of comparing and aligning relational structure across trials, as Kotovsky and Gentner showed in a second study.

In that study, 4-year-olds were given a progressive alignment sequence, as shown in Table 7: first 8 same-dimension (and same polarity) triads, which were relatively easy to align; and then 8 cross-dimension triads (also same-polarity). A control group received 8 initial size-change triads (so that they did not experience easy alignments over the saturation dimension. The progressive alignment group performed better on the subsequent cross-dimensional triads than did the control group). This suggests that successfully aligning the same-dimension triads led children to see the higher order patterns that they had formerly missed—that is, to re-represent the stimuli. We simulate Experiments 1 & 2 of Kotovsky & Gentner (1996) as follows.

4.3.1. Simulation Experiment 1

Recall that the four types of triads (ordered in terms of predicted difficulty) are:

1. Same dimension/same polarity (SDSP)
2. Same dimension/different polarity (SDDP)
3. Different dimension/same polarity (DDSP)
4. Different dimension/different polarity (DDDP)

We created two ordered sets of 16 triads grouped by polarity, shuffled so that there would be no more than two of the same triad types consecutively, as in Experiment 1 of Kotovsky & Gentner (1996). In particular, as in that study, same-dimension triads (like the top left triad in Figure 22) and cross-dimension triads (bottom left, Figure 22) were mixed semi-randomly across the study.

We evaluated our model on the two sets. The model performs the triads task sequentially following the determined order. The model uses three parameters. The assimilation threshold (0.95) is used by SAGE WM to determine when to assimilate winning mappings into generalization. It is also applied during the reminding phase to choose the most similar generalization. The re-representation threshold (0.55) controls when a mapping between a base and a target looks promising enough to attempt re-representation. The size-limit (5) determines the maximum number of items in the SAGE-WM interim generalization pool.

When the model has no clear choice, it does not make a decision, unlike the children, who always had to make a choice. (Importantly, the children were not given feedback as to whether their choices were correct or not.) The Kotovsky & Gentner experiments measured the proportion of relational responses. Table 8 shows the results for four year olds along with the model's responses. As noted above, the correct choice is always the relational choice, so the children were above chance only for the SDSP case.

	Children Relational Response %	Model Relational Response %	Model Non-Relational Response%	Model No-choice
SDSP	68%	100%	0%	0%
SDDP	49%	0%	87.5%	12.5%
DDSP	49%	37.5%	0%	62.5%
DDDP	48%	12.5%	12.5%	75%

Table 8: Proportions of choice types for Kotovsky & Gentner (1996) Experiment 1 Simulation.

The results of the model are qualitatively consistent with the children's behavior. First, the SDSP cases are easiest. The model gets 100% of these correct because the automatic encoding process, using

CogSketch, is deterministic and uniform, whereas children (68% correct) are likely to vary more in their encodings. Second, when the no-choice model answers are randomly distributed between the two possible choices, the model is at chance for DDDP, somewhat better than chance for DDSP, and far worse than chance for SDDP.

In the SDDP stimuli, there is sufficient relational overlap between even a non-relational standard to make the base-normalized comparison scores different enough to satisfy the system that it has a reasonable answer. We suspect that increasing the required difference in similarity between the two alternatives would eliminate this behavior. In the DDSP case, while the same dimension triads were not consecutive, they were sometimes close enough that occasionally interim generalizations were getting created. This suggests that our model can form interim generalizations a bit more readily than children do.

4.3.2. Simulation Experiment 2

Experiment 2 was designed to test the Progressive Alignment hypothesis, i.e. that children who first received highly similar (i.e., highly alignable) closely spaced trials could then do tasks that were beyond them previously. The stimuli consisted of only same polarity triads. There were two conditions.

1. **Experimental condition:** Eight same dimension triads followed by eight cross dimension triads. The same dimension triads consisted of both saturation-change and size-change triads. To encourage progressive alignment, the triads were ordered as shown in Table 7. The children received two of each type.
2. **Control condition:** Same as in the progressive alignment condition, but (as in the Kotovsky & Gentner 1996 study) the eight same dimension triads are all size-change triads, with no saturation-change triads.

The procedure is the same as in Simulation Experiment 1. The proportions of relational choices are shown in Table 9. Consistent with the human pattern, the model was extremely accurate on the same-dimension triads in both conditions.

Experiment 2 conditions	Dimension	Relational choice	Non-relational choice	No choice
Experimental (PA) Condition	Same	100%	0%	0%
	Different	100%	0%	0%
Control Condition	Same	100%	0%	0%
	Different	50%	0%	50%

Table 9: Proportions of choice types for Kotovsky & Gentner (1996) Experiment 2 Simulation.

Also consistent with the human data, the model was far more accurate on the subsequent cross-dimensional triads in the experimental (progressive alignment) condition than in the control condition. In the progressive alignment condition, the model formed four interim generalizations for size-change vs saturation-change based on the type of change, symmetry or monotonicity. Examples of how the model forms and uses interim generalization are illustrated in Figure 23 and Figure 24 respectively.

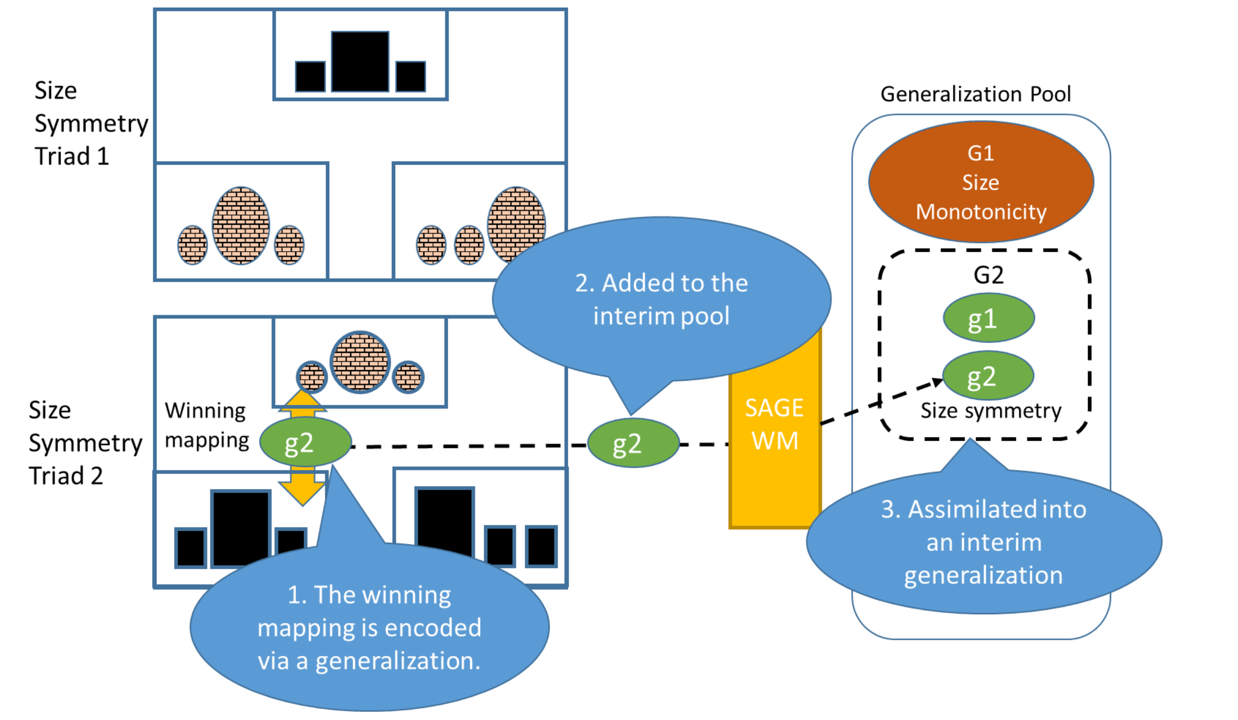

Figure 23: Creation of Interim Generalizations, during SDSP triads in Experiment 2 PA condition.

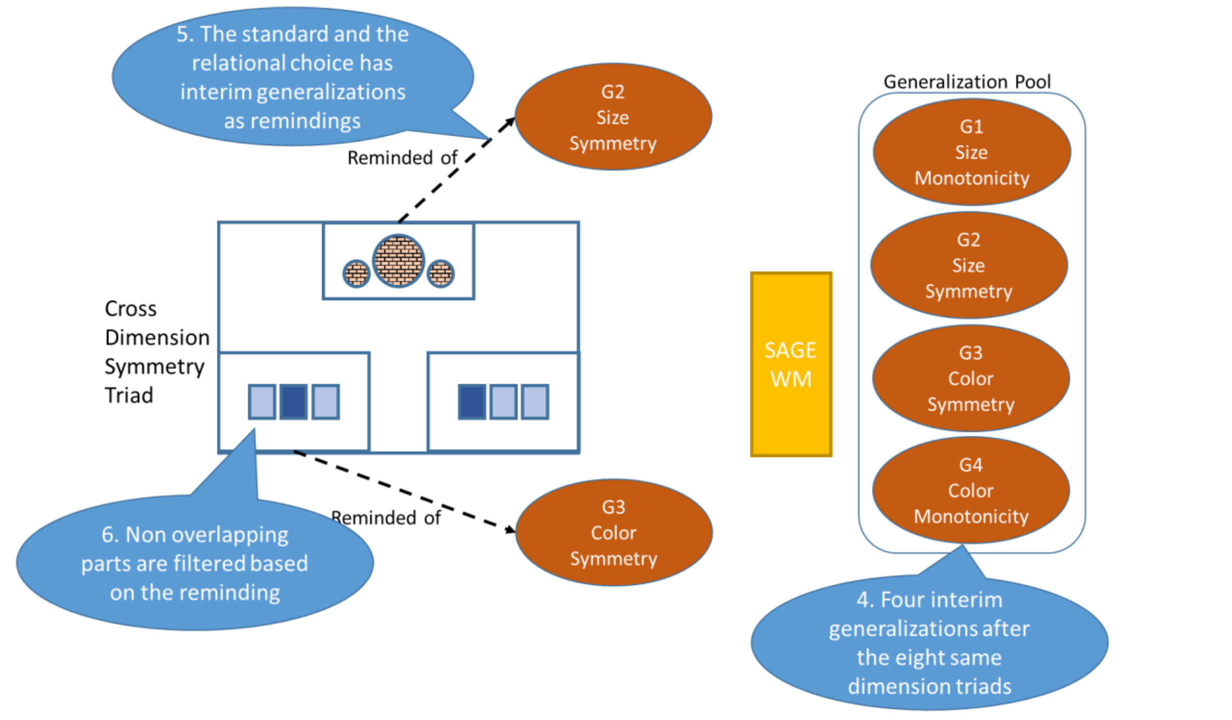

Figure 24: Remindings from interim generalization pool, for a cross dimension triad in Experiment 2 (PA) condition.

When the model retrieves an interim generalization as a reminding, only the overlap between the reminding and the portion of the stimulus is kept. The average-normalized-score of the relational choices increase as the non-contributing object attributes are filtered out. This drove re-representation, leading to relational choices being preferred. Figure 25 shows the representational change to the size-symmetry standard. First, the object attributes are filtered out because of the reminding and then, the dimension centric relations (e.g. `(biggerThan A B)` is re-represented into a dimension independent for, (e.g. `(greaterThan (Area A) (Area B))`.

By contrast, in the control condition, the model did not form any interim generalizations involving saturation-change. These results are qualitatively consistent with the results of Kotovsky & Gentner (1996). Like the children, the model performed better on cross-dimensional triads after progressive

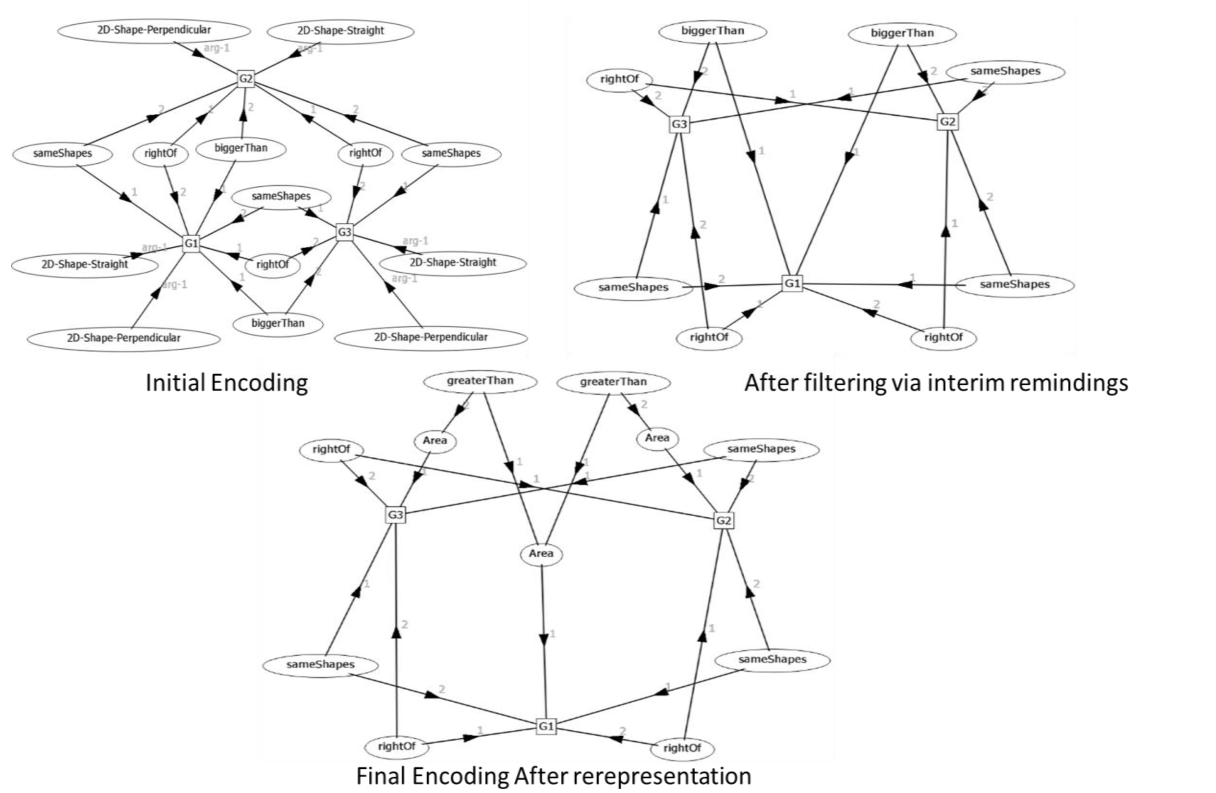

Figure 25: Graphical representation of the relational choice during different stages.

alignment on both dimensions than after progressive alignment only on the size/area dimension. Thus, the simulation has shown that the forced choice model can simulate the progressive alignment effects on 4 year olds found in (Kotovsky & Gentner 1996).

However, there are some discrepancies. First, the simulation performs too well, especially on the same-dimension triads. The model’s high degree of uniform encoding, and aggressive use of re-representation, appears to be going beyond what the children are doing. Perhaps, the differences may be partly due to the variability in encoding among children. Second, our model currently does not do several things that children probably do during the course of development. For example, it does not change its encoding strategy to shift to a more abstract comparative relation, nor does it introduce a new higher-order relationship (symmetry or monotonicity) to encode the newly-discovered pattern.

Since the model's behavior is qualitatively consistent with 4 year olds without these operations, it may be that the children are not doing this, but there is insufficient evidence to tell one way or the other.

Finally, we note that the simulation's responses are uniform and performance improves rapidly, whereas children exhibit a wider range of behavior. For example, even the 8 year olds in the original experiment were not at ceiling in this task. We predict that expanding the range of re-representation operations available, as well as looking for re-representation opportunities in both pairs, would widen the search space of the model and perhaps capture the more gradual improvement trajectory of children.

In summary, the results of the simulations support the claims of the thesis embodied in the forced choice model. We have demonstrated that the model can simulate comparison driven representational change, and the subsequent improvement in relational responding, observed in young children in the forced choice tasks.

Chapter 5: Related work

Representations are important for modeling cognition, and changes in cognitive capabilities are often tied to changes in representations (Dietrich & Markman, 2000). Gentner and her colleagues have done many studies regarding the dynamics of relational representations and its importance to analogy. This thesis builds on their work to provide a computational, algorithmic level explanation of representational change driven by language and progressive alignment, within the domain of forced choice tasks.

This chapter compares other relevant psychological and computational research. First, we summarize three psychology theories that addresses the issue of representational change, Karmiloff-smith's representational redescription (Karmiloff-Smith, 1995), Siegler's overlapping-waves theory (Siegler,

1998) and a dynamic systems approach for representational change (Stephen, Dixon, & Isenhower, 2009). Second, we discuss DORA, a computational model of learning relational representations (Doumas, Hummel, & Sandhofer, 2008), and a Bayesian model for learning relational representations (Kemp & Jern, 2009).

5.1 Cognitive psychology theories

5.1.1. Representational redescription hypothesis

The Representational Redescription hypothesis (RR) (Karmiloff-Smith, 1995) outlines an iterative process of how representations change during development. The RR framework posits that representational competency progresses through four different levels: Implicit (I), Explicit-1 (E1), Explicit-2 (E2) and

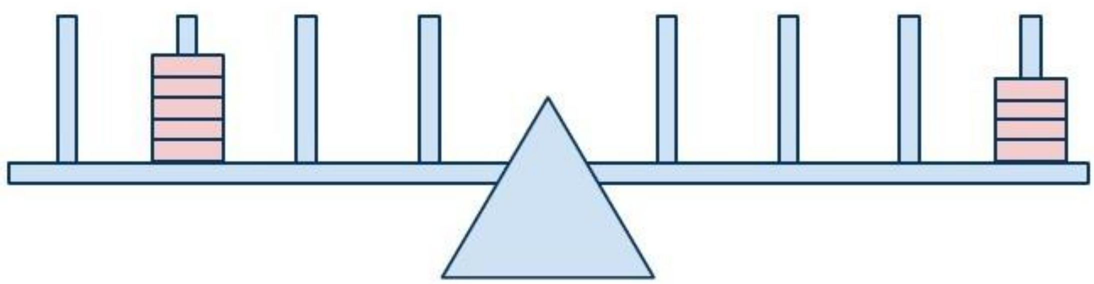

Figure 26: Balance Beam Task

Explicit-3 (E3). The flexibility of representations and their access to verbal explanation increases with each level. The implicit level (I) is the most inflexible and the explicit-3 level (E3) is the most flexible. RR predicts that the progression through the levels follow a U-shaped curve i.e. children's performance in the task drops when they progress through intermediate levels of representation due to overgeneralization.

For example, Pine & Messer (2003) tested the RR-model via children's progression through a balance beam task, illustrated in Figure 26, consisting of symmetrical and asymmetrical weights. As predicted by RR-model, initially children were eventually able to balance both types of beams but were unable to give any verbal explanation. This is interpreted as indicating that they were operating at the implicit level (I). Children then go through a phase where they appear to focus on the importance of weight for balancing, but overgeneralization causes them to perform poorly with asymmetrical beams. This level is soon followed by weight-based verbal explanations of why the beams did not balance. Eventually, they attain mastery with weight-and-distance based explanations of why the beams balance (E3).

There are a number of studies that provide support for the RR model (Butler, 2007). RR attempts to explain cognitive development in phenomenological terms rather than specifying precise mechanisms of change (Karmiloff-Smith, 1995). In contrast, our model gives a detailed description of a set of operations and representations sufficient for the computational implementation of representational change.

5.1.2. Overlapping Waves Theory

Siegler's overlapping waves theory (Opfer & Siegler, 2007) focuses on the development of problem solving abilities in young children. According to this theory, children know and use a variety of problem solving approaches involving different strategies, rules and representations. The theory posits that children formulate new rules by noticing and including potential explanatory variables whenever their observations conflict with expectations. Thus a child might start out using a rule based on only the variable of weight to solve a balance scale task. With experience and feedback, they learn that rules involving both length and weight are better predictors of the observations. In other words, cognitive development involves meta-cognitive processes that change which strategies and representations are used, and enables discovery of new ones by adaptation. But observing a conflict or other means of feedback is not always necessary for learning. For instance, children in Kotovsky & Gentner (1996) studies showed improvement in performance on cross-dimension tasks after experiencing progressively alignable same-dimension tasks. They did this without any feedback from the experimenter. Overlapping waves theory does not account for such changes. It does not provide an account of the mechanisms behind the discovery of potential explanatory variables nor the mechanisms that drive strategy change.

5.1.3. Dynamic systems approach

Stephen, Dixon, & Isenhower (2009) approach representational change from a dynamic systems theory perspective. They propose that new representations emerge from self-organization as explained by the theory of non-linear dynamics. Self-organization occurs via spontaneous breaking and reforming of constraints binding the parts of the system. The change is triggered by a critical instability i.e. when the entropy of the system reaches a critical point.

They tested their hypothesis in a gear system domain, where the participant is tasked with resolving the direction of rotation of a target-gear given the rotation of a driving gear. Their results are consistent with their hypothesis. However, they do not provide a computational model, nor is it clear that it could model any of the phenomena modeled in this work.

Unsurprisingly, none of these psychological models provide detailed representation or process models of representational change. This thesis, while not directly handling the same range as the models above, provides an algorithm-level account, which leads to new insights (e.g. postulating interim generalizations) and provides new capabilities for cognitive architectures and AI systems.

5.2 Computational models

5.2.1. Discovery of Relations by Analogy (DORA)

Doumas et al. (2008) present a theory of how structured relational representation can be learned from unstructured nonrelational examples. DORA learns new relations by detecting relational and featural invariants across experience. DORA is a symbolic-connectionist model. It is based on distributed semantic network activation and uses a set of algorithmic operations for discovery and predication of new relations (see Figure 27 from Doumas et al. (2008)).

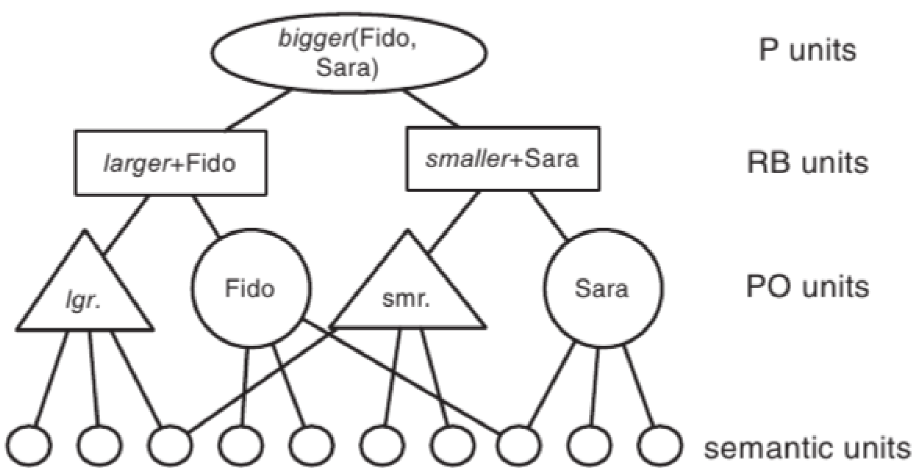

Figure 27: DORA's relational representations.

1. At the bottom level, there are distributed semantic units that code features of objects and relations.
2. Next layer has Predicate-Object (PO) units that code for individual predicates and objects (Example: Larger, Fido)
3. The third layer, called the Role Binding layer (RB), binds the attributes to objects to create single place predicates (example: larger (Fido))

4. On the top, we have the predicate layer (P), which binds predicates into multi-place relations. (example: bigger (Fido, Sara))

DORA represents relational structures as linked-set of role-filler pairs. DORA starts with a holistic representation of hand-coded object features that roughly corresponds to perceptual invariants. For example, the visual invariants include ‘round’, ‘square’, ‘shiny’ and relational invariants include ‘more’, ‘less’, ‘same’.

When propositions enter working memory, DORA sequentially activates a series of semantic units corresponding to the propositions. For example, assume that there are two objects in current focus: a red truck and a grey elephant. The semantic units corresponding to the object (elephant and truck) and their attributes such as grey, big, red, etc. will get activated. The role-filler binding is disambiguated by the temporal proximity of activation i.e. the binding of elephant to grey (grey elephant) instead to red is achieved by activating ‘grey’ right before or close enough to the activation of ‘elephant’.

DORA’s operation can be explained using four sets of propositions: the driver, the recipient, the emerging-recipient and the long-term memory (LTM). The semantic units are shared by all them. The driver is the set of propositions that are current focus of attention. Activation of semantic units that corresponds to the driver results in activations in LTM and the propositions are retrieved into the recipient. Roughly, driver represents the new situation, which reminds it of a similar situation encountered before (from LTM) and now have access to the remindings in working memory represented by the recipients (Hummel & Holyoak, 1997).

DORA, like its predecessor LISA, champions combining analogical retrieval with analogical mapping. After retrieval, the units in the recipients are mapped to the same type of units in the driver, creating mapping-hypotheses. The mapping-hypotheses are strengthened based on a simple Hebbian learning rule, as a product of activations of both the driver unit and the recipient unit.

DORA learns new single-place predicates using a simple algorithm for intersection discovery. For example, when DORA compares an elephant and a bear, it attempts to map them. Thanks to the mapping the units in driver can activate the units in the recipient which in turn passes excitation to the semantic units. Thus semantic units connected to both elephant and bear gets more excitation than semantic units unique to just one. DORA uses this higher excitation to hypothesize a potential single-place predicate and recruits a new PO unit i.e. the semantic units corresponding to ‘big’ will be more excited than others, hence a new PO for ‘big’ is predicated. Initially, the PO for ‘big’ may have unwanted featural overlaps, but gets refined progressively on encountering multiple examples.

Consider a scenario where the driver contains propositions about a Dog which is ‘big’ and ‘brown’, and a Cat which is ‘small’ and ‘furry’. Assume that DORA has already seen a similar configuration before of (say) a bear and a fox that are big and small respectively. DORA will be reminded of the bear-fox scenario when presented with the dog-cat scenario, and will map them to each other, such that POs match POs and RBs to RBs. This results in a mapping between ‘big’ and ‘small’ from both scenarios. DORA keeps track of activations and thus notices a systematic temporal pattern of activation between the units in both the driver and recipient. Consequently, DORA hypothesizes a new double-place relation such as ‘bigger-than’. In simple words, activation of ‘big’ in driver (dog-cat) will activate the ‘big’ in the recipient (bear-fox) and likewise for ‘small’ which eventually leads to the predication of ‘bigger-than’.

On encountering more examples, DORA refines the semantic features for the learned predicates using the same intersection discovery algorithm. Even though, learning higher-order relations, such as “cause”, was not part of the simulations, Doumas et al. (2008) claim DORA can learn higher-order relations the same way as it learned the ‘bigger-than’ relation. Additionally, they claim that the predicate refinement applied iteratively will eventually lead value-independent relation such as greater-than (size (a), size (b)) to emerge from value-dependent relation like bigger-than (a, b).

Despite an interesting attempt to model relation discovery and predication, DORA leaves many questions unanswered. DORA is based on a prior model of analogy called LISA, which has issues with scalability (Gentner & Forbus, 2011). It departs from LISAs approach by using unique semantic units and by addressing relation discovery. However, it suffers from limitations in terms of scalability similar to LISA. The original version of LISA is incapable of handling higher-order relational structures, unless the working memory limitations are alleviated via a new mechanism called group units (Hummel, Licato, & Bringsjord, 2014).

DORA recruits new PO unit whenever two existing predicates or objects are compared. There is a potential for explosion for PO units. DORA alludes that external constraints imposed by verbal labeling or adult instruction will keep the search space manageable. We agree that language provides benefits, but still it is far from likely that in the absence of language humans would exhibit a similar behavior. Additionally, the claim of using unstructured inputs must be viewed with caution due to the hand-coded nature of the inputs.

Furthermore, in our view, we question the idea that children start with the concept ‘big’ and ‘small’ before the discovery of ‘bigger-than’ relation. Studies suggest that infants are sensitive to identity relations (Tyrrell, Stauffer, & Snowman, 1991) and can acquire rules about physical events (Baillargeon, 2002). Nine-month-old infants are capable of generalizing and over hypothesizing (Dewar & Xu, 2010). Infant relational capabilities might be highly restricted to modality, dimension and context. Likewise, their representation could be more object-centric. Nevertheless, they must possess some rudimentary relational capabilities, as observed in other primates and even ducklings (Martinho & Kacelnik, 2016). Finally, the claim that simple intersection based predicate-refinement will eventually lead to the formation of dimension-independent relation from dimension-specific relation needs further explication and evidence.

5.2.2. Bayesian approach for learning relational categories

Kemp & Jern (2009) present a Bayesian model for learning and using relational categories. The model is based on the generative theory of similarity (Kemp, Bernstein, & Tenenbaum, 2005), which states that similarity judgements are inferences about generative processes. Every object is the outcome of a generative process, and two objects are similar if they are likely to have been generated by the same process.

The model starts with schemata represented using a logical language. The schemata correspond to abstractions created after encountering multiple instances of a concept. According to Kemp & Jern (2009), the schemata could be created using a Hierarchical Bayesian Approach described in (Gelman, Carlin, Stern, & Rubin, 2014). Each category is associated with a schema, that helps specify valid instances of the category. Figure 28 illustrates a schema *s*, the group *g* (randomly sampled valid instances) and the observation: o partially observed version of *g*. Given the observations, the model can select the most probable schema from the hypothesis space.

The hypotheses space is constructed using the templates in Figure 29. The templates are instantiated by substituting three dimensions: size, color and ball-position for D_i and three values along each dimensions {1, 2, 3}. Consequently, the hypothesis space consists of 1568 distinct schemas and roughly one million conjunctions.

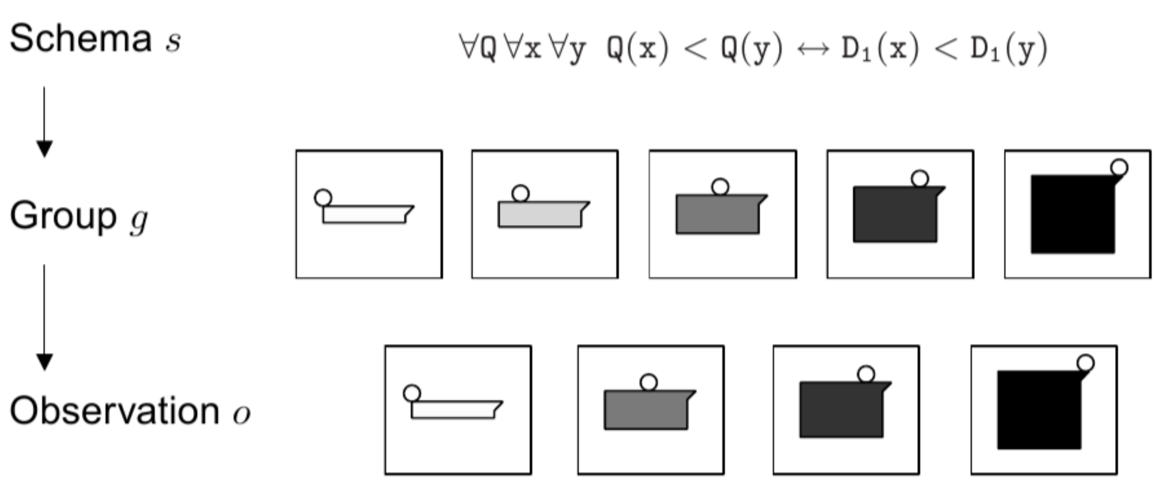

Figure 28: The Schema *s*, randomly sampled group *g* and partial observation *o* (Kemp & Jern 2009)

$$1\quad \left\{\begin{matrix}\forall x\\ \exists x\end{matrix}\right\} D_i(x)\,\{=,\neq,<,>\}\, v_k$$

$$2\quad \left\{\begin{matrix}\forall x\\ \exists x\end{matrix}\right\} \left\{\begin{matrix}\forall y\; x\neq y \rightarrow\\ \exists y\; x\neq y \wedge\end{matrix}\right\} D_i(x)\,\{=,\neq,<,>\}\, D_i(y)$$

$$3\quad \forall x\; D_i(x)\,\{=,\neq\}\, v_k \left\{\begin{matrix}\wedge\\ \vee\\ \leftrightarrow\end{matrix}\right\} D_j(x)\,\{=,\neq\}\, v_l$$

$$4\quad \forall x \forall y\; x\neq y \rightarrow \left(D_i(x)\,\{=,\neq,<,>\}\, D_i(y) \left\{\begin{matrix}\wedge\\ \vee\\ \leftrightarrow\end{matrix}\right\} D_j(x)\,\{=,\neq,<,>\}\, D_j(y)\right)$$

$$5\quad \left\{\begin{matrix}\forall Q\\ \exists Q\end{matrix}\right\} \left\{\begin{matrix}\forall x\\ \exists x\end{matrix}\right\} \left\{\begin{matrix}\forall y\; x\neq y \rightarrow\\ \exists y\; x\neq y \wedge\end{matrix}\right\} Q(x)\,\{=,\neq,<,>\}\, Q(y)$$

$$6\quad \left\{\begin{matrix}\forall Q\; Q\neq D_i \rightarrow\\ \exists Q\; Q\neq D_i \wedge\end{matrix}\right\} \forall x \forall y\; x\neq y \rightarrow \left(Q(x)\,\{=,\neq,<,>\}\, Q(y) \left\{\begin{matrix}\wedge\\ \vee\\ \leftrightarrow\end{matrix}\right\} D_i(x)\,\{=,\neq,<,>\}\, D_i(y)\right)$$

$$7\quad \left\{\begin{matrix}\forall Q\\ \exists Q\end{matrix}\right\} \left\{\begin{matrix}\forall R\; Q\neq R \rightarrow\\ \exists R\; Q\neq R \wedge\end{matrix}\right\} \forall x \forall y\; x\neq y \rightarrow \left(Q(x)\,\{=,\neq,<,>\}\, Q(y) \left\{\begin{matrix}\wedge\\ \vee\\ \leftrightarrow\end{matrix}\right\} R(x)\,\{=,\neq,<,>\}\, R(y)\right)$$

Figure 29: Templates used to construct the hypothesis space

They used a triad task for the experiments. They varied the dimensions of size, color and position-of-the-ball to create instances based on schemas. In the first experiment, the participants were shown a standard, three cards from the same category i.e. instances created using the same schema, and asked to decide which of the two choice groups (three cards each) belong to the same category. The model was run with the same stimuli given to humans.

The model works as follows. Given the standard group g_e and the two choice groups g_1 and g_2. The model computes the relative probability of the two hypotheses.

***h1**: g_e and g_1 are instances of the same schema, and g_2 is sampled randomly from other groups.*

***h2**: g_e and g_2 are instances of the same schema, and g_1 is sampled randomly from other groups.*

The priors for h1 and h2 are set uniformly i.e. P(h1) = P(h2) = 0.5. The model computes the conditional probability $P(h1| g_e, g_1, g_2)$ and $P(h2| g_e, g_1, g_2)$ by integrating over all schemata in the hypothesis space. The conditional probability of h1 will be higher if g_e and g_1 are instances of the same schema. Likewise, probability of h2 will be higher if g_e and g_2 are instances of the same schema. Based on which one is higher, the appropriate choice is chosen indicating the model's preference. They compared model's predictions to human response and found that for nine out of ten cases the model prefers the same choice as humans.

One serious limitation of the model is that there is no clear explanation of the origin of the hypothesis space. Also, their explanation of how the schemata could be learned via encountering instances, does not describe the mechanisms and the representations used for learning. Kemp & Jern (2009) note that two of their triads are similar to Kotovsky & Gentner (1996) triads and the model, as well as humans in their experiments, preferred the relational match. However, it is unclear how this model could account for the children's improvement in cross dimension triad tasks, after seeing progressively alignable same dimension triads. Furthermore, Bayesian approaches are typically considered to apply at Marr's computational level, rather than algorithmic level. Consequently, they provide less insight at the level of representations and processes.

Chapter 6: Conclusion & Future directions

This thesis presents a computational model of forced choice tasks. It proposes mechanisms of comparison driven representational change, evaluating them via simulating cognitive psychology studies. Chapter 1 presented the claims of the dissertation, and a summary of the studies simulated. Chapter 2 reviewed the background of the Companion cognitive architecture and the extensions we made to it. Chapter 2 also introduced SAGE-WM, our model of interim generalizations. Chapter 3 presented our model of forced choice tasks and a general analysis of our model's response to relational match forced choice tasks. Chapter 4 described the three simulations performed using the model, providing empirical evidence to support the claims of the dissertation. In chapter 5, we compared our approach to other psychology theories and cognitive models of representational change.

Here we start by revisiting our claims in light of the evidence provided by the simulations and close with a discussion of limitations and outline opportunities for future work.

6.1 Claims Revisited

Here we discuss each claim of the thesis.

<u>Claim 1:</u> Recent experiences affect how new problems or tasks are encoded. This can be modeled using interim generalizations.

The idea that recent experience affects how a new situation is encoded is not a new one. It is the algorithm by how this is done that is a novel contribution. We implemented SAGE-WM, a model of interim generalizations, to support this claim.

According to the claim, only positive (successful) experiences are assimilated in the interim generalization pool for recent experiences. This means that remindings from this pool tend to highlight representational structure that were useful for successful completion of the task, and irrelevant elements filtered out. For example, filtering out non-contributing attributes in cross-dimension triads of Kotovsky & Gentner (1996) simulation enabled the model to converge on the relational choice.

Claim 2: Forced choice tasks can be modeled using structure mapping comparisons.

1) The difference in structural evaluation scores between a standard and the alternatives is used for selecting a winning choice.
2) Verification of candidate inferences produced by comparisons is used to improve mappings.
3) When comparisons result in scores that are too close to discriminate, re-representation is triggered to attempt to differentiate between the alternatives.

The psychology experiments simulated here had no experimenter feedback. Nevertheless, classifying experiences as positive or negative is important for the model to categorize and utilize representations. The model determines success or failure for itself, based on whether or not the similarity between the standard and the choices are sufficiently different. That is, discriminability is measured by the differences in base normalized structural evaluation scores of the mappings. The success of the three simulations all provide evidence for this claim.

The second part of the claim is not a new idea and is probably the least controversial. Structure mapping comparison results in the projection of inferences based on the structural overlap. The inferences are only surmises and need verification before use.

The descriptions computed by people are likely to contain both perceptual and conceptual knowledge. However the initial coding is produced, it is highly unlikely that it includes everything that is known about an entity. Thus allowing the model to draw on additional background knowledge when verifying candidate inferences enables it to be influenced by knowledge not in the initial encoding. For instance, the Namy & Gentner (2002) simulations show how candidate inference validation tilts the comparison in favor of the taxonomic choice for the One-Kind condition.

Re-representation is important for adding flexibility to analogical matching. However, since an arbitrary number of re-representations are in principle possible for any description, this process must be tightly controlled. Our third claim argues that when the choice to select is not clear, re-representation is invoked on the standard and the choice which is nearest, as judged by the average-normalized structural evaluation score. For example, in Kotovsky & Gentner (1996) simulation, for the cross dimension stimuli in the progressive alignment condition of experiment 2, the remindings from the interim generalization pool helped in selecting the right candidate for re-representation. After re-representation, the relational choice wins, as rewriting dimension specific concrete relations into dimension-independent relations enabled the model to recognize the common relational structure. Thus, we showed how impasses can be resolved using re-representation.

<u>Claim 3:</u> Labeling two examples the same triggers a comparison for the purpose of understanding the meaning of the label. This can be simulated using structure mapping comparisons and generalizations.

1) The examples that are labeled the same are compared and assimilated into a generalization.

2) The generalization highlights commonalities and deemphasizes dissimilarities.
3) Examples are augmented with conceptual commonalities if possible, resulting in enhanced generalizations.

This claim addresses the importance of language/labels in representational change. There is a considerable literature showing that labeling sets of objects can facilitate children's forming categories around those objects. Subsequently, when objects share a label, there is an increased tendency to assimilate them into a generalization. The commonalities are highlighted when the objects are compared, and is captured via analogical generalization. We simulate this using SAGE and SAGE-WM. The increased tendency to assimilate is simulated by lowering the assimilation-threshold. The Christie & Gentner (2010) simulation shows how labeling and comparing objects results in generalizations that highlight relational commonalities while de-emphasizing (mostly perceptual) differences. This enabled the model to notice the relational commonalities between the standard(s) and the relational choice, and hence selects the relational choice as the winning choice.

The third part of the claim states that comparing familiar objects that are labeled the same results in an enriched generalization. The model enriches the generalization by augmenting the descriptions with background knowledge. Instead of bringing in all that is known about the objects compared, only the commonalities, if they exist, are considered. This is done by instantiating and comparing the schemas corresponding to the objects. The Namy & Gentner (2002) simulation shows how comparing familiar objects, which are from same category results in enriching the generalization with conceptual commonalities. As a result, the model chose the taxonomic choice more often than the perceptually similar choice, thanks to the enriched conceptual content in the generalization.

We revisit the forces of representational change with respect to the relational match forced choice task and the response pattern of the model. The model, similar to the children in the experiments, initially favors the perceptual or non-relational choice. However, over experience and learning, the model starts responding relationally, similar to the children in the experiments. This, as shown via the simulations, happens as the representations of the base and the targets change. Accordingly, another way to understand the forces of representational change is as forces acting along the dimensions of commonalities, as shown in Figure 30. This allows us to understand the forces that shifts/alters the representations somewhat independent of our representational assumptions.

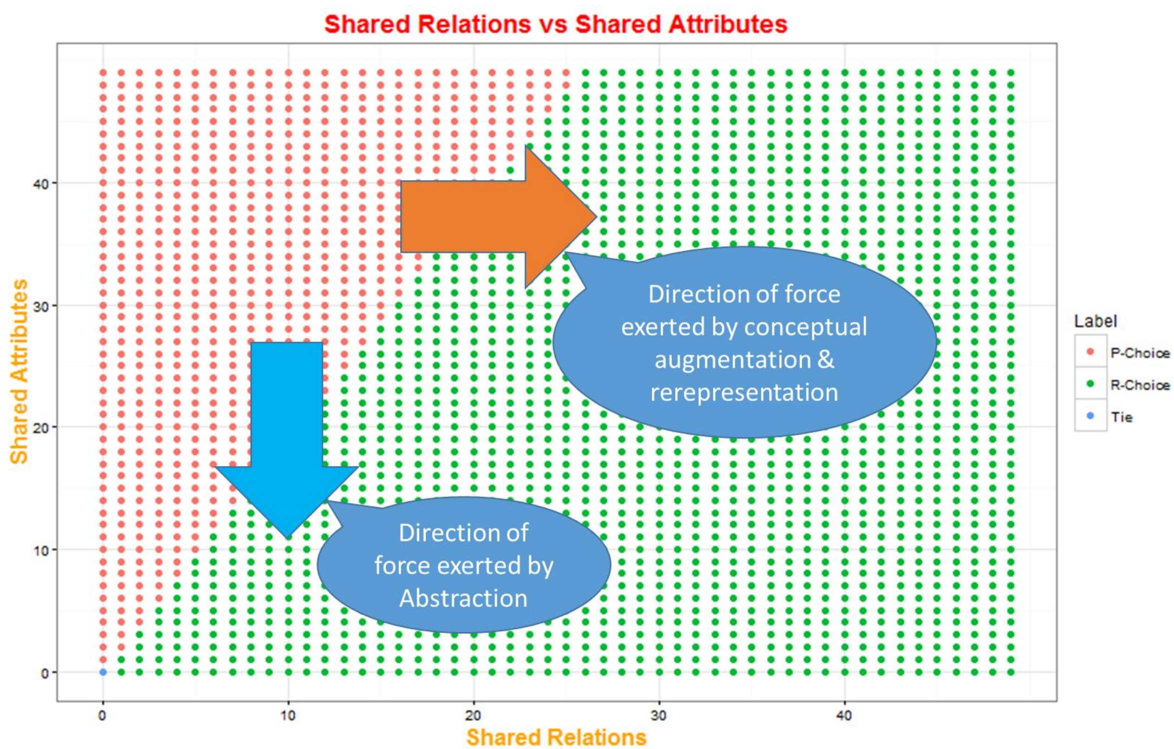

Figure 30: Forces of representational change

For example, in Christie & Gentner (2010) simulations, the standards share relational but not perceptual commonalities. The standards are labeled the same and hence generalized. The generalization is chosen as the base of the forced choice comparison, as opposed to a single standard in the Solo-Condition. Note that the base of Comparison-Condition (the generalization) has the same relational content as the base of Solo-Condition (single standard). But the generalization has lower perceptual commonalities with the perceptual choice. Thus we can view the effect of abstraction as exerting a downward force, by reducing shared attributes, changing the response of the model towards the relational (green) side.

Conceptual augmentation happens in two ways, enriching the generalization via schema comparison and via candidate inference validation. For example, in the simulation of Namy & Gentner (2002) One-Kind condition, the conceptual augmentation (both via enriching generalization and via candidate inference validation) increases the commonalities between the standard and the taxonomic choice, thereby exerting a rightward horizontal force.

Similarly, re-representation enables the model to notice relational commonalities that were not noticed earlier due to representational variance. For example, in Kotovsky & Gentner (1996) simulation, re-representation increases the relational commonality between the standard and the relational choice, thus, exerting a rightward force and changing the response of the model accordingly.

The simulation results presented here provide evidence that our model shows a plausible account of representational change within the relational match forced choice tasks.

6.2 Limitations and future work

Representational change is long-term. In Kotovsky & Gentner (1996) Experiments, adults and older children, but not 4 year olds, were able to appreciate relational similarity even if not supported by lower order similarity. Likewise, several studies show that adults exhibit exceptional relational competence compared to young children. This clearly indicates that representations evolve over years of experience and the representational knowledge is retained over time. In our model, we capture representational change that happens within-task and between-task within a single experimental session. We realize that there is much to be done to model long-term representational learning.

6.2.1. Simulating long term representational change

To capture human-like representational change over long term, capabilities to gather statistics about the use of representational elements and encoding strategies will be needed. Ultimately, this information leads to changes in encoding strategies. For instance, (say) a forced choice task that often requires comparing descriptions involving two different dimensions. The relations that are dimension-general (e.g. (greaterThan (Size A) (Size B))) should be preferred over dimension specific relation (e.g. (biggerThan A B)). This could be derived by maintaining statistics over situations which specifically used this transformation.

The process of applying transformations to elements of working memory is analogous to the rule matching operation of a production rule engine. Many production rule engines utilize rule-mapping algorithm like RETE (Forgy, 1982) which uses a tree-structured network for indexing production rules (Laird, 2012). We believe that analogical retrieval could also be a very effective tool for retrieving and applying transformation rules. The model can maintain transformation rules as generalizations in SAGE and retrieve it using MAC/FAC. In a production system, a rule is triggered when its left hand side matches elements in the working memory. This ensures that the rule is appropriate for the context. The results of SAGE retrievals would have to be tested to ensure that the antecedents of candidate inferences are valid in the current situation.

6.2.2. Learning new relations

An important theme that has emerged from the study of analogy, both empirical and theoretical, is that solutions to problems depend critically on what is represented and how it is represented. Models based on relational representations, including ours, contribute to understanding the nature of representations and how they change. However, little work, has been done to address the problem of how new relations are acquired in the first place.

We believe that the model will greatly benefit from a mechanism to learn new predicates. For example, the model could have introduced new predicates for "symmetry" and "monotonicity" in Kotovsky & Gentner (1996) stimuli based on the reoccurrence of relational patterns. Adding predicates to concisely summarize such patterns could help create higher-order structure to represent the stimuli, thereby making mapping and transfer more efficient.

6.3 Simulating other tasks

We modeled relational match forced choice tasks to provide evidence for our hypotheses about comparison driven representational change. There are many other tasks that has been successfully used in developmental studies to investigate representational change (Pine & Messer, 2003) (Opfer & Siegler, 2006). In the future, we plan to extend the model to support a wide range of tasks used in developmental psychology experiments. Given the current capabilities of our system, there are two tasks that are within immediate reach, the balance scale task and the card sorting task.

6.3.1. Balance Scale Task

The balance-scale task of naïve or intuitive physics has been of interest to developmental psychologists since its introduction by Piaget (e.g., Inhelder & Piaget, 1958). The task involves making a prediction about the state of a two-armed balance, which side will tip or whether it stays balanced, based on a configuration of weights (pegs) at particular distances from the fulcrum. The task is appealing because of

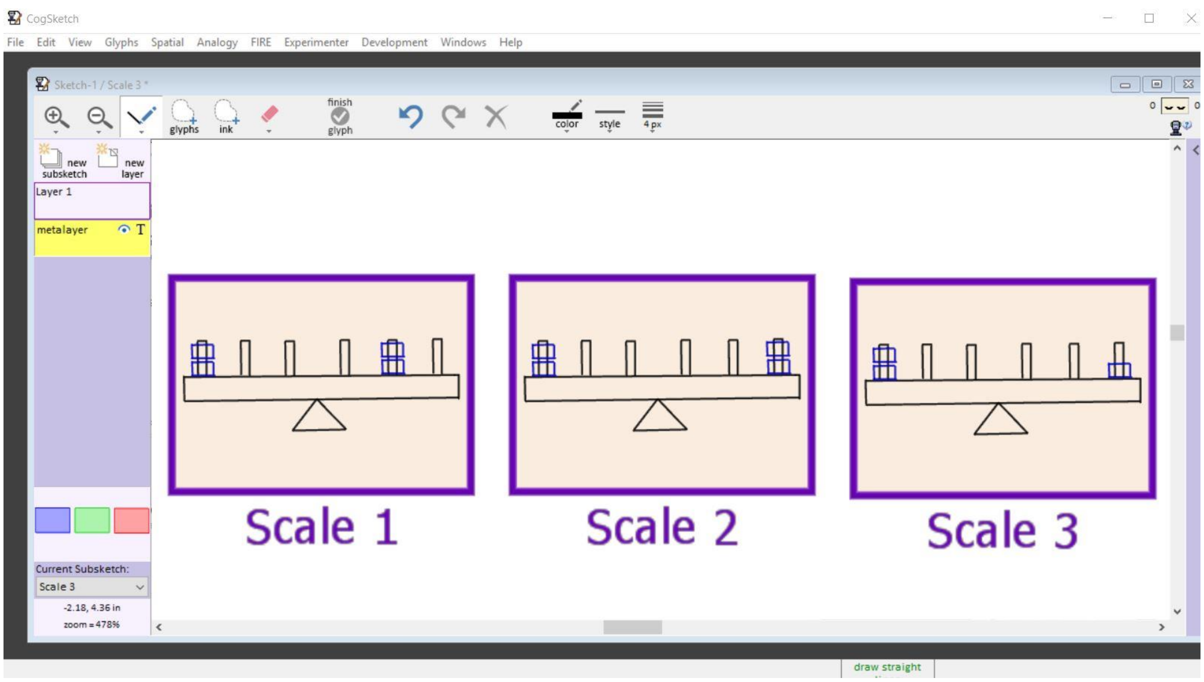

Figure 31: An example of sketched version of Balance Scale Task.

age-related trends in performance and has emerged as an important problem for researchers attempting to model cognitive/ representational dynamics (Siegler & Chen, 1998).

CogSketch, our sketch understanding system, has been used successfully as a part of a variety of computational experiments in a range of domains. CogSketch computes qualitative representations, but if necessary, relevant quantitative information can be computed (e.g. “distance of a peg from the fulcrum in balance-scale-task) and included into the descriptions (Chang & Forbus, 2012). Pine & Messer (2003) used a picture based balance scale task, where the participant was provided with a set of pictures and asked to choose the one that they think is balanced. This can be simulated with little modification to the current infrastructure. The balance scales in different configurations can be represented as CogSketch subsketches. Then the companion can be asked to pick "which one of the scale is balanced?". We can enable CogSketch to include distance & size information of the pegs in the descriptions for the subsketch. An example is illustrated in Figure 31.

Even though balance-scale problems can be viewed as a forced-choice task, the model presented here is not sufficient to cover them. However, we conjecture that extending our model to use Friedman's Assembled Coherence Theory of conceptual change (Friedman, 2012) would be able to handle such learning trajectories, since it includes mechanisms for uncovering causal variables based on observed behaviors.

6.3.2. Card Sorting Tasks

Card sorting is another task widely used by developmental psychologists to study a range of phenomenon such as executive functions (Zelazo, 2006), causal reasoning (Rottman, Gentner, &

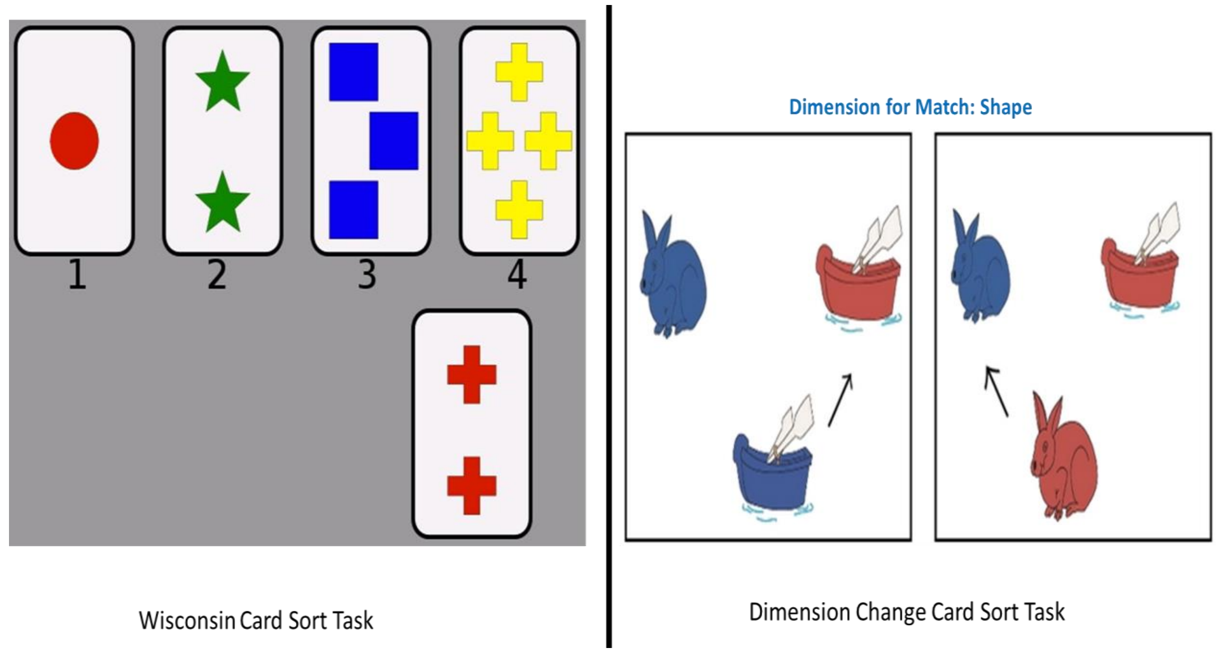

Figure 32: Card Sorting Tasks

Goldwater, 2012), etc. We are interested in two versions of the task, the dimensional change card sort task (DCCS) and the Wisconsin card sort task (WCST) (Figure 32).

In dimensional change card sort (DCCS), children are required to sort a series of bivalent test cards, first according to one dimension (e.g., color), and then according to the other (e.g., shape). DCCS measures cognitive and representational flexibility. The children are given clear instruction about what they are to do in every trial. In contrast, the Wisconsin card sort task (WCST) requires the children to discern the sort criterion by themselves based upon “correct” versus “incorrect” feedback given by the experimenter. After correctly matching a card according to a stimulus feature (color, form, or number) for N consecutive trials, the matching feature changes. Successful performance on the Card Sorting

Tasks requires determining the correct response in dimension and then maintain responding to that dimension. This requires switching the encoding strategy based on the dimension to focus. Card sorting tasks might be modeled by representing cards using subsketches, with similarity used to determine which subsketches should be placed together.

6.4 Alternative strategies for triggering and using re-representation

Our model uses re-representation to resolve failures of discriminability. As noted above, re-representation can be computationally expensive and hence must be under tight control. A different strategy from what was used in this work might be utilized to introduce higher-order structure to produce more systematic descriptions.

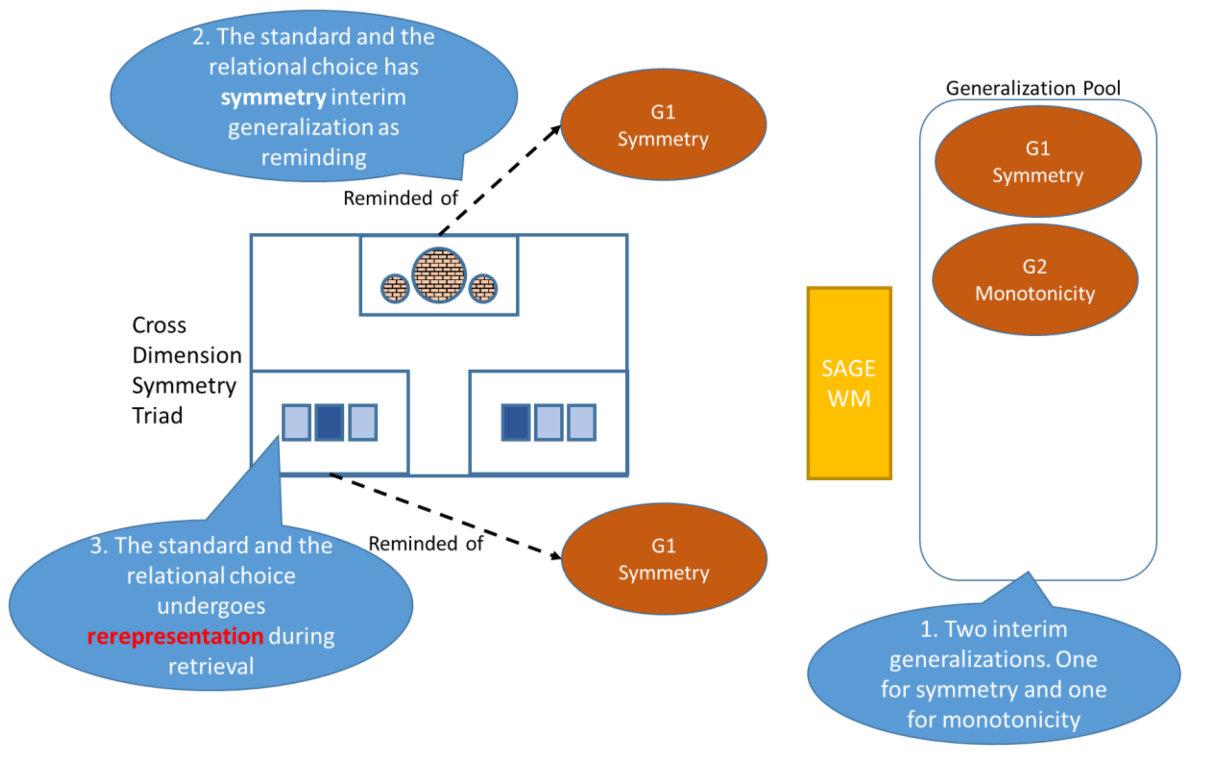

Figure 33: Symmetry generalization as reminding for cross dimension triad's standard and relational choice.

Studies have shown that analogical reasoning depends on executive resources of working memory. For example, dual-task experiments provide evidence that there is interference in performance from a concurrent WM task (Cho, Holyoak, & Cannon, 2007). The results of (Bor, Duncan, Wiseman, & Owen, 2003) suggest that effective reorganization of working memory content can decrease task difficulty. Participants who received a structured sequence of stimuli, designed to encourage reorganizing them into higher level chunks, performed significantly better in a spatial span task. Likewise, there is evidence that the gestalt principle of similarity, which leads to grouping similar elements, benefits visual working memory (Peterson & Berryhill, 2013). While working memory capacity is still murky, reducing working memory load seems like a psychologically plausible signal for re-representation.

Here is a potential way to extend our model to cover this. The basic idea is to allow re-representation during the assimilation of winning mappings and also during the application of remindings from interim generalization. We illustrate with an example from Kotovsky & Gentner (1996) simulation. After encountering the size symmetry triads, the model has a size symmetry interim generalization in the pool. When the model successfully completes the color symmetry triad, it adds the winning mapping to the interim pool. The pressure for re-representation could come from the necessity to keep working memory load small. That is, instead of adding a new element into the pool, the model attempts to assimilate it into an existing interim generalization. That way the model could reduce the number of elements in the pool and as a result, the size of the working memory. This strategy for triggering re-representation could result in the creation of a symmetry generalization and a monotonicity generalization after encountering all same dimension triads.

Later, when the model is given the first cross symmetry triad, it attempts to retrieve remindings from the pool for the standard and the choices. We propose that re-representation could be allowed during this phase. Instead of having to encode and construct a new description, the model could prefer using a description that already exists in the working memory. Hence, during reminding the model could allow re-representation as necessary, as shown in Figure 33. This would result in retrieving the symmetry generalization as reminding for the standard and the relational choice. As before, only the intersection is kept, but note that thanks to re-representation the encoding of relations (of both the standard and the relational choice) may already be in a dimension-independent form and thus, the relational choice would be selected as the winning choice.

Chapter 7: REFERENCES

Baillargeon, R. (2002). The acquisition of physical knowledge in infancy: A summary in eight lessons. *Blackwell Handbook of Childhood Cognitive Development*, *1*, 46–83.

Bor, D., Duncan, J., Wiseman, R. J. & Owen, A. M. (2003). Encoding strategies dissociate prefrontal activity from working memory demand. *Neuron*, *37*(2), 361–7.

Butler, C. (2007). *Evaluating the utility and validity of the representational redescription model as a general model for cognitive development*. Diss. University of Hertfordshire, 2007.

Chalmers, D. J., French, R. M. & Hofstadter, D. R. (1992). High-level perception, representation, and analogy: A critique of artificial intelligence methodology. *Journal of Experimental \& Theoretical Artificial Intelligence*, *4*(3), 185–211.

Chang, M. D. & Forbus, K. D. (2012). Using Quantitative Information to Improve Analogical Matching Between Sketches. In *Twenty-Fourth IAAI Conference*.

Cho, S., Holyoak, K. J. & Cannon, T. D. (2007). Analogical reasoning in working memory: Resources shared among relational integration, interference resolution, and maintenance. *Memory \& Cognition*, *35*(6), 1445–1455.

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

Day, S. B. & Gentner, D. (2007). Nonintentional analogical inference in text comprehension. *Memory \& Cognition*, *35*(1), 39–49.

Dewar, K. M. & Xu, F. (2010). Induction, overhypothesis, and the origin of abstract knowledge evidence from 9-month-old infants. *Psychological Science*.

Dietrich, E. & Markman, A. B. (2000). Cognitive dynamics: Computation and representation regained. *Cognitive Dynamics: Conceptual and Representational Change in Humans and Machines*, 5–30.

Doumas, L. A., Hummel, J. E. & Sandhofer, C. M. (2008). A theory of the discovery and predication of relational concepts. *Psychological Review*, *115*(1), 1.

Falkenhainer, B. (1988). Learning from physical analogies: a study in analogy and the explanation process (Doctoral Dissertation). Illinois University at Urbana (No. UIUCDCS-R-88-1479).

Falkenhainer, B., Forbus, K. D. & Gentner, D. (1986). *The structure-mapping engine*. Department of Computer Science, University of Illinois at Urbana-Champaign.

Falkenhainer, B., Forbus, K. D. & Gentner, D. (1989). The structure-mapping engine: Algorithm and examples. *Artificial Intelligence*, *41*(1), 1–63.

Forbus, K. D., & Gentner, D. (1989, August). Structural evaluation of analogies: What counts. In *Proceedings of the eleventh annual Conference of the Cognitive Science Society* (Vol. 34, pp. 341-348).

Forbus, K. D., Gentner, D. & Law, K. (1995). MAC/FAC: A model of similarity-based retrieval. *Cognitive Science*, *19*(2), 141–205.

Forbus, K. D., Gentner, D., Markman, A. B. & Ferguson, R. W. (1998). Analogy just looks like high level perception: Why a domain-general approach to analogical mapping is right. *Journal of Experimental & Theoretical Artificial Intelligence*, *10*(2), 231–257.

Forbus, K. D. & Hinrichs, T. R. (2006). Companion cognitive systems: a step toward Human-Level AI. *AI Magazine*, *27*(2), 83.

Forbus, K. D., Klenk, M. & Hinrichs, T. (2009). Companion cognitive systems: Design goals and lessons learned so far. *Intelligent Systems, IEEE*, *24*(4), 36–46.

Forgy, C. L. (1982). Rete: A fast algorithm for the many pattern/many object pattern match problem. *Artificial Intelligence*, *19*(1), 17–37.

Friedman, S. E. (2012). *Computational conceptual change: An explanation-based approach*.

Gelman, A., Carlin, J. B., Stern, H. S. & Rubin, D. B. (2014). *Bayesian data analysis* (Vol. 2). Chapman \& Hall/CRC Boca Raton, FL, USA.

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

Gentner, D. (1988). Metaphor as structure mapping: The relational shift. *Child Development*, 47–59.

Gentner, D. (2003). Why We're So Smart. *Language in Mind*, 195.

Gentner, D. & Forbus, K. D. (2011). Computational models of analogy. *Wiley Interdisciplinary Reviews: Cognitive Science*, *2*(3), 266–276.

Gentner, D. & Namy, L. L. (2006). Analogical processes in language learning. *Current Directions in Psychological Science*, *15*(6), 297–301.

Gentner, D. & Rattermann, M. (1991). 7. Language and the career of similarity. *Perspectives on Language and Thought: Interrelations in Development*, 225.

Hofstadter, D. R. (2008). *Fluid concepts and creative analogies: Computer models of the fundamental mechanisms of thought*. Basic books.

Holyoak, K. J., Gentner, D. & Kokinov, B. N. (2001). Introduction: The place of analogy in cognition. *The Analogical Mind: Perspectives from Cognitive Science*, 1–19.

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

Hummel, J. E., Licato, J. & Bringsjord, S. (2014). Analogy, explanation, and proof. *Frontiers in Human Neuroscience*, *8*.

Inhelder, B. & Piaget, J. (1958). The growth of logical thinking.

Kamp, H. & Reyle, U. (1993). *From discourse to logic: Introduction to modeltheoretic semantics of natural language, formal logic and discourse representation theory*. Springer.

Kandaswamy, S., Forbus, K. D. & Gentner, D. (2014). Modeling Learning via Progressive Alignment using Interim Generalizations.

Karmiloff-Smith, A. (1995). *Beyond modularity: A developmental perspective on cognitive science*. MIT press.

Kemp, C., Bernstein, A. & Tenenbaum, J. B. (2005). A generative theory of similarity. In *Proceedings of the 27th Annual Conference of the Cognitive Science Society* (pp. 1132–1137).

Kemp, C. & Jern, A. (2009). Abstraction and relational learning. In *Advances in neural information processing systems* (pp. 934–942).

Kotovsky, L. & Gentner, D. (1996). Comparison and categorization in the development of relational similarity. *Child Development*, *67*(6), 2797–2822.

Laird, J. (2012). *The Soar cognitive architecture*. MIT Press.

Lockwood, K., Lovett, A. & Forbus, K. (2008). Automatic classification of containment and support spatial relations in English and Dutch. In *Spatial Cognition VI. Learning, Reasoning, and Talking about Space* (pp. 283–294). Springer.

Markman, A. B. & Dietrich, E. (2000). In defense of representation. *Cognitive Psychology*, *40*(2), 138–171.

Martinho, A. & Kacelnik, A. (2016). Ducklings imprint on the relational concept of "same or different." *Science*, *353*(6296), 286–288.

McLure, M., Friedman, S., & Forbus, K. (2010). Learning concepts from sketches via analogical generalization and near-misses. In *Proceedings of the 32nd Annual Conference of the Cognitive Science Society (CogSci)*.

McLure, M. D., Friedman, S. E., Lovett, A., & Forbus, K. D. (2011). Edge-cycles: A qualitative sketch representation to support recognition. In *Proceedings of the 25th International Workshop on Qualitative Reasoning*.

Namy, L. L. & Gentner, D. (2002). Making a silk purse out of two sow's ears: young children's use of comparison in category learning. *Journal of Experimental Psychology. General*, *131*(1), 5–15.

Opfer, J. E. & Siegler, R. S. (2007). Representational change and children's numerical estimation. *Cognitive Psychology*, *55*(3), 169–195. doi:10.1016/j.cogpsych.2006.09.002

Peterson, D. J. & Berryhill, M. E. (2013). The Gestalt principle of similarity benefits visual working memory. *Psychonomic Bulletin & Review*, *20*(6), 1282–9. doi:10.3758/s13423-013-0460-x

Peura, M. & Iivarinen, J. (1997). Efficiency of simple shape descriptors. *Aspects of Visual Form*, 443–451.

Pine, K. & Messer, D. (2003). The development of representations as children learn about balancing. *British Journal of Developmental Psychology*, *21*(2), 285–301.

Rattermann, M. J. & Gentner, D. (1998). More evidence for a relational shift in the development of analogy: Children's performance on a causal-mapping task. *Cognitive Development*, *13*(4), 453–478.

Rottman, B. M., Gentner, D. & Goldwater, M. B. (2012). Causal systems categories: Differences in novice and expert categorization of causal phenomena. *Cognitive Science*, *36*(5), 919–932.

Siegler, R. S. (1998). *Emerging minds: The process of change in children's thinking*. Oxford University Press.

Stephen, D. G., Dixon, J. A. & Isenhower, R. W. (2009). Dynamics of representational change: entropy, action, and cognition. *Journal of Experimental Psychology: Human Perception and Performance*, *35*(6), 1811.

Taylor, J. L., Friedman, S. E., Forbus, K., Goldwater, M. & Gentner, D. (2011). Modeling structural priming in sentence production via analogical processes. In *Proceedings of the 33rd Annual Conference of the Cognitive Science Society* (pp. 2916–2921).

Tomai, E. R. (2009). *A pragmatic approach to computational narrative understanding*. Evanston, IL.

Tyrrell, D. J., Stauffer, L. B. & Snowman, L. G. (1991). Perception of abstract identity/difference relationships by infants. *Infant Behavior and Development*, *14*(1), 125–129.

Yan, J., Forbus, K. D. & Gentner, D. (2003). A theory of rerepresentation in analogical matching. In *Proceedings of the Twenty-fifth Annual Meeting of the Cognitive Science Society*.

Zelazo, P. D. (2006). The Dimensional Change Card Sort (DCCS): a method of assessing executive function in children. *Nature Protocols, 1*(1), 297–301. doi:10.1038/nprot.2006.46

APPENDIX A: BACKGROUND SCHEMA EXAMPLES

Object	WordSmyth Definition
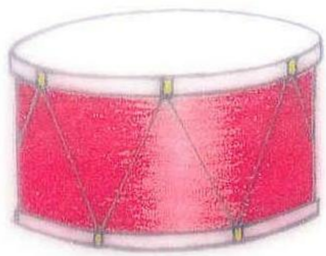 Drum	An instrument for playing music that has a hollow round shape and a tight covering over an open end. You play a drum by hitting its surface with your hands or sticks.

Schema

```
(in-gcontext (ConceptSchemaFn DrumInstrument))
(isa (GenEntFn 0 0 (ConceptSchemaFn DrumInstrument)) DrumInstrument)
(isa (GenEntFn 0 0 (ConceptSchemaFn DrumInstrument))
MusicalInstrument)
(isa (GenEntFn 1 0 (ConceptSchemaFn DrumInstrument)) Drumstick)
(isa (GenEntFn 2 0 (ConceptSchemaFn DrumInstrument))
(InstrumentPlayingFn DrumInstrument))
(deviceUsed (GenEntFn 2 0 (ConceptSchemaFn DrumInstrument))
            (GenEntFn 0 0 (ConceptSchemaFn DrumInstrument)))
(deviceUsed (GenEntFn 2 0 (ConceptSchemaFn DrumInstrument))
            (GenEntFn 1 0 (ConceptSchemaFn DrumInstrument)))
```

Object	WordSmyth Definition
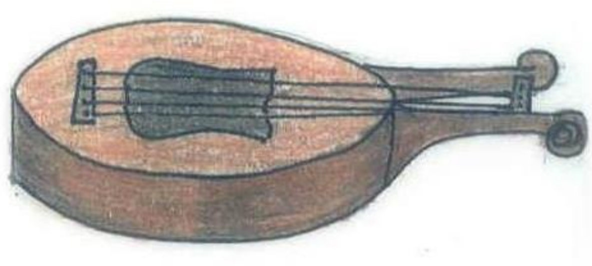 Rounded Lyre	a stringed instrument of ancient Greece that is like a harp.

Schema

```
(in-gcontext (ConceptSchemaFn RoundedLyre-TheStringInstrument))
(isa (GenEntFn 0 0 (ConceptSchemaFn RoundedLyre-TheStringInstrument))
RoundedLyre-TheStringInstrument)
```

```
(isa (GenEntFn 0 0 (ConceptSchemaFn RoundedLyre-TheStringInstrument))
MusicalInstrument)
(isa (GenEntFn 1 0 (ConceptSchemaFn RoundedLyre-TheStringInstrument))
(InstrumentPlayingFn RoundedLyre-TheStringInstrument))
(isa (GenEntFn 2 0 (ConceptSchemaFn RoundedLyre-TheStringInstrument))
Finger)
(deviceUsed (GenEntFn 1 0 (ConceptSchemaFn RoundedLyre-
TheStringInstrument))
            (GenEntFn 0 0 (ConceptSchemaFn RoundedLyre-
TheStringInstrument)))
(bodyPartsUsed (GenEntFn 1 0 (ConceptSchemaFn RoundedLyre-
TheStringInstrument))
               (GenEntFn 2 0 (ConceptSchemaFn RoundedLyre-
TheStringInstrument)))
```

Object	**WordSmyth Definition**
 Cake	a sweet food made of batter and cooked in an oven. {birthday cake}

Schema

```
(in-gcontext (ConceptSchemaFn BirthdayCake))

(isa (GenEntFn 0 0 (ConceptSchemaFn BirthdayCake)) BirthdayCake)
(isa (GenEntFn 0 0 (ConceptSchemaFn BirthdayCake)) Baked)
(tasteOfObject (GenEntFn 0 0 (ConceptSchemaFn BirthdayCake))
SweetTaste)
(hasPreparationStyle (GenEntFn 0 0 (ConceptSchemaFn BirthdayCake))
Baked)
```

Object	WordSmyth Definition
 Bucket	an open container with round sides, a flat bottom, and a curved handle at the top.

Schema

```
(in-gcontext (ConceptSchemaFn Bucket))

(isa (GenEntFn 0 0 (ConceptSchemaFn Bucket)) Bucket)
(isa (GenEntFn 0 0 (ConceptSchemaFn Bucket)) ContainerArtifact)
(isa (GenEntFn 1 0 (ConceptSchemaFn Bucket)) Handle)
(physicalParts (GenEntFn 0 0 (ConceptSchemaFn Bucket))
               (GenEntFn 1 0 (ConceptSchemaFn Bucket)))
```

Object	WordSmyth Definition
 Flute	An instrument for playing music. It is a long tube made of metal or wood that you play by blowing into a hole at one end.

Schema

```
(in-gcontext (ConceptSchemaFn Flute))

(isa (GenEntFn 0 0 (ConceptSchemaFn Flute)) Flute)
```

```
(isa (GenEntFn 0 0 (ConceptSchemaFn Flute)) MusicalInstrument)
(isa (GenEntFn 1 0 (ConceptSchemaFn Flute)) (InstrumentPlayingFn
Flute))
(deviceUsed (GenEntFn 1 0 (ConceptSchemaFn Flute))
            (GenEntFn 0 0 (ConceptSchemaFn Flute)))
```