Analogical Dialogue Acts: Supporting Learning by Reading Analogies in Instructional Texts

David M. Barbella & Kenneth D. Forbus

Qualitative Reasoning Group, Northwestern University
2133 Sheridan Road, Evanston, IL, 60208, USA
barbella@u.northwestern.edu; forbus@northwestern.edu

Abstract

Analogy is heavily used 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.

Introduction

Instructional texts and other written explanations often use analogy to convey new concepts and systems of related ideas to learners. Any learning by reading system must ultimately understand such analogies. Here we combine Gentner's (1983) structure-mapping theory with ideas from dialogue act theory (e.g. Traum 2000) to describe a catalog of analogical dialogue acts (ADAs) which capture the functional roles that discourse elements play in instructional analogies. Our goal is to ultimately provide human-like capabilities for understanding such analogies, in order to understand human learning better, and to make better learning by reading systems. We view these goals as closely aligned, since, after all, natural language texts are intended for human readers. We outline criteria for identifying ADAs in text but mainly focus on what operations they imply 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 relevant aspects of dialogue act theory and structure-mapping theory, theories that our theory draws from. We then describe our catalog of analogical dialogue acts, based on a theoretical analysis of what is required to understand instructional analogies. A prototype implementation is described, with an experiment showing ADAs support answering questions based on material learned via analogy, suggesting that ADA processing can be used by reading system to assist in learning. We close with a discussion of related and future work.

Background

Dialogue act theories (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?", but does not have to do so. The surface manifestation might instead be a statement, or an indirect question, e.g. "Do you have the time?" 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 recognized, and a set of implied commitments and obligations for the dialogue participants. There are a wide variety of dialogue act models, but all include some variation of *Inform* (Traum, 2000), which indicates the intent to describe the state of the world. Analogical dialogue acts can be viewed as specializations of Inform.

Understanding instructional analogies requires that the reader set up and perform the analogical mapping intended by the author. In the structure-mapping theory of analogy, 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 *correspondences* identifying how entities and relationships in the base align with the target and a *score*

Barbella, D., & Forbus, K. (2011) Analogical Dialogue Acts: Supporting Learning by Reading Analogies in Instructional Texts. In Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence (AAAI 2011), San Francisco, CA.

indicating the structural quality of the mapping. Mappings also include *candidate inferences*, base statements that are projected onto the target, or vice-versa, suggested by the correspondences of the mapping. In discourse, candidate inferences are often used to convey new information to the learner, or to highlight differences.

The Structure-Mapping Engine (SME) provides a simulation of analogical matching (Falkenhainer et al 1989; Forbus et al 1994; Forbus et al 2002). SME also allows constraints on matches as input. Given a base item b_i and target item t_j, either entities or statements, *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.

In understanding instructional analogies, a learner is expected to draw upon their existing world knowledge, particularly about the base domain. 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. In dynamic case construction, a seed entity or concept is provided as a starting point, and facts which mention it are gathered, typically filtered by some criterion and then further expanded. For example, "The economy of India" might have India as its seed, and facts filtered based on their relevance to economic matters.

A hot brick leaks heat to a cool room. The temperature difference between the brick's temperature and the room's temperature pushes the heat from the brick. The heat escapes until the temperature difference disappears.	*Extend Target*
The hot brick is like a bucket of the water.	*Introduce Comparison*
There is a hole in the bucket. The water exits the bucket through the hole. The water's depth exceeds the hole's height. A volume of water flows from the bucket. The depth difference between the water's depth and the hole's height causes the flow. After the water leaves, the depth difference disappears. When the depth difference disappears, the volume does not exit through the hole.	*Extend Base*
The temperature difference is like the depth difference.	*Introduce Correspondence*

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

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 successfully work through an instructional analogy, a reader must understand that analogy is being employed and what goes into the base and what goes into the target. Formulating the base and target can be complex because what is stated in the text needs to be combined with what the reader already knows. Often readers know more than is relevant to a particular analogy, and there can be multiple mappings between the base and the target. Hence instructional analogies often walk readers through the intended mapping and conclusions to be drawn from it. Analogical dialogue acts identify these constraints, so that they can be used in the reader's understanding process.

Next we describe our proposed analog dialogue acts. We focus here on analyzing what processing needs to be done in response to such acts, laying out a space of possibilities based on human observations and logical possibility, with a specific set of processing choices instantiated in a computer model. We also outline some identification criteria, although that is not the focus of this paper (see future work), based on informal analyses. The first three acts are concerned with identifying the representations to be compared, and the rest are concerned with correspondences and candidate inferences.

Introduce Comparison: Identifies the base and target to be compared. For example, in "We can understand convection by comparing it to water leaking from a bucket." the base is a situation involving leaking water, and the target is the phenomenon of convection. The comparison is often not introduced first, e.g. in Figure 1 the target is described before the comparison introduction.
In Figure 1 the comparison is introduced explicitly in a single sentence, but more complex cases involve combining information across multiple sentences, e.g. parallel sentence structure in subsequent sentences. Determining which of the domains is the base and which is the target requires a non-local assessment about what the 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 candidate inferences can be constructed bidirectionally, an incorrect assessment is not fatal.

Processing an Introduce Comparison act requires producing appropriate cases for the base and target. The target 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 possible construals of a case, 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

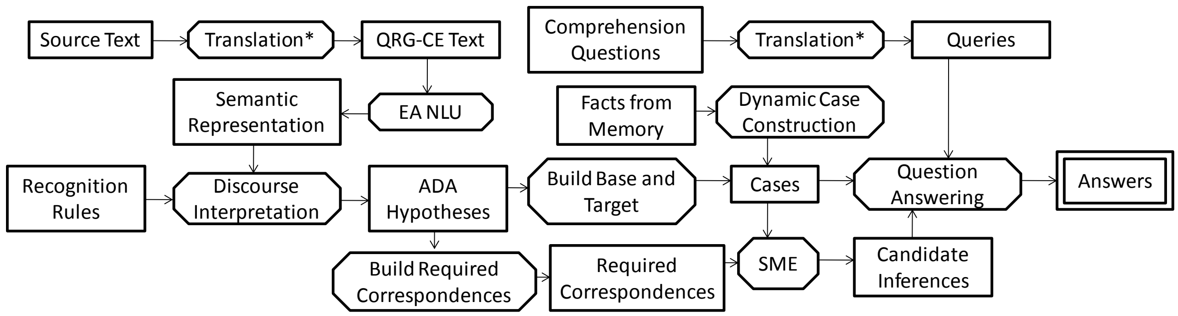

Figure 2: Architecture of the experimental prototype. Processes performed by hand are marked with an asterisk.

entity in the base or target, but which is not a statement about correspondences or candidate inferences. For example, "There is a hole in the bucket." extends the base, and "A brick leaks heat to a cool room." extends the target. Entities mentioned in these acts are added to the construal of the case by retrieving additional knowledge about them, focusing on statements involving other entities in the current construal. If the 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 temperature difference is like the depth difference." indicates that those two entities correspond. Sometimes Introduce Correspondence acts are expressed as identity statements, e.g. "The glass is the atmosphere." in the standard greenhouse/atmosphere analogy. Sometimes these acts are signaled by pairs of sentences, one expressing a fact about the base followed immediately by one about the target, with similar 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.

Block Correspondence: These acts are provided by the author to block a correspondence that a reader might otherwise find tempting, e.g. "The greenhouse door is not like the hole in the ozone layer." In our experience these acts are rare in written text, but show up more frequently as a form of feedback in interactive dialogue, where a learner has the opportunity to describe their current mapping.

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 response is to add an open exclusion constraint (e.g. *excludedBase* or *excludedTarget*, which are versions of the match constraints that do not mention an item from the other description). The excluded item may also simply be removed from the case, 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.

Introduce Candidate Inference: These acts alert the reader to information that the author intended to convey via the analogy. An example is "Just as water leaks faster from a fuller bucket, heat leaks faster from a warmer brick." 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 credence. If the candidate inference cannot be found, then there are several possibilities that a reader should explore: Their construal of the base or target might be too different from what the author expects, or they should generate a different mapping.

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, heat flow that occurs by conduction is unaffected by color." If the candidate inference is part of the reader's mapping, this act indicates that the reader should ignore them. Aggressive readers who did not generate this inference might explore modifications of their base or target to see if they can generate it, 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. In our experience these acts, like Block Correspondence, occur relatively infrequently.

A prototype implementation

To explore the utility of this theory, we implemented a simple computational model (Figure 2) which uses ADAs to learn from instructional texts and then answers questions based on incorporating what it learned with what it already knows. The knowledge base contents are extracted from ResearchCyc[1] and extended with other knowledge, including an analogy ontology (Forbus et al 2002). In addition to the lexical information already in ResearchCyc, we also use the COMLEX lexicon (Macleod et al 1998) for

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

part of speech and subcategory information. For natural language understanding, we use EA NLU (Tomai & Forbus, 2009). 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 experiment because it focuses on generating rich semantic representations. It does so at the expense of syntactic coverage: it is restricted to QRG-CE (Kuehne & Forbus, 2004), a form of simplified English much like Computer-Processable Language (Clark et al 2005). For example, complex sentences are broken up into a number of shorter, simpler sentences (Figure 3). 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 (Barker et al 1998). This keeps tailorability low, while allowing the system to process more complex texts. Due to the large lexicon of the system, the average number of lexical disambiguations per example is 76.1, or 5.5 per sentence. The average number of choices per disambiguation is 2.9. Anaphora resolution, required to identify distinct mentions of concepts as the same entity, is handled automatically (Kamp & Reyle, 1993). Aside from the simplification process and disambiguation, the model runs autonomously from start to finish.

Original: 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

Simplified: A hot brick leaks heat to a cool room. The temperature difference between the brick's temperature and the room's temperature pushes the heat from the brick.
The heat escapes until the temperature difference disappears.

Figure 3: Example of simplification. Complex syntax is eliminated, preserving the original meaning as much as possible.

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 simplification, 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, simplifying the recognition process to focus instead on their use in learning by reading. 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 2). Discourse representation theory forms the basis for how information from sentences is stored and informs how coreference is handled; these steps are automatic. EA NLU works well with non-analogical texts already; for example, it has been used to handle stimulus texts from psychological experiments on blame assignment and moral decision-making (Tomai & Forbus, 2009) and for learning by reading with simplified English and diagrams (Lockwood & Forbus, 2009).

Example	#B	#A
Rubber Ball/Sound Echo	8	5
Gold mine/Collecting Solar Energy	21	32
Bucket of water/Hot brick/	12	13
Water storage/Heat storage	10	13
Phase change materials	8	12
Faucet/Thermostat and furnace	14	18
Stopping a leak/Insulation	7	8
Rain/Sunlight	9	10
Stagnation depth/Stagnation temperature	12	12
Intensity of rain/sunlight	9	12
Power plant/Mitochondrion	6	12
Greenhouse/Atmosphere (Text 1)	9	15
Greenhouse/Atmosphere (Text 2)	5	14
Mean	9.8	13.7

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

We believe that there are significant individual differences in processing strategies for these acts. Consequently, we have started with a relatively simple approach. Here is what our simulation currently does for each type of act:

Introduce Comparison: Builds initial construals of the base and the target by retrieving relevant facts from the knowledge base. We use a case constructor similar to `CaseFn` from (Mostek et al 2000), but including automatic expansion of Cyc's rule macro predicates and using microtheory information for filtering.

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 the analogy. This forces the entities mentioned to be aligned in any mapping that is generated.

Introduce Candidate Inference: The information in these statements is currently treated as a fact about the target domain, i.e., we do not change the mapping if a candidate inference in text is not included.

Block Correspondence/Candidate Inference: Not implemented currently, because examples of these did not show up in our initial corpus. These will result in the addition of an excluded correspondence constraint for the entities mentioned, to be used when SME is run.

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. For simplicity, 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 initial base and target cases, they are expanded via dynamic case construction. This 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 for granted that the reader is aware of this. The expanded base and target, plus the required correspondences between them, are given to SME, which is used to compute a mapping as its interpretation of the analogy. The mapping includes candidate inferences, which are provisionally accepted by the system. Figure 4 illustrates the correspondences from the Figure 1 analogy that the system detects. Note that the Hole-Room correspondence is incorrect; these are aligned by the system because they are loosely structurally similar and neither has a better corresponding entity. (Neither the text nor background knowledge mention any “portal” through which heat leaves the brick, nor the place where the water goes after leaving the bucket.) Incorrect correspondences can lead to incorrect candidate inferences, as some of the inferences in Figure 5 illustrate, e.g. (Bore room9622), the surmise that the room is a hole. Typically candidate inferences are tested for consistency with what is known about the target, and inconsistent instances simply ignored. This model currently does not do such checking, but we intend to add it in the future.

Base	Target
Leave	Leak
Water	Heat
Disappear	Disappear
Hole	Room
Height	Temperature
Difference (Depth)	Difference (Temperature)
Bucket	Brick

Figure 4: Entity correspondences from the analogy in Figure 1 that are detected by the system.

Experiment

An essential test for a theory of analogy dialogue acts is whether or not they can be used to construct new knowledge from instructional analogies in text. To test this, we extracted a small corpus of 13 instructional analogies from a variety of instructional texts for young people (Buckley 1979; Lehr et al 1987; Pfeiffer et al 1964; Scott 1973; Hoff & Rogers 1995) covering topics from a range of fields. We simplified the syntax of the original texts using the appropriate surface forms for the analogy dialogue acts that we perceived in the text. Table 1 summarizes properties of the original texts and the results of the simplification process. One of the analogies is illustrated in Figure 1, with part of its translation is shown in Figure 3. Figure 5 shows the candidate inferences computed by the system for the example. SME computes a structural support score for each inference, as a heuristic for evaluating them. Here the candidate inferences are listed in decreasing order of structural support. (Hence “the brick is a bucket” appears at the bottom.) Note that many of the inferences are simple mappings of features, but others but other others hypothesize useful relationships between the entities aligned.

```
(interrupts (:skolem stop14548) push10389)
(qpropEvent (:skolem stop14548) disappear11616)
(inputsDestroyed (:skolem disappear14302) temperature-difference10213)
(fromLocation leak9149 leak9149)
(depthOfObject heat9226 (:skolem depth13473))
((HypothesizedPrepositionSenseFn Between-TheWord Preposition) temperature-difference10213 temperature10363)
((HypothesizedPrepositionSenseFn Between-TheWord Preposition) temperature-difference10213 (:skolem depth13473))
(DepthDifference temperature-difference10213)
(Height temperature10363)
(heightOfObject room9622 temperature10363)
(causes-Underspecified temperature-difference10213 push10389)
(possessiveRelation heat9226 (:skolem depth13473))
(objectMoving (:skolem flow13242) heat9226)
(from-Generic (:skolem flow13242) brick9105)
(depthOfObject heat9226 (:skolem depth13012))
(heightOfObject room9622 (:skolem height13139))
(possessiveRelation heat9226 (:skolem depth13012))
(possessiveRelation room9622 (:skolem height13139))
(Bore room9622)
(trajectoryPassesThrough push10389 room9622)
(fromLocation push10389 brick9105)
(PassingThroughPortal push10389)
(portalPassedThrough push10389 room9622)
(doneBy push10389 heat9226)
(ExitingAContainer push10389)
(in-UnderspecifiedContainer (DemonstrativeFn (:skolem there)) brick9105)
((LiquidFn Water) heat9226)
(Bucket brick9105)
```

Figure 5: Candidate inferences for the text of Figure 1. The skolems are hypothesized entities.

To test the effectiveness of knowledge capture, 22 comprehension questions similar to those found in middle-school science texts were generated by independent readers of the texts (see Figure 6 for an example). All questions were designed to require understanding the analogy in order to answer them. Moreover, some of the questions require combining pre-existing information from the knowledge base with knowledge gleaned from the text.

Four experimental conditions were run, based on a 2x2 design. 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 information from the text alone is sufficient to answer only one question, with or without information from the KB (-A, -K/+K). Understanding the analogy using just knowledge from the text enables just over a third of the questions to be answered (+A, -K), and allowing the system to add background knowledge from the knowledge base to the analogy raises this to over 80% (+A, +K). This demonstrates that ADAs can help a system learn from an analogy, including harnessing existing knowledge better to improve performance. By incorporating existing knowledge – relational knowledge in particular – the system is able to build a much fuller construal of the base domain. The background relational knowledge can fill in relationships between the entities mentioned in the base that might be "common knowledge", but which are not spelled out in the text for precisely that reason. For example, a text is unlikely to state the relationship between a greenhouse and its roof – that the roof covers the greenhouse – as this information is common knowledge to anyone who knows what a greenhouse is at all. However, this relational knowledge can assist in correctly matching the structure to the structure of the target, where a similar relationship holds, so access to that information allows for more complete mappings, which in turn leads to more and better candidate inferences.

Question: What disappears as the heat leaks from the brick?
Question: The disappearance of what causes the heat to stop exiting the brick?

Figure 6: Questions for analogy of Figure 1. The first question can be answered by understanding the text, the second requires using the analogy to infer causality.

Condition	-K	+K
-A	1 (4.5%)	.1 (4.5%)
+A	8 (36.3%)	18 (81.8%)

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

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 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 in these analogies is beyond the number of relationships it can currently handle (2 or 3). (The average number of relationships in each case was 16.0; the average number of features was 22.8.)

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.

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 further expand the corpus, including more examples of all the ADAs, and to implement full support for blocking acts to better test our model. We also intend to experiment with a wider range of processing strategies, e.g. how valuable is aggressively modeling the domain in terms of better knowledge capture? We are also exploring strategies for filtering implausible candidate inferences.

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, since a listener can explain their mapping as they go.

Our goal is to incorporate these ideas into a larger-scale learning by reading system (e.g., Barker et al 2007; Forbus et al 2007; Forbus et al 2009), along with other dialogue processing, to better interpret larger-scale texts, including texts with diagrams (e.g., Lockwood & Forbus, 2009).

Acknowledgments

This research was supported by the Intelligent & 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).

Barker, Ken, Sylvain Delisle & Stan Szpakowicz (1998). "Test-Driving TANKA: Evaluating a Semi-Automatic System of Text Analysis for Knowledge Acquisition." Proc. of the 12th Canadian Conference on Artificial Intelligence (LNAI 1418). Vancouver, 60-71.

Barker, K., Bhalchandra, A., Chaw, S., Fan, J., Friedland, N., Glass, M., Hobbs, J., Hovy, E., Israel, D. Kim, D., Mulkar, R., Patwardhan, S., Porter, B., Tecuci, D., and Yeh, P. 2007. Learning by reading: A prototype system, performance baseline, and lessons learned. Proceedings of AAAI07.

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., 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.

Hoff, M., & Rogers, M. M. (1995). Our Endangered Planet: Atmosphere. Minneapolis, MN: Lerner Publications Company.

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.

Kuehne, S. and Forbus, K. (2004). Capturing QP-relevant information from natural language text. *Proceedings of the 18th International Qualitative Reasoning Workshop*, Evanston, Illinois, August.

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.

Scott, J. (1973). What is Sound? New York, NY: Parents' Magazine Press.

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.