Analogical Dialogue Acts: Supporting Learning by Reading Analogies

David Barbella
Qualitative Reasoning Group
Northwestern University
2133 Sheridan Road, Evanston, IL, USA
barbella@u.northwestern.edu

Kenneth D. Forbus
Qualitative Reasoning Group
Northwestern University
2133 Sheridan Road, Evanston, IL, 60201, USA
forbus@northwestern.edu

Abstract

Analogy is heavily used in written explanations, particularly in instructional texts. We introduce the concept of *analogical dialogue acts* (ADAs) which represent the roles utterances play in instructional analogies. We describe a catalog of such acts, based on ideas from structure-mapping theory. We focus on the operations that these acts lead to while understanding instructional texts, using the Structure-Mapping Engine (SME) and dynamic case construction in a computational model. We test this model on a small corpus of instructional analogies, expressed in simplified English, which were understood via a semi-automatic natural language system using analogical dialogue acts. The model enabled a system to answer questions after understanding the analogies that it was not able to answer without them.

1 Introduction

People use analogy heavily in written explanations. Instructional texts, for example, use analogy to convey new concepts and systems of related ideas to learners. Any learning by reading system must ultimately include the capability of understanding such analogies. Here we combine Gentner's (1983) structure-mapping theory with ideas from dialogue act theory (Traum, 2000) to describe a catalog of analogical dialogue acts (ADAs) which capture the functional roles that discourse elements play in instructional analogies. We outline criteria for identifying ADAs in text and describe what operations they suggest for discourse processing. We provide evidence that this model captures important aspects of understanding instructional analogies via a simulation that uses knowledge gleaned from reading instructional analogies to answer questions.

We start by reviewing the relevant aspects of structure-mapping theory and dialogue act theory. Then we describe our catalog of analogical dialogue acts, based on a theoretical analysis of the roles structure-mapping operations can play in language understanding. A prototype implementation of these ideas is described next, followed by an experiment illustrating that these ideas can be used to understand analogies in text, based on answering questions. We close with a discussion of related and future work.

2 Background

Dialogue act theories (also called speech acts (Allen & Perrault, 1980)) are concerned with the roles utterances play in discourse and the effects they have on the world or on understanding. An utterance identified as a Requesting Information, for example, might take the syntactic form of a question that makes the information requested explicit, e.g. "What time is it?" The surface manifestation might instead be a statement, or an indirect question, e.g. "Do you have the time?" In other words, its classification is based on its function in the dialogue and the set of operations it suggests for the recipient to undertake. We claim that there exists a set of analogical dialogue acts that are used in communicating analogies. Like other dialogue acts, they have criteria by which they can be rec-

Proceedings of the NAACL HLT 2010 First International Workshop on Formalisms and Methodology for Learning by Reading, pages 96–104,
Los Angeles, California, June 2010.

ognized, and a set of implied commitments and obligations for the dialogue participants. This paper focuses on instructional analogies in texts, both because they are an important phenomenon and because it allows us to factor out follow-up questions, making it a useful starting point.

There are a wide variety of dialogue act models, but all of them include some variation of acts like Inform (Traum, 2000), which indicate the intent to describe the state of the world. The analogical dialogue acts we discuss here can be viewed as specializations of Inform.

The organization of analogical dialogue acts follows directly from the concepts of structure-mapping theory. In structure-mapping, analogical matching takes as input two structured, relational representations, the *base* and *target*, and produces as output one or more *mappings*. Each mapping consists of a set of *correspondences*, identifying how entities and statements in the base align with entities and statements in the target. Mappings include a *structural evaluation score* providing an estimate of their overall quality. This estimate is based on s*ystematicity*, i.e., the amount of nested relational structure in the mapping, especially higher-order relations that serve as inferential connections between other statements. Causal, logical, and mathematical statements are all examples of higher-order relations. Systematicity thus serves as a local heuristic measure of the explanatory promise of a mapping.

Mappings can also contain *candidate inferences*, statements in the base that are projected onto the target, using the correspondences of the mapping. The candidate inferences represent conjectures about the target, and constitute a source of analogy's generative power. Whether or not the candidate inferences are in fact correct is evaluated outside of the matching process. In discourse, candidate inferences are often used to convey new information about the target to the learner. Candidate inferences can be forward, from base to target, or reverse, from target to base. Candidate inferences also represent differences between two representations, when they cannot be consistently projected from one description to the other.

The Structure-Mapping Engine (SME, Falkenhainer et al 1989) provides a simulation of analogical matching. SME typically produces only one mapping, but can produce a second or third mapping if they are sufficiently close to the best mapping. SME can accept input about the base and target incrementally, updating its mappings as new information becomes available (Forbus et al 1994), which can be important for modeling the incremental nature of discourse. One cost of SME's greedy match algorithm and incremental operation is that matches can go awry. Consequently, SME also supports a small set of constraints, optionally specified as part of the matcher's input, which guide it based on task constraints. Here the relevant constraints are those concerning correspondences. That is, given a base item b_i and target item t_j, either entities or statements, the following constraints are defined: *required*$(b_i\ t_j)$ means that b_i must correspond to t_j in every mapping, and *excluded*$(b_i\ t_j)$ means that b_i cannot correspond to t_j in any mapping. The following open constraints are also defined: *requiredBase*(b_i), means that something in every mapping must correspond to b_i, with *requiredTarget*(t_j) defined similarly. *excludedBase*(b_i) means that b_i cannot participate in any correspondence, with *excludedTarget*(t_j) defined similarly.

An important problem in understanding analogy in discourse concerns how the representations provided to SME are constructed. As described below, the representations that constitute an understanding of the text are produced in our model via a semi-automatic natural language understanding system, which reduces tailorability. In understanding instructional analogies, a learner is expected to draw upon their existing world knowledge. In some situations, whole cases representing a prior experience are retrieved from memory. In other situations, cases seem to be constructed dynamically from one's general knowledge of the world. We use *dynamic case construction* methods (Mostek et al 2000) to model this process. In dynamic case construction, a seed entity or concept is provided as a starting point, and facts which mention it are gathered, perhaps filtering by some criterion. For example, "The economy of India" might have India as its seed, and facts filtered based on their judged relevance to economic matters. When a reader is processing an instructional analogy, we believe that something like this process is used to create representations to be used in their understanding of the analogy.

Heat flows from one place to another because the temperature of the two places is different. A hot brick loses heat to a cool room. The temperature difference - the brick's temperature minus the room's temperature – drives the heat from the brick. Heat leaks from the brick until the temperature difference is gone. No more heat flows from the brick when it becomes as cool as the room it is in.
Similarly, a full can of water will leak volume from a hole in the side of the can. The depth of the water is higher than the depth of the hole, so the depth difference drives volume out through the hole.
Eventually, all the volume that can leak out does so. When this happens, the water depth has fallen so that it is the same as that of the hole. There is no more depth difference, so no more volume flows out through the hole. Just as a difference in temperature causes heat to flow, so a difference in depth causes volume to flow. When there is no temperature difference, heat flow ceases; when there is no depth difference, volume flow ceases.

Extend Target
Extend Base
Introduce Comparison
Candidate Inference

Figure 1: An analogy from our test corpus, hand-annotated with analogical dialogue acts.

3 Analogical Dialogue Acts

Our model of analogical dialog acts is based on an analysis of how the functional constraints on performing analogical mapping and case construction interact with the properties of discourse. To carry out an analogy, a reader must be able to infer that an analogy is required. They must understand what goes into the base and what goes into the target, which can be complex because what is stated in the text typically needs to be combined with the reader's own knowledge. Since readers often know quite a lot to begin with, figuring out which subset of what they know is relevant to the analogy can be complicated. Finally, they have to understand how the author intends the mapping to go, since there can be multiple mappings between the same domains. Analogical dialogue acts, we argue, provide readers with information that they need to perform these tasks.

Let us examine this process in more detail. To carry out an analogy, the contents of the base and target representations must be identified. A fundamental problem is that the reader must figure out an appropriate construal of the base and target, i.e., what subset of their knowledge should be brought to bear in the current comparison? A reader's starting knowledge may or may not be sufficient to guide the mapping process correctly, in order to reconstruct the mapping that the author intended. This is especially true in instructional analogies, of course. We believe that this is why one commonly finds explicit information about intended correspondences provided as part of instructional analogies. Such information provides a source of constraints that can be used to guide case construction and mapping. Similarly, and we believe for similar reasons, the desired inferences to be drawn from the analogy are often highlighted. Since there can be multiple construals (i.e., specific sets of facts retrieved) for the given base and target, mentioning candidate inferences explicitly provides clues to the reader about how to construe the base and target (i.e., the given candidate inference should be derivable) as well as information about its validity.

Next we describe our proposed analogy dialogue acts. For each act, we give an example, some criteria for identifying them, and describe what operations a reader might do when they detect such an act has occurred. At this point our focus has been on developing the basic set and the operations they entail, rather than on developing a comprehensive set of identification criteria. The first three acts are concerned with introducing the representations to be compared, and the rest are concerned with correspondences and candidate inferences. We use a greenhouse/atmosphere analogy as a source of examples.

Introduce Comparison: Introduces a comparison by providing both base and target. For example, in "We can understand the greenhouse effect by comparing it to what goes on in an actual greenhouse." the base is a greenhouse, and the target is the Earth's atmosphere. Recognizing an Introduce Comparison can require combining information across multiple sentences. In Figure 1, for example, the target is described in the paragraph above the point where the comparison is introduced. Sometimes this intent must be inferred from parallel sentence structure in subsequent sentences and other sophisticated rhetorical devices, while in other cases, like this example, the comparison is introduced explicitly.

What is the base and what is the target requires a non-local assessment about what the containing text is about. (This particular example is drawn from a book on solar energy, and the rest of the chapter makes clear that heat is the domain being taught.) Since we assume that candidate inferences can be constructed bidirectionally, an incorrect assessment is not catastrophic.

Processing an Introduce Comparison act requires finding appropriate construals of the base and target. The target, as in this case, is constrained by what has already been introduced in the text. The base, unless it has been used before in the same text and is being used in a consistent manner, must be constructed from the reader's knowledge. Whether this is done aggressively or lazily is, we suspect, a strategy that is subject to individual variation. Ambiguity in linguistic cues can lead to the need to explore multiple construals, to find combinations with significant overlap.

Extend Base, **Extend Target:** These acts add information to the base or target of a comparison, respectively. Such acts are identified by relationships and/or entities being mentioned in the same statement as an entity in the base or target, but which is not a statement about correspondences or candidate inferences. For example, "The glass of a greenhouse lets the short solar rays through." is extending the base, and "The earth's atmosphere admits most of the solar radiation." is an example of extending the target. Entities that are mentioned in these acts are added to the construal of the case, if not there already, by retrieving additional knowledge about them, focusing on statements involving other entities in the current construal. If the specific facts mentioned are not already known to the reader, they are provisionally accepted as being true about the base or target, as appropriate.

Introduce Correspondence: These acts provide clues as to the author's intended mapping. For example, "The Earth's atmosphere is like the glass in the greenhouse." indicates that "Earth's atmosphere" corresponds to "glass in greenhouse". Distinguishing these acts from introducing a comparison can be tricky, since "is like" is a syntactic pattern common to both. The first occurrence of "is like" in such cases is typically the introduction of the base and target, with subsequent statements introducing correspondences. Sometimes Introduce Correspondence acts are expressed as identity statements, e.g. "The glass is the atmosphere." Sometimes these acts are signaled by pairs of sentences, one expressing a fact about the base followed immediately by one about the target, with identical syntax.

When an Introduce Correspondence act is detected, the base and target are checked to see if they already contain the entities or relationships mentioned. If they do not, then the descriptions are extended to include them. The final step is introducing a *required* constraint between them as part of the input to SME. If mappings have already been generated that are not consistent with this constraint, they are discarded and new mappings are generated.

Block Correspondence: These acts are provided by the author to block a correspondence that a reader might otherwise find tempting. An example is "The greenhouse door is not like the hole in the ozone layer." We believe that these acts are relatively rare, and especially in written text compared with spoken dialogue, where there are opportunities for feedback, a matter discussed later.

When both a base and target item are mentioned, an *exclude* constraint is introduced between them. When only one of them is mentioned, the minimal operation is to add an open exclusion constraint (e.g. *excludedBase* or *excludedTarget*). The reader may decide to simply remove the excluded item from the construal, along with all of the facts that mention it. This would prevent it from being mapped, but it would also prevent it from appearing in any candidate inferences, and hence is more extreme.

Introduce Candidate Inference: These acts alert the reader to information that the author intended to convey via the analogy. An example is "Just as heat is trapped by the greenhouse roof, heat is trapped by the Earth's atmosphere." Phrases such as "just as" and "just like", or even "Like <*base statement to be projected*>, <*resulting candidate inference*>." are clues for identifying such acts. If the candidate inference can be found in the mapping that the reader has built up so far, then that surmise should be given additional weight as being true. (If it is already known by the reader, it may already be part of a mapping. This does not indicate failure, only that it is uninformative for that reader.) If the candidate inference cannot be

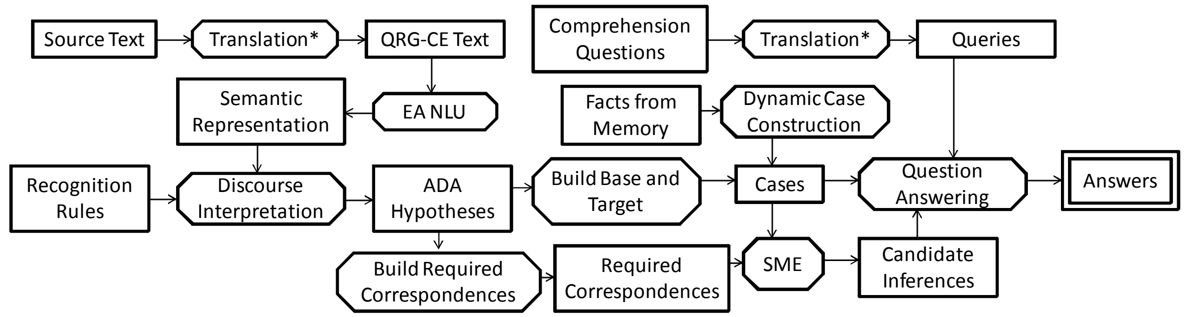

Figure 2: Architecture of the experimental prototype. Processes performed by hand are marked with an asterisk.

found, then there are several possibilities that a reader should explore: Their construal of the base and/or target might be too different from what the author expects, or they should generate a different mapping.

It is important to note that whether a statement combining information from the base and target is considered an intended correspondence versus an intended candidate inference depends to some degree on the reader's state of knowledge. If the target information is unknown, then for that reader, a candidate inference is being introduced. A very active reader may ponder whether it would be a correspondence for a more informed reader, and conversely, whether something an active and well-informed reader views as a correspondence might have been intended as a candidate inference. In both cases, considering the alternate classification would affect the reader's judgment of informativeness, so the distinction between these two types of acts is useful to make. Candidate inferences represent the point of the analogy, what it was set up to convey, and hence distinguishing them seems important.

Block Candidate Inference: These acts alert the reader that an inference that they are likely to make is not in fact correct. For example, "Unlike solar radiation, radiation heat flow reacts in the same way to different colors." If the candidate inference is part of the reader's mapping, then these acts indicate that the reader should mark them as incorrect. A reader with an aggressive processing style who did not generate this inference might explore modifications of their base and/or target to see if they can generate that inference, thereby ensuring they are more in sync with the author's intentions and thus better able to process subsequent statements. These acts are sometimes identifiable by terms such as "unlike," "however," or "you might expect… but" which include one clause expressing information about the base and one clause expressing information about the target. We believe that, like Block Correspondence, these occur relatively infrequently.

4 A prototype implementation

To explore the utility of our analogical dialogue acts theory, we implemented a simple computational model which uses ADAs to learn from instructional texts and answer questions based on what it has learned, synthesized with what it already knows (Figure 1). Our model uses the FIRE reasoning engine, which incorporates SME. The knowledge base contents are extracted from ResearchCyc[1] and extended with other knowledge, including an analogy ontology that lets analogy operations and other forms of reasoning be freely mixed (Forbus et al 2002). In addition to the natural language lexical information built into ResearchCyc, we also use the COMLEX lexicon (Macleod et al 1998) for part of speech and subcat information. For natural language understanding, we use EA NLU (Tomai & Forbus, 2009), which also uses FIRE and the same knowledge base. EA NLU uses Allen's (1994) parser for syntactic processing and construction of initial semantic representations. It uses Discourse Representation Theory (Kamp & Reyle, 1993) for dealing with tense, quotation, logical and numerical quantification, and counterfactuals.

EA NLU is useful for this type of learning by reading experiment because it focuses on generating rich semantic representations. It does so at the expense of syntactic coverage: We restrict inputs syntactically, using QRG-CE (Kuehne & Forbus, 2004), a form of simplified English much like CPL (Clark et al 2005). For example, complex sen-

[1] http://research.cyc.com

tences are broken up into a number of shorter, simpler sentences. Explicit object references (e.g. “the greenhouse greenhouse12” every time the same greenhouse is mentioned) are used to factor out the difficulty of anaphora resolution. EA NLU provides facilities for semi-automatic processing; In this mode, the ambiguities it cannot resolve on its own are presented as choices to the experimenter. This keeps tailorability low, while allowing the system to process more complex texts.

As noted above, we do not yet have a robust model of identification criteria for analogical dialogue acts, so we extended EA NLU’s grammar to have at least one naturally occurring pattern for every ADA. As part of the translation to QRG-CE, texts are rewritten to use those patterns when we view an analogical dialogue act as being present. This allows the system to automatically classify ADAs during processing. Here our goal is to model the processing that must take place once such acts are recognized, since identifying such acts is irrelevant if they are not useful for reasoning. EA NLU’s parsing system produces semantic representations used in its discourse interpretation processing. The ADA recognition rules are used along with EA NLU’s standard discourse interpretation rules to generate ADA hypotheses as part of its discourse representations (Figure 1).

We believe that there are significant individual differences in processing strategies for these acts. For example, some people seem to be quite aggressive about building up mappings, whereas others appear to do minimal work. Consequently, we have started with the simplest possible approach. Here is what our simulation currently does for each of the types of acts:

Introduce Comparison: Builds initial construals of the base and the target by retrieving relevant facts from the knowledge base[2].

Extend Base/Extend Target: The understanding of the sentence is added to the base or target, as appropriate. This decision is made by keeping track of the concepts that are mentioned by statements in each domain, starting with the Introduce Comparison act.

Introduce Correspondence: A required correspondence constraint is introduced for the entities involved, to be used when SME is run for this analogy.

Introduce Candidate Inference: The information in these statements is simply treated as a fact about the target domain. We do not currently change the mapping if a candidate inference in text is not part of the mapping computed.

Block Correspondence/Candidate Inference: Not implemented currently, because examples of these did not show up in our initial corpus.

Analogical dialogue acts are identified via inference rules that are run over the discourse-level interpretation that EA NLU produces. Analogical mapping occurs only at the end of processing a text, rather than incrementally. Statements about the base and target are accepted uncritically, rather than being tested for inconsistencies against background knowledge. These simplifications represent one point in the possible space of strategies that people seem likely to use; plans to explore other strategies are discussed below.

Once the ADA hypotheses are used to construct the base and target domain and the required correspondences between them, this information is used by SME to generate candidate inferences - statements that might be true on the basis of the analogy constructed. The base and target case are expanded using dynamic case construction, which adds knowledge from the KB to fill in information that the text leaves out. For example, a text may not explicitly mention that rain falls from the sky to the earth, taking it as a given that the reader is aware of this.

Example	#O	#A
Gold mining/Collecting solar energy	8	11
Water flow/heat flow	11	12
depth of water in bucket/temperature of house	8	16
Bucket with hole/house leaking heat	4	10
Bucket/Solar collector	5	8
Earth’s atmosphere/greenhouse	7	14
Mean	7.2	11.8

Table 1: Corpus Information. #O/#A = # sentences before/after translation to QRG-CE

5 Experiment

An essential test for a theory of analogy dialogue acts is whether or not it can be used to construct new knowledge from instructional analogies in text. To test this, we extracted a small corpus of 6 instructional analogies from a book on solar energy (Buckley, 1979) and a book on weather (Lehr et al

[2] We use a case constructor similar to `CaseFn` from Mostek et al 2000, but including automatic expansion of rule macro predicates and using microtheory information for filtering.

1987). We simplified the syntax of the original texts into QRG-CE, using the appropriate surface forms for the analogy dialogue acts that we perceived in the text. One of the analogies is illustrated in Figure 1, with part of its translation is shown in Figure 3. Table 1 summarizes properties of the original texts and the simplification process.

Original: Similarly, a full can of water will leak volume from a hole in the side of the can.
QRG-CE: A hot brick brick005 is like a can can001 of water water001. There is a hole hole001 in can can001. The water water001 exits can can001 through hole hole001.

Figure 3: Example of translation to QRG-CE. The specific individuals are added to factor out anaphora processing. Cues to analogical dialogue acts spread across multiple sentences in the original text are combined into single sentences during the translation process.

To test the effectiveness of knowledge capture, 12 comprehension questions similar to those found in middle-school science texts were generated by independent readers of the texts (see Figure 4 for an example). All questions were designed to require understanding the analogy in order to answer them. Moreover, some of the questions require combining information from the knowledge base with knowledge gleaned from the text.

Question: What disappears as the heat leaks from the brick?
Predicate calculus version:

```
(and
 (inputsDestroyed ?d ?ourAnswer)
 (after-Underspecified ?d ?leaving)
 (objectMoving ?leaving heat005)
 (isa ?heat ThermalEnergy)
 (isa ?leaving LeavingAPlace)
 (fromLocation ?leaving brick005))
```

Figure 4: A question for the analogy of Figure 1, in English and the hand-generated predicate calculus generated from it.

Four experimental conditions were run, based on a 2x2 design here the factors were whether or not analogy was used (+A) or not used (-A), and whether what was learned from the text was augmented with information from the knowledge base (+K) or not (-K).

Table 2 shows the results. The system was able to answer all twelve questions when it understood the analogy and combined what it learned by reading with information from the knowledge base.

Condition	# correct	%
-A, -K	0	0
+A, -K	7	58
-A, +K	0	0
+A, +K	12	100

Table 2: Results for Q/A. +/- means with/without, A means analogy, K means facts retrieved from KB

That this was due to understanding the analogy can be seen from the other conditions. The information from the text alone is insufficient to answer any of the questions (-A, -K), as is the information from the KB alone (-A, +K). Analogy by itself over what was learned by reading the passages can handle over half the questions (+A, -K), but the rest require combining facts learned by reading with facts from the KB (+A, +K).

6 Related Work

There has been very little work on modeling analogies in dialogue. One of the few efforts has been Lulis & Evans (2003), who examined the use of analogies by human tutors for potential extensions to their intelligent tutoring system for cardiac function. Recently they have begun incorporating analogies into their tutor (Lulis, Evans, & Michael, 2004), but they have not focused on understanding novel analogies presented via language.

Because EA NLU is designed to explore issues of understanding, it is focused more on semantic coverage than on syntactic coverage. The most similar system is Boeing's BLUE (Clark & Harrison, 2008), which also uses simplified syntax and focuses on integrating language with a knowledge base and reasoning.

Aside from SME, we suspect that the only other current widely tested model of analogy that might be able to handle this task is IAM (Keane & Brayshaw 1988). CAB (Larkey & Love 2003) does not model inference, and hence could not model this task. Although LISA (Hummel & Holyoak, 2003) can model some analogical inferences, the number of relations (see Table 3) in these analogies is beyond the number of relationships it can currently handle (2 or 3).

The first simulation of analogy to use natural language input was Winston's (1982, 1986), which used a simple domain-specific parser in modeling the learning of if-then rules and censors. EA NLU benefits from subsequent progress in natural language research, enabling it to handle a wider range of phenomena.

Example	#S	#BA	#BR	#TA	#TR
Gold mining/Collecting solar energy	8	26	32	4	4
Water flow/heat flow	11	14	21	13	16
depth of water in bucket/temperature of house	8	12	19	9	12
Bucket with hole/house leaking heat	4	14	20	8	6
Bucket/Solar collector	5	13	15	4	4
Earth's atmosphere/greenhouse	7	12	19	11	14
Mean	7.2	15.2	21	8.2	9.3

Table 3: Statistics of base and target domains produced by EA NLU. #S = number of sentences, B/T = Base, Target; A/T = Attributes/Relations

7 Discussion and Future Work

Modeling the roles that analogy plays in understanding language is an important problem in learning by reading. This paper is an initial exploration of how analogy can be integrated into dialogue act theories, focusing on instructional analogies in text. We presented a catalog of analogical dialogue acts, based on an analysis of how the functional constraints of analogical mapping and case construction interact with the properties of discourse. We showed that a simulation using these ideas, combined with a natural language understanding system to semi-automatically produce input representations, can indeed learn information from simplified English analogies, which is encouraging evidence for these ideas.

The next step is to expand the corpus substantially, including more examples of all the ADAs, to better test our model. We also need to implement the rest of the ADAs, and experiment with a wider range of processing strategies.

To better model how ADAs can be identified in natural texts, we plan to use a large-scale web-based corpus analysis. We have focused on text here, but we believe that these ideas apply to spoken dialogue as well. We predict more opportunities for blocking in spoken dialogue, due to opportunities for feedback.

Our goal is to incorporate these ideas into a 2nd generation learning by reading system (e.g., Forbus et al 2007; Forbus et al 2009a), along with other dialogue processing, to better interpret larger-scale texts (e.g., Lockwood & Forbus, 2009). This will be built using the Companions cognitive architecture (Forbus et al 2009b), to more easily model a wider range of processing strategies, and so that the system can learn to improve its interpretation processes.

Acknowledgments

This research was supported by the Intelligent and Autonomous Systems Program of the Office of Naval Research.

References

Allen, J.F. (1994). Natural Language Understanding. (2nd Ed.) Redwood City, CA: Benjamin/Cummings.

Allen, J. F. & C. R. Perrault (1980). Analyzing Intention in Utterances. Artificial Intelligence 15(3).

Buckley, S. (1979). From Sun Up to Sun Down. New York: McGraw-Hill.

Clark, P. & Harrison, P. (2008). Boeing's NLP system and the challenges of semantic representation

Clark, P., Harrison, P., Jenkins, T., Thompson, J. & Wojcik, R. (2005). Acquiring and using world knowledge using a restricted subset of English. 18th International FLAIRS Conference.

Falkenhainer, B., Forbus, K. & Gentner, D. (1989). The Structure-Mapping Engine: Algorithms and Examples. Artificial Intelligence, 41, 1-63.

Forbus, K., Ferguson, R. & Gentner, D. (1994) Incremental structure-mapping. Proceedings of CogSci94.

Forbus, K., Lockwood, K. & Sharma, A. (2009). Steps towards a 2nd generation learning by reading system. AAAI Spring Symposium on Learning by Reading, Spring 2009.

Forbus, K., Klenk, M., & Hinrichs, T. , (2009). Companion Cognitive Systems: Design Goals and Lessons Learned So Far. IEEE Intelligent Systems, vol. 24, no. 4, pp. 36-46, July/August.

Forbus, K., Mostek, T. & Ferguson, R. (2002). An analogy ontology for integrating analogical processing and first-principles reasoning. Proceedings of IAAI-02, July.

Forbus, K. Riesbeck, C., Birnbaum, L., Livingston, K., Sharma, A., & Ureel, L. (2007). Integrating natural language, knowledge representation and reasoning, and analogical processing to learn by reading. Proceedings of AAAI-07 Vancouver, BC.

Gentner, D. (1983). Structure-Mapping: A Theoretical Framework for Analogy. Cognitive Science, 7: 155-170.

Gentner, D., Bowdle, B., Wolff, P., & Boronat, C. (2001). Metaphor is like analogy. In Gentner, D., Holyoak, K., and Kokinov, B. (Eds.) The analogical mind: Perspective from cognitive science. pp. 199-253, Cambridge, MA: MIT Press.

Hummel, J. E., & Holyoak, K. J. (2003). A symbolic-connectionist theory of relational inference and generalization. Psychological Review, 110, 220-264.

Kamp, H. & Reyle, U. (1993). From Discourse to Logic: Introduction to Model-theoretic Semantics of Natural Language. Kluwer Academic Dordrecht: Boston.

Keane, M., and Brayshaw, M. (1988). The Incremental Analogy machine: A computational model of analogy. European Working Session on Learning.

Larkey, L. & Love, B. (2003). CAB: Connectionist Analogy Builder. Cognitive Science 27,781-794.

Lehr, P. E., Burnett, R. W., & Zim, H. S. (1987). Weather. New York, NY: Golden Books Publishing Company, Inc.

Lockwood, K. & Forbus, K. 2009. Multimodal knowledge capture from text and diagrams. Proceedings of KCAP-2009.

Lulis, E. & Evans, M. (2003). The use of analogies in human tutoring dialogues. AAAI Technical Report SS-03-06.

Lulis, E., Evans, M. & Michael, J. (2004). Implementing analogies in an electronic tutoring system. In Lecture Notes in Computer Science, Vol 3220, pp. 228-231, Springer Berlin/Heidelberg.

Macleod, C., Grisham, R., & Meyers, A. (1998). COMLEX Syntax Reference Manual, Version 3.0. Linguistic Data Consortium, University of Pennsylvania: Philadelphia, PA.

Mostek, T., Forbus, K, & Meverden, C. (2000). Dynamic case creation and expansion for analogical reasoning. Proceedings of AAAI-2000. Austin, TX.

Tomai, E. & Forbus, K. (2009). EA NLU: Practical Language Understanding for Cognitive Modeling. Proceedings of the 22nd International Florida Artificial Intelligence Research Society Conference. Sanibel Island, Florida.

Traum, David R. (2000). 20 Questions on Dialogue Act Taxonomies. Journal of Semantics, 17, 7-30.

Winston, P.H. 1982. Learning new principles from precedents and exercises. Artificial Intelligence 23(12).

Winston, P. 1986. Learning by augmenting rules and accumulating censors. In Michalski, R., Carbonell, J. and Mitchell, T. (Eds.) Machine Learning: An Artificial Intelligence Approach, Volume 2. Pp. 45-62. Morgan-Kaufman.