

Machine and Human Analogical Reasoning for a Case-Method Intelligent Tutoring System

Susann Luperfoy
Eric Domeshek
Eli Holman
David Struck
Brian Glidewell
Ryan Houlette
Stottler Henke Associates, Inc.

{Luperfoy, Domeshek, Holman, Struck, Glidewell,
Houlette}@stottlerhenke.com

ABSTRACT

This paper concerns a project using the case method of instruction to develop an Intelligent Tutoring System (ITS) for military decision-making skills. Our technical objective is actually the authoring shell that lets domain experts and educators enter new cases into a library and then use stored cases in creating lessons that foster analogical encoding in students, a mental process shown empirically to improved acquisition, retention, and transfer of domain knowledge. We have adopted a collaborative view of human-machine interaction in order to construct an integrated cognitive system in which analogical reasoning by the machine supplements and enhances analogical reasoning by the human. In this instance the supported forms of human reasoning are *analogical encoding* by students at instruction delivery time and *analogical recall* by experts at authoring time. The focus of this paper is the design and implementation of the case-authoring component of the authoring shell that assists domain experts in creating new cases and integrating them into the case library.

ABOUT THE AUTHORS

Susann Luperfoy is a Principal Scientist at Stottler Henke Associates where she builds interactive dialog agents for intelligent applications. Her work in applied Artificial Intelligence (AI) research for academia and industry has concentrated on human-machine and human-human dialog interaction, knowledge representation and knowledge acquisition, Machine Translation, Intelligent Tutoring System, and standards for software interoperability and reuse.

Eric Domeshek is an AI Project Manager at Stottler Henke. His PhD from Yale University involved research on cognitive modeling and technology, with a focus on Case Based Reasoning. As Research Faculty at the Georgia Institute of Technology he helped launch the EduTech Institute and directed work on educational applications of AI and CBR. He continued this work as faculty at Northwestern University's Institute for the Learning Sciences and at Stottler Henke where he currently leads a variety of AI and ITS development projects.

Elias Holman is an AI Software Engineer at Stottler Henke. He received his BA in Music Technology from Oberlin College, and Masters in Educational Technology from the Harvard University School of Education. Over the last four years, Mr. Holman has worked on several ITS projects at Stottler Henke, as well as projects focused on web-based collaboration.

David Struck is an AI Programmer at Stottler Henke. He received his BA in Computer Science from Boston College, and has contributed to implementation of several ITS projects at Stottler Henke including substantial work on the human interfaces used by this project.

Brian Glidewell is an Intern at Stottler Henke where he is implementing and/or integrating language-processing components of the case author and synthetic dialog subsystems. He is a student at Massachusetts Institute of Technology working towards a Bachelor of Science in Computer Science and Engineering with concentration in AI.

Ryan Houlette is a Project Manager at Stottler Henke where he is currently lead software engineer on an intelligent track identification and analysis system for the Navy. He holds an M.S. in Computer Science from Stanford University with concentration in AI. He has helped develop a wide range of AI systems, with a particular focus on autonomous agents and intelligent human-machine interfaces.

Machine and Human Analogical Reasoning for a Case-Method Intelligent Tutoring System

Susann Luperfoy
Eric Domeshek
Eli Holman
David Struck
Brian Glidewell
Ryan Houlette
Stottler Henke Associates, Inc.

{Luperfoy, Domeshek, Holman, Struck, Glidewell,
Houlette}@stottlerhenke.com

INTRODUCTION

We are constructing an Intelligent Tutoring System (ITS) that employs the case method of instruction. A key component of the system is the authoring shell that lets educators create new lessons. The authoring shell operates in two primary modes. The case-creation mode lets domain field experts enter and annotate instructional vignettes, stories, or cases, and the lesson-creation mode lets educators turn cases into lessons for distance learning. The focus of this paper is the case-authoring component of the authoring shell and the way it supports the domain expert in assigning semantic attributes to a given case so that an analogical retrieval engine based on the MAC/FAC algorithms of Gentner and Forbus (1989) can recall it later and map it to analogous cases.

BACKGROUND

The case method of instruction is widely used in law, business, medicine, and military training institutions as a form of pseudo experience; cases expose students to domain problems and solutions through stories that illustrate valued domain principles. Each case is a story that describes the circumstances of a problem, the decisions made, the outcome, and often contains an analysis with take-away message from the author. Case method lessons are delivered to students in a sequence of four modes: individual students study the case, small groups prepare for classroom discussion, the professor leads a classroom discussion, and the professor guides a reflective session to review the lesson and offer feedback (Williams, 1991; Hammond, 2002; Benard, 1990).

The customary practice in the schoolhouse version of case instruction is to present each case in isolation encouraging students to assume the mindset and restricted information of the case protagonist. Each

student formulates their own diagnosis and course of action (COA) recommendation for the decision makers in the case, and constructs arguments in defense of that recommendation. In this way, case method lessons offer a form of vicarious practice in high-level decision making through cognitive role-play.

The case method assumes that students can acquire abstract principles from cases and that they can recall them from memory later, when as practitioners they encounter an analogous problem in the field. In actuality however, humans often fail to transfer useful knowledge from learned cases due to the way we encode case experience at learning time and the way we retrieve analogous cases in field practice situations.

This failure to transfer learned knowledge has been attributed, in part, to weaknesses in standard case method instruction. Empirical results (Loewenstein, et al., 1999) show that without explicit guidance, humans do a poor job of recalling useful analogies, and that we do better at analogical storage and retrieval if encouraged to construct abstract analogies at learning time. This means that by presenting the instructional cases in isolation, the conventional approach misses an opportunity to support the sort of analogical encoding that could help the student derive the intended abstract principle from the lesson case.

This project aims to surpass the training effectiveness of conventional case instruction while delivering lessons in a distance-learning context. To accomplish this we are building an authoring tool that assists the expert in creating lessons to foster productive *analogical encoding* in students at runtime. These lessons are meant to guide students through a sequence of modes similar to those experienced in the schoolhouse context—case-preparation by the student, simulated classroom discussion, and a reflective review session—with the addition of coaching by the ITS in

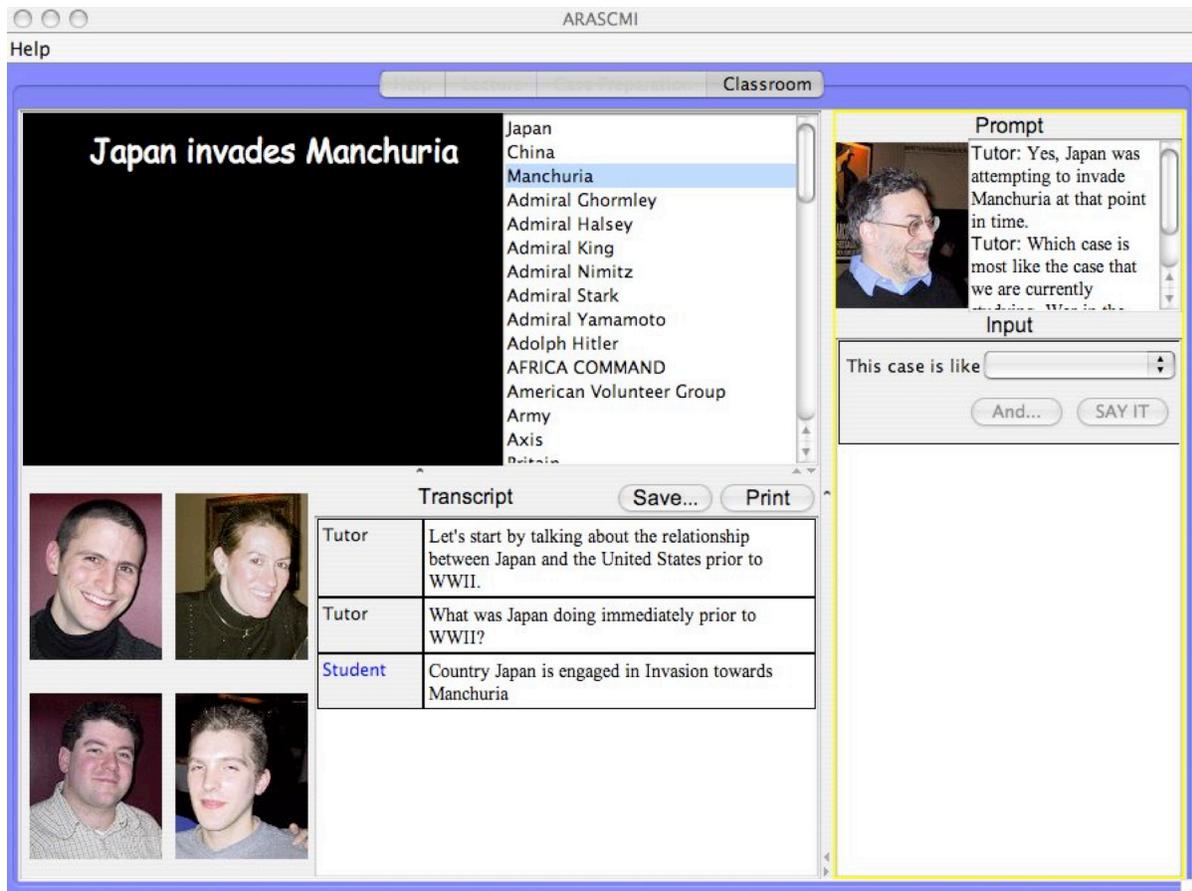


Figure 1: Screenshot from Simulated Classroom Discussion

case comparison as an extension to the reflective session. At lesson creation time, the authoring shell helps instructors overcome their own difficulty with *analogical recall* by locating analogous cases in the case library.

HUMAN-MACHINE COLLABORATION

The Intelligent Tutoring System (ITS) described here is modeled after the schoolhouse protocol for case method instruction. In this ITS, each stage in the life cycle of a case method lesson is accomplished through a collaboration between human and machine so as to benefit from their complementary cognitive strengths.

Case Authoring

In the first stage of development, a domain expert authors a case by importing a document in the form of text, video, or mixed media representation of a story account. Example sources include the lessons-learned collection *66 Battle Command Stories* (hereafter “66 Stories”) published by the Army Research Institute for Behavioral and Social Sciences (Frame, et al., 2000), and the recent Robert MacNamara documentary film on the war with Vietnam, *Fog of War*. In addition to

importing raw media for the story object, the case author must supply sufficient semantic information to enable the machine to find the story in the case library during analogical retrieval based on the abstract principles illustrated in the case.

Whether the case is a published report or a first-person account entered with a computer keyboard by an active field practitioner, there is an unavoidable impedance mismatch between the language of human thought and any formal language that can be manipulated by a machine. This hurdle, which is the bane of any knowledge acquisition endeavor, is illustrated in the following excerpt from the first battle command story in the *66 Stories* collection, written by Major General William S. Wallace.

It just so happened there was one tank crew in a wadi on the north side of Red Pass. One tank crew who had reconnoitered to the nth degree and knew every inch of that wadi. As the OPFOR assault elements and engineer assets approached the obstacle, that one tank crew started taking them under fire and killing

every one of them. They would fire two rounds, duck down in the wadi, move along the wadi, come up to another firing position, fire a couple of rounds and so on. So as far as the OPFOR was concerned, they thought they had a whole platoon, maybe a company, off on their left flank, and it was one single tank that was using the terrain very, very well, and had boresighted their weapon system.

In this short paragraph we encounter numerous challenges for any state-of-the-art Natural Language Understanding (NLU) systems that might be used to perform or aid knowledge acquisition of its semantic content. There are proper names and acronyms (*Red Pass, OPFOR*); definite and indefinite pronominal references (*they, them, their, it, every one of them*); esoteric or domain-specific lexical items (*wadi, reconned, boresighted, platoon*); domain-specific usage of familiar lexical items (*fire, rounds, tank, company, flank, duck*); idiomatic phrasal elements (*to the nth degree, very very well*); an incomplete sentence (*One tank crew who had reconned to the nth degree and knew every inch of that wadi*); a nominal reference that is syntactically singular but semantically plural¹ (*crew*); and pleonastic uses of *it* and *there* (*it just so happened, there was one tank*). This is to say nothing of the references to events, states, and processes, or the relations implied or asserted among them, including shifts in tense and aspect (e.g. from past to narrative present; habitual usage in *they would fire*). Nor have we mentioned metaphor, metonymy, and other figurative linguistic devices that occur elsewhere in this two-page story.

A collaborative view of case authoring allows us to lower the height of the NLU hurdle by recognizing that full understanding of the story meaning is neither possible nor necessary for the ITS application in question; it is sufficient to encode only those properties of a case that are essential for analogical retrieval. Further, we gain a leg up on the hurdle by empowering the human domain expert to decide what those essential properties are. The human performs the open ended NLU, as it were, and determines what is salient about the case. The machine captures those determinations in the form of semantic metadata, stored with the case document in a machine accessible formal language that supports later retrieval based on semantic content.

The role of the machine during case authoring is to hide the low-level processing of case documents, elicit

¹ In American English dialects we say that the committee *has* reached *its* decision, whereas some British dialects would say that the committee *have* reached *their* decision.

human contributions of essential case semantics, and capture the case story and semantic metadata in a persistent database, using a formal knowledge representation that supports machine reasoning for later retrieval and evaluation.

The domain expert identifies a story document and the machine extracts from the text component of the document all references to people, locations, organizations, and events. These entities and events are presented to the user in the GUI of the case-authoring tool for selection or discard. The selected entities are then loaded into a matrix that lets the user visualize them as rows and columns and assign unary and binary relations in the cells of the matrix. The information content of the matrix is reflected in the underlying knowledge base representing the agents, goals, plans, actions, and consequences portrayed in the story text, in order to capture the narrative sequence of the account. Since the first version of this paper was submitted we have added GUI features that allow the author to follow back pointers from extracted events and entities to the original text. This was a prerequisite to our current development effort to let the user assert higher-order relations (e.g., causation, dependence, equivalence) between events and/or entities.

Lesson Authoring

The second stage of collaboration involves a human in the role of educator. This individual may be a course instructor, a curriculum developer, or the same domain expert who has authored the relevant cases using case authoring mode. With learning objectives in mind, the educator creates a lesson by selecting from the case library one or more cases that illustrate the lesson concept. The lesson author designs the student experience of these cases by specifying study questions, good analogies and false analogies, dialog behaviors of the synthetic agents in the simulated classroom discussion, and ITS interaction with the student during the reflective session.

This project takes as givens that humans can learn from stories (Domeshek, 1992; Schank, 1991) and that for higher-level skills we learn even better if presented with two stories that share an abstract analogy (Loewenstein, et al., 1999). To take advantage of this phenomenon in our runtime ITS, it is not sufficient to have captured a library of cases that can be presented in isolation. We need to also provide for authoring of pairs (or sets) of case stories that share an abstract principle in common. This, too, is best handled as a collaboration between human and machine.

In composing lessons from a library of cases, the human author supplies educational expertise including which concepts are essential to the course syllabus and

the immediate lesson, which cases can be used to teach those concepts, which cases are analogous to each other with respect to the relations in question, and which counter-analogies are likely to arise during discussion as a result of predictable student misconceptions. The machine's role in the lesson authoring collaboration is to help the human to manage the case library by presenting to the user, candidate cases that appear to be analogous with respect to the learning concept. The retrieval engine can justify its conclusion by pointing back from each metadata assertion to the corresponding segment of text that gave rise to it, allowing the human to evaluate the strength of the analogy.

It should be noted that a story could contain learning potential beyond what was intended by the author or even the instructor who presents the story as a case study. Indeed, the unique background of a student can make it possible to gain more from a case story than even the instructor intended. This means that a given case may serve in multiple lessons for diverse courses designed by different educators. It also means that case reusability can benefit from continued semantic metadata encoding even after a case is in use, with distributed authors contributing to a shared case repository available for refinement, correction, and extension.

Lesson Delivery to Student

In the final stage of case instruction the lesson is delivered to the distance-learning student. The student first calls up the lesson, reads the case(s) and develops quantitative and qualitative analyses to support a recommended course of action. With arguments in hand, the student then moves to a simulated classroom discussion where synthetic dialog agents play the roles of classmates and professor. Figure 1 contains a screen captured from the prototype ITS in the midst of a simulated classroom discussion.

The human analogical reasoning facility is a powerful learning mechanism and is often under-exploited by pedagogical methods (Gentner, 1983). To help students acquire and retain the desired abstract principles, this ITS simply prompts them to compare selected cases. There are students who may look for analogies on their own without prompting. For the majority who are not so predisposed, the role of the machine is to encourage analogical encoding by presenting comparisons to the student and offering feedback when the student proposes analogs recalled from memory.

This guidance is offered at the end of the simulated classroom discussion in a reflective dialog session. As

part of concluding the case discussion the ITS, via the persona of the synthetic professor, guides the student through consideration of the case in a larger domain context. The professor can first present any information from the domain world beyond the case narrative, or can pose hypothetical alternatives to the facts of the case, prompting the student to revisit their COA decisions in light of the new information. Then, to promote analogical encoding, the student is encouraged to consider analogous cases through any of the following analogical encoding tasks (in descending order of expected difficulty and learning value):

1. Prompt the student to name another case that is analogous to the current case, and have them explain their answer in terms of structural relations between the two cases
2. Present another case and ask whether/how it is analogous to the current case
3. Present two external cases and ask which is a better analog for the current case
4. Present two analogous cases and invite the student to compare them
5. Present two cases and tell the student to identify the analogy between them
6. Present the key principle(s) of the lesson/cases(s) without requesting a student response.

The structure and content of this reflective session will have been defined at lesson-creation time, and the full version of the runtime ITS will be instrumented to record student input selections and timing information. This information is aggregated for use by the educator in assessing the training effectiveness of the lesson they have designed, or its appropriateness for a specific training audience.

Summary of Machine Contributions

In the various use cases involving authors and students over the life cycle of a case-method lesson, the machine serves several roles. These roles for runtime modes of case preparation, simulated classroom discussion, reflective session, and system help are embodied as distinct user-interfaces to achieve each of the intended student experiences. For the ITS to engage the student in dialog about the relationships between cases during the reflective session, it must be able to present good analogies and assess proposals offered by the student. This ITS behavior requires human-like dialog capability and places several demands on the underlying knowledge representations and the authoring procedures that create them.

First, the case library must contain cases that share instructionally valuable relational structures. For instance, drawing from the *66 Stories* collection, there may be two or more cases in which the commander failed to empathize fully with the enemy, or there may be multiple cases in which proportional response was called for, one of which demonstrates failure to do so. The case library should also contain cases that are only superficially similar and thus fail to constitute useful analogies. For example, two mission stories that lack relational similarity but both involve simple references to the OPFOR, disruptive changes in weather, or unexpected terrain features can have instructional value as counter-analogies.

Furthermore, it is not enough for the cases with instructional value to merely exist in the case library; the lessons presented to the student must index pairs (sets) of cases that represent analogies and counter-

analogies. This indexing occurs at lesson creation time. This is when the machine's role is to retrieve and suggest analogs and then help the author design an engaging lesson around the cases by providing examples and editing tools for entering study questions, key talking points, and dialog behavior of synthetic agents.

IMPLEMENTATION STATUS

We have in place the design for an end-to-end ITS and authoring shell with four author modes (case authoring, lesson authoring, lesson review, and system help) and four student modes (case preparation, classroom discussion, reflective session, and system help). The diagram in Figure 2 illustrates the three stages of a case lesson and the three user roles, where Domain Expert and Educator can refer to the same individual. There is as of this writing a running prototype with a preliminary version of each of the four student modes,

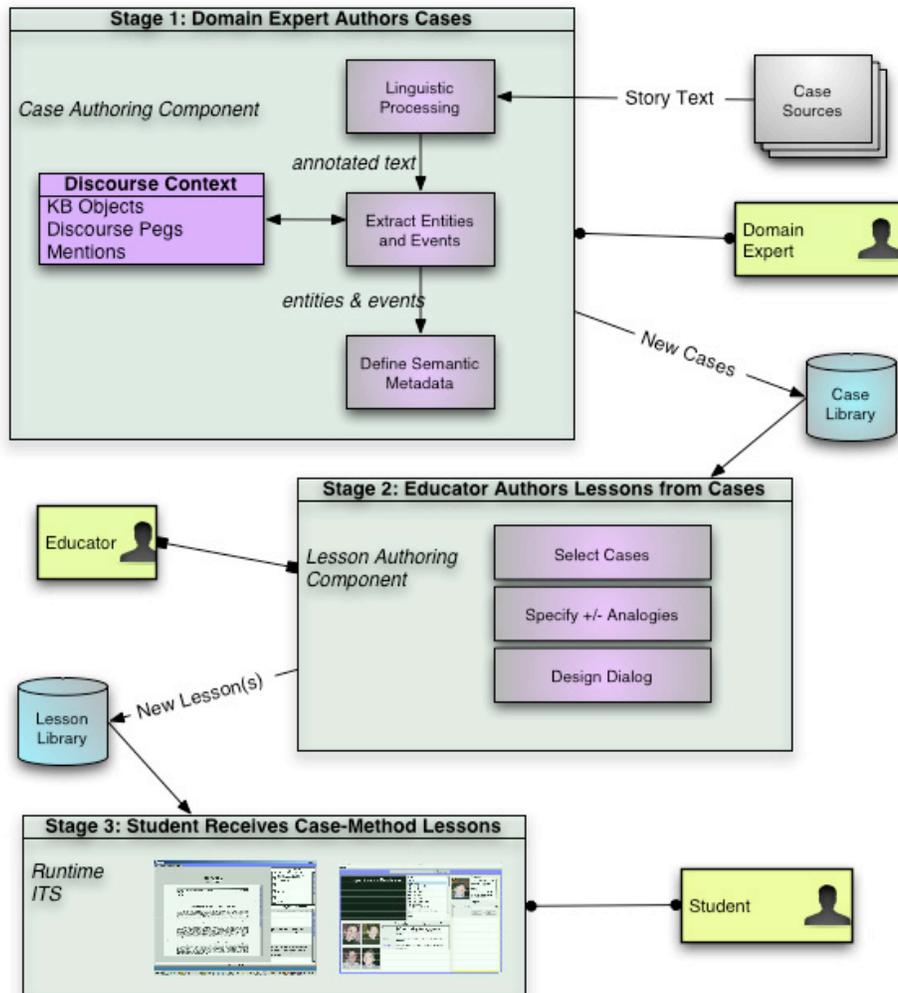


Figure 2: Architecture for Authoring and Instruction

and all but the lesson review mode for authors. The case library contains three cases from the Joint and Combined Warfare School Intermediate Program (from the Joint Forces Staff College) and a single lesson serves as the prototype for exercising the knowledge representation and reasoning algorithms.

We are using the *66 Stories* collection as a test data corpus for the case authoring process. This source is particularly useful because it provides a large amount of naturally generated text containing first-person accounts known to have instructional value for the military decision-making domain. Furthermore, each of the *66 Stories* includes highlighted text encapsulating the author's primary message. These can be used as benchmarks for testing the ability of the human-machine team to extract the defining information for each case and encode it into the semantic metadata.

Figure 2 depicts the process that transforms stories into cases and cases into lessons. For the linguistic processing inside the case-authoring component (the topic of this paper) we are using a combination of in-house and open source modules. We use the GATE system to extract Part Of Speech (POS) annotations from any raw text, pdf, or html document (<http://gate.ac.uk/>). The gazetteer feature of GATE performs lexical insertion to give estimations of animacy, personhood, and other semantic features of referring expressions. The Link Parser (Temperley, et al. 2004) takes a different approach to text analysis than GATE. Link is our current choice for syntactic analysis of student input during Stage 3 instruction delivery. The WordNet lexical database for English serves as the primary lexical resource for both case text processing and dialog and it is accessed using the Java WordNet Library (JWNL) interface. (See <http://jwordnet.sourceforge.net/>) WordNet data is helpful in resolution of nominal references and the distributive/collective semantic distinction of nominal references to groups (*crew, enemy, platoon*).

For formal representation of semantic metadata on stored cases, as well as interpretations of language at runtime, we are using our in-house General Representation and Inference Storage Tool (GRIST). GRIST (Domeshek, 2004) contains the underlying domain ontology that supports case semantic metadata representation and analogical retrieval that is conducted according to the Structure Mapping Engine (SME) algorithm (Forbus, 1997; Gentner, et al., 1989; Faulkenheimer et al., 1989). The framework used for extraction, representation and updating of discourse context during analysis of a story text, is outlined Luperfoy, et al. (1998, 2003).

SUMMARY OF FINDINGS

The case method is especially useful for teaching higher-level skills that could be characterized as intuitions for the domain. These are skills for which no procedural recipe can be given. Rather, case lessons offer problem-solving practice with performance feedback to help develop these more abstract skills that are required for thinking about problems in the way an expert practitioner would.

This project was motivated by recent observations in empirical psychology indicating that students acquire, retain, and transfer understanding of abstract principles more effectively when they are encouraged to compare analogous cases at learning time. To capture value from this result, the Intelligent Tutoring System (ITS) we are developing presents cases to students one at a time, as in the schoolhouse setting, but then goes a step further, by engaging students in dialog that encourages them to consider multiple cases and encode analogies across cases share an abstract principle.

Having accepted the importance of analogical encoding for learning, this project requires a method for machine retrieval of analogous cases during lesson delivery. But a machine can only retrieve what it has been given and the task of converting human knowledge into formalized machine knowledge is an unsolved problem in AI. In the process of designing an authoring tool to overcome the knowledge-acquisition hurdle and allow course developers to construct the desired sort of case method lessons, we have discovered and acted on some additional realizations:

- Authors themselves will have difficulty retrieving useful analogies from memory and can benefit from a similar sort of machine assistance for analogical retrieval as afforded to students at runtime.
- Authoring involves two distinct activities that require distinct human expertise as well as tools optimized to the task activities. The case author can be any domain expert or field practitioner with a story to tell whereas the lesson author is accountable for turning cases into lessons with demonstrable educational value
- The authoring tool for importing case documents, stories, video clips, and even first-person accounts entered on computer keyboard should perform surface lexical and syntactic analysis of large corpora of language, and it should do this in the background. The result should be presented to the user for semantic analysis, fully indexed to the source text so that the human can access the original text as context for determining salience.
- A single case can be useful for illustrating diverse principles in multiple lessons. So authoring of

cases and lessons can be pursued as a collaborative human endeavor involving participation of domain experts and educators distributed geographically and temporally, who contribute as loosely coupled teams enhancing and reusing the cases and lessons originated by others.

- The runtime system should be engaging enough to keep students involved without neglecting to reproduce the attested value of the schoolhouse case method. With that in mind, we are working to incorporate peer teaching/learning and practice formulating logical arguments.

Success of this project depends on several factors yet to be fully understood. First, nothing about our human-machine collaboration approach is claimed to completely solve the knowledge acquisition bottleneck. The resulting case authoring system is not a solution to Natural Language Understanding (NLU). Rather, we are building tools that assist the human in processing the text component of a story, report, or emotional account of an incident in Baghdad. The robust text processing algorithms and the user interface design for extraction help import the text and extract references to entities and events. Still, the human expert retains control over decisions on what is salient about the case relative to their objectives for the domain and the students. It is the lesson author who decides how to package the lesson and where the case fits into a lesson. Finally, the system is being developed based on user-centered engineering including a cognitive task analysis of lesson authors, but until it has been fully validated with educators and students *in situ*, the claims and assumptions presented here are to be interpreted as intermediate research results.

ACKNOWLEDGEMENTS

The authors would like to acknowledge the contributions of James Ong who has been an advisor to this project from the outset. Anonymous reviewers made useful suggestions that have helped clarify the presentation of ideas.

REFERENCES

- Benard, B. (1990). A Case for Peers. Portland, OR: Northwest Regional Educational Laboratory. ED 327 755.
- Domeshek, E. (1992) Do the Right Thing: A Component Theory for Indexing Stories as Social Advice. PhD dissertation, Yale University
- Domeshek, E. (2004) Phase-2 Final Report on an Intelligent tutoring System for Teaching battlefield Command Reasoning Skills. Technical report #1143, US Army Research Institute for the Behavioral and Social Sciences.
- Faulkenheimer, B., K. Forbus and D. Gentner (1989) The Structure-Mapping Engine: Algorithm and Examples. *Artificial Intelligence* 41 pp 1-63.
- Forbus, K.D. (1997) MAC/FAC Manual and API Version 10.24.97 pp. 1-11.
- Frame, A. and J.W. Lussier (eds) (2000) 66 Stories of Battle Command. *US Army command and general Staff college Press. Fort Leavenworth, Kansas.*
- Gentner, D. (1983) Structure Mapping:A theoretical framework for analogy, *Cognitive Science* 7(2).
- Gentner, D. and Forbus, K.D. (1989) MAC/FAC: A Model of Similarity-based Retrieval Report from ONR Contract No. N00014-89-J-1272.
- Hammond, J.S. (2002) Learning by the Case Method. Harvard Business School case #9-376-241.
- Loewenstein, J., Thompson, L., & Gentner, D. (1999). Analogical encoding facilitates knowledge transfer in negotiation. *Psychonomic Bulletin & Review*, 6, 586-597.
- Luperfoy, S., D. Loehr, D. Duff, K. Miller, F. Reeder, and L. Harper (1998) "An Architecture for Dialog Management, Context Tracking, and Pragmatic Adaptation in Spoken Dialog Systems". In proceedings 36th Annual Meeting of the Association for Computational Linguistics.
- Luperfoy, S., E. Domeshek, E. Holman, D. Struck (2003) "An Architecture for Incorporating Spoken Dialog Interaction with Complex Simulations" In proceedings Interservice/Industry Training, Simulation, and Education Conference.
- McAllister, Elizabeth (1990). Peer Teaching and Collaborative Learning in Language Arts. Bloomington, In: EdInfo Press. ED 325 818.
- Schank, R. and R. Abelson (1977) Scripts, Plans, Goals, and Understanding. Lawrence Erlbaum Associates, Hillsdale, NJ.
- Schank, R. (1991) Tell Me a Story. Charles Scribner and Sons, New York, NY.
- Temperley, D., D. Sleator and Lafferty, J. Link Grammar Syntactic Parser for the English Language (2004). <http://www.link.cs.cmu.edu/link/>
- Williams, Susan M. (1991) Putting case-based instruction into context: Examples from legal and medical education. Technical Report, Learning Technology Center, Vanderbilt University. (Republished (1992) *Journal of the Learning Sciences*, 2, 367-427.)