Position Paper for AAAI 2006 Fellows Symposium
Kevin D. Ashley, University of Pittsburgh

Computational modeling using AI techniques has long been urged as a tool for empirically investigating issues of interest to non-AI domain experts in a variety of fields, such as biochemistry, medicine, law, ethics, and philosophy. AI, it was hoped, would add tools to domains already susceptible of scientific methods, or introduce scientific methodologies to domains that never had them. When the final history of AI is written, it will be interesting to see how well-founded these hopes have been. It may be intriguing for interested Fellows to discuss the attempts that have been made so far, their successes and failures, and even the criteria for evaluating their success. For instance, have the results of AI investigations been published in non-AI research journals, have they been accepted by non-AI domain experts, to what extent have non-AI-related funders provided support for the work, etc.?

Although my efforts of this type have achieved only modest results, I have been interested in using AI computational models of case-based reasoning empirically to investigate semantic relationships between abstract normative principles and fact-specific cases. Moral and legal philosophers have long observed a dialectical relationship between them: the abstract principles inform the decisions of specific cases, but the decisions, in turn, elaborate the principles' meaning. I thank my dissertation adviser, Edwina Rissland, for "turning me on" to this insight a long time ago (See Ashley & Rissland, 2003, p. 31).

An article in *Jurimetrics* (Ashley, 2004) summarized one foray in using AI computational modeling as an empirical methodology to investigate a normative phenomenon. In the SIROCCO program, my former Ph.D. student Bruce McLaren and I represented over 180 ethics cases in which a professional engineering association's board of ethical review applied ethical principles from the association's code of ethics in making decisions. Given new problems, SIROCCO retrieved code provisions and past cases relevant to analyzing the new problems. Experiments with SIROCCO demonstrated empirically that over time the board's case decisions operationalized the abstract code principles, fledging out the meaning of the code provisions in a way that the program could use to retrieve relevant cases more effectively. By comparing SIROCCO's ability to retrieve relevant ethics principles and cases with and without such operationalizing information, we objectively demonstrated that cases make a dialectical contribution to abstract principles, improving a program's representation of what the principles and cases mean. *Jurimetrics*, the ABA Section of Science and Technology journal, is a peer-reviewed law journal.

In a second foray, I used another AI computational model to investigate a philosophical argument concerning the meaning of intermediate legal concepts like "ownership" or "trade secret" that are employed in legal rules. In response to a philosophical skeptic's provocative argument, moral scholars have argued that such legal concepts have a creative or extending function by which they provide guidelines for the handling of new and problematic cases. Using another computer program, the Issue-Based-Prediction Algorithm (IBP), my current Ph.D. advisee Stefanie Brüninghaus and I tested this hypothesis empirically. IBP employs a hypothesis-testing methodology that is somewhat analogous to testing a scientific hypothesis. The hypotheses, however, concern the legal outcomes of factual disputes, and the "experimental data" are the past decided cases in a database developed originally for CATO, a tutoring program for case-based argument. Given the factors of a new dispute, IBP identifies legal issues, determines which party is favored for each issue, combines the analyses into a prediction, and explains its reasoning. In the process, IBP formulates hypotheses about who should win issues, tests the hypotheses against cases in its database, and attempts to explain away counterexamples to the hypothesis. In an experimental evaluation, IBP attained significantly higher accuracy (91%) than a variety of alternative statistical, machine learning, and CBR prediction approaches (Brüninghaus & Ashley, 2003a; b; Ashley & Brüninghaus, 2006 (also in *Jurimetrics*)). When we "turned off" the issue-based concepts in IBP's computational model of the legal domain, we demonstrated that legal concepts make a significant contribution to predictive accuracy in framing the most useful hypotheses to test, empirical evidence of the concepts' guidance function (Ashley & Brüninghaus, 2003; 2006). Conversely, by turning off the case-based procedures for theory-testing and explaining-away counterexamples, we demonstrated cases' contribution of a means for resolving conflicting reasons (Brüninghaus & Ashley, 2003a; b).

As reported in a Nectar paper at AAAI-06, in the SMILE project, Steffi Brüninghaus has gone on to test how augmented text representations can enable a machine learning program to learn case indexing concepts so that IBP can analyze cases and predict their outcomes *directly from a textual* description of the case's facts. She has developed techniques for using domain knowledge, such as the roles of parties and

objects in a case and schematic representations of who did what to whom, to augment the representation of the textual case fact descriptions. (Brüninghaus & Ashley, 2006).

In other recent work (Pinkwart, et al., 2006) I examine the practice and teaching of skills of legal argumentation and moral imagination as evidenced in oral arguments before the United States Supreme Court, where the Justices famously pose hypotheticals to explore and test the competing advocates' normative propositions, their application to current problems and consistency with past decisions. Initially, my CMU colleague and former Ph.D. student, Vincent Aleven, and I are examining how law students self-explain selected Supreme Court hypotheticals. We are experimenting with techniques for helping students to better understand, reason with, and respond to hypotheticals in legal arguments and as a tool for legal analysis. We will explore adapting for this purpose the CATO program, a computerized instructional environment for teaching skills of case-based legal argument. In work with another graduate student, Ravi Desai, I was able to show that by engaging law students in more realistic, dialectical arguments based on CATO's model, the students were better able to transfer their argumentation skills to an unfamiliar legal domain (Ashley et al., 2002). In other words, what they learned about analogizing and distinguishing cases in one legal domain, they were able to transfer to a new, unfamiliar legal domain. This is a key result in legal pedagogy where laws change and professors struggle to prepare students to apply their skills in new legal areas.

In two other projects, I have developed empirical tools for assessing improvements in students' abilities to analyze problems normatively. In the first project with my colleague Matt Keefer, we developed a coding scheme and tutoring intervention designed to assess and improve the extent to which engineering students considered alternative factual conditions that would impact the moral analysis (Keefer and Ashley, 2001). In a subsequent, more detailed study of students' moral analyses of problems, my ethicist colleague Rosa Pinkus and graduate student Ilya Goldin and I have developed a coding scheme for identifying the extent to which students label, define, and use important ethical concepts as well as employ professional knowledge to frame moral issues, identify different perspectives in the case and move flexibly among them, and identify analogous cases and articulate why the cases are analogous. (Goldin, et al., in preparation).

In summary, in my work I have tried to demonstrate how AI computational models can supply a scientific methodology for investigating phenomena of interest to legal scholars and ethicists. I would be glad to learn of other Fellows' efforts to do the same in these and other domains.

References

Ashley, K. (2004) "Capturing the Dialectic between Principles and Cases." In *Jurimetrics* 44:229.

Ashley, K. and Brüninghaus, S. (2006) "Computer Models for Legal Prediction." *Jurimetrics*. 46:309.

Ashley, K. and Brüninghaus, S. (2003) "A Predictive Role for Intermediate Legal Concepts," *Proc. 16th Annual Conf. on Legal Knowledge and Information Systems*, Jurix-03. 153. IOS Press, Amsterdam.

Ashley, K., Desai, R. and Levine, J.M. (2002) "Teaching Case-Based Argumentation Concepts using Dialectic Arguments vs. Didactic Explanations". *Proc. Intelligent Tutoring Systems Conf.*, ITS '02 (S.A. Cerri, G. Gouardères, F. Paraguaçu, ed.) LNCS 585-595. Springer-Verlag: Berlin.

Ashley, K. & Rissland, E. (2003) "Law, Learning and Representation". *Artificial Intelligence.* 150, 17-58.

Brüninghaus, S. and Ashley K. (2006) "Progress in Textual Case-Based Reasoning: Predicting the Outcome of Legal Cases from Text", 21st Nat. Conf. on Artificial Intelligence (AAAI-06) Boston. July.

Brüninghaus, S. and Ashley, K.D. (2003a) "Combining Case-Based and Model-Based Reasoning for Predicting the Outcome of Legal Cases". *Proc. 5th Int'l Conf. on CBR* (ICCBR-03). 65. Springer, Berlin.

Brüninghaus, S. and Ashley, K.D. (2003b) "Predicting Outcomes of Case-Based Legal Arguments." *Proc. 9th Int'l Conf. on Artificial Intelligence and Law* (ICAIL-03). pp. 234. ACM Press.

Goldin, I., Pinkus, R. and Ashley, K. (in preparation) "Sensitivity and Reliability of an Instrument for Assessing Case Analyses in Bioengineering Ethics Education."

Keefer, M. and Ashley, K.D. (2001) "Case-based Approaches to Professional Ethics: a systematic comparison of students' and ethicists' moral reasoning." *Journal of Moral Education*. 30:4 377-398.

Pinkwart, N., Aleven, V., Ashley, K., Lynch, C. (2006) "Toward Legal Argument Instruction with Graph Grammars and Collaborative Filtering Techniques" 8th Int'l Conf. on *Intelligent Tutoring Systems*, (ITS 2006). Springer, Berlin.