Electronic Commerce is an Intriguing Domain for AI Learning Theory

(A Statement of Interest)

Leona F. Fass

Gedanken-Experiments for Cyberspace Learning

As a theoretician confronted with cyberspace we do not ourselves build systems or implement applications but we *black box experiment* with them as an informed and sometimes surprised user. Our specific interest in the area of Artificial Intelligence for Electronic Commerce comes from this "experimenting user" mode. As we employ EC systems that now exist, we cannot help but analyze them. Based on our work in behavioral modeling and learning, we may celebrate EC system successes *or* detect what we hope to be correctable system flaws.

Adaptation and learning, and criteria for assessing results obtained, are issues we have considered in our theoretically-oriented research. Within that context we have investigated the representation and construction of complex systems and informational domains. We have sought to identify necessary components that might interact to produce a specified behavioral goal. At first we expected "perfect" implementations and strove to determine processes and processors for behavior that would be "correct". But our current work is more (realistically) concerned with achievement of goals and results that are *feasible.* We do not expect perfect results (which may not exist), and we may be satisfied to test a process or system and determine it is *in*correct. Discovery of (inevitable) errors and adaptive correction is just another approach to learning and intelligence.

Positive EC Learning: (Potentially) Successful Constructive Deployments

Reactive agents with goal-oriented behavior exhibit just such the intelligent activity that we have described, above. We ourselves have observed this in our daily interactions with EC.

> **Example 1.** We have successfully trained and retrained the speech recognition component of an electronic stock-brokering system to understand our oral requests for information about the stocks we own. When we began our efforts the system could not distinguish "AVISTA" from "Assist Technologies", but now it can.

By considering such examples we have come to believe that our theoretical research into positive, constructive approaches to learning indeed is applicable in the very practical EC domain. There masses of information must be represented, analyzed and processed, with results (stock sales, plane ticket issuance, dinner orders, . . .) obtained. Such EC considerations are just like the issues we have dealt with in the field of learning.

Applications to EC that are related to our learning research include data mining and knowledge discovery (as in link analysis); intelligent user modeling; information retrieval and integration (especially inference and machine learning); and intelligent agent development, as we just noted. We expect our research to be applicable to the design of reactive EC agents that learn from experience, and to designing cooperating agents and federated agent systems, that decompose EC problems and solve their components (as needed) autonomously or interactively.

> **Example 2.** Such an EC instance we have encountered is an airline system that makes reservations, issues electronic tickets, bills credit cards, and orders special meals. Upon request the system confirms all of the foregoing by fax to the user.

Practical EC agent deployment most certainly has similarity to our theoretical construction-of-behavioral-realizations work.

Adversarial EC Learning: Identifying Correctable Problems and Perils

Our original constructive learning research was extended into the "adversarial" approach of error detection, and so too can this research be applied to EC system design. A complex process(or) is expected to have defects, and their discovery should lead to creating a better, "less flawed" system. In our "experimenting user" mode we have inadvertently discovered some EC *problems and perils* in deployed systems. These could be alleviated through the use of AI learning techniques.

We have identified *a need for adequate testing* to resolve premature deployment problems:

> **Example 3.** We have found that the "cutting edge" voice-broker system for a major mutual fund company, rather than positively impressing shareholders, almost never works. It requires almost constant human intervention and loss of staff/customer patience and time, possibly doing more harm than (its intended) good.

We have identified *a need for development of protection-, back-up- and integrity-preserving-agents*:

> **Example 4.** A well-known brokerage house that features electronic trading lost all records of security-trading orders made before 10 a.m. last September 28, due to a "system upgrade". As an AI system analyzer, we were amused, although as a stockholder, we were not.

> **Example 5.** When the airline system discussed in Example 2 works, it works very well. But when one of

its "cooperating agent" components fails, so may most others. We have lost our ticketing record, seats and meal orders, all at the same time.

We have identified *privacy, security and legal issues* raised by an EC site's planting of "smart cookies", and similar issues for "too smart" browsers:

Example 6. We have permission to engage in electronic Web-trading of securities via a discount brokerage house, and can access our own account, using our own password, from any public machine. We did so, from a public library computer one day, shortly after a new "remembering" browser version was installed. Nine hours later, we returned to a library workstation and began to type in the URL for a site whose URL began with the same few symbols as that of the brokerage house. We immediately found ourselves "transported" erroneously, not just into the home page of the brokerage house but, without a password, directly into our account! *Anyone* using that machine throughout the day could have done the same: accessing our account, trading and making bank transfers and cash transactions.

The tech support person we contacted at the brokerage house told us that the "cookie" planted that day to "help us" (i.e., to track our Web-page activities) should have expired in 15-minutes of use. But it did not.

We concluded that a *really* smart cookie should be developed to recognize its environment and to "know" when it should not implant. A "browsing agent" should be developed to recognize and honor secure areas, particularly when it finds itself within a multi-user machine.

Conclusions

It was the positive discoveries, and the identified problems and perils such as those discussed above, that have led our generally theory-directed thinking into considering the issues that arise in the very practical EC domain. As a user we may be pleased or frustrated by EC systems. As an AI learning theorist, we should be able to design and improve them.

References

Anonymous, 1998. Private communications between a prefers-to-remain-anonymous electronic-trading brokerage-house technical support staffer and the author, September 28, 1998–November 4, 1998.

Bienkowski, M., 1998. A Reader's Guide to Agent Literacy. *SIGART Bulletin*, Vol 9. No 2 (Fall 1998), 23–29.

Cherniavsky, J.C., 1987. Computer Systems as Scientific Theories: a Popperian Approach to Testing. In *Proceedings of the Fifth Pacific Northwest Software Quality Conference*, Portland OR, October 1987, 297–308.

Fass, L. F., 1992. Software Design as a Problem in Learning Theory. In Notes of the *AAAI-92 Workshop on Automating Software Design*, San Jose CA, July 1992, 48–49. Notes also released as Technical Report FIA-92-18, NASA Ames Research Center AI Branch.

Fass, L. F., 1994. Modeling Perfect Behavior: a Machine-Theoretic Approach. In *Proceedings on Summary of the Joint Conference on Information Science: Computer Theory and Informatics*, Duke University/Pinehurst NC, November 1994, 141–144.

Fass, L., F., 1999. Establishing Software "Correctness" by Logical and Algebraic Means. In Notes of the *AAAI-99 Workshop on Intelligent Software Engineering*, Orlando FL, July 1999. Forthcoming.

Jensen, D., and Goldberg, H., eds. 1998. *Papers from the AAAI Fall Symposium on Artificial Intelligence and Link Analysis*, Orlando FL, October 1998, AAAI Report FS98-01.

Kiernan, V., 1998. Use of "Cookies" in Research Sparks a Debate over Privacy. *The Chronicle of Higher Education*, September 25, 1998, A31–A32.

Special Issues

AI Magazine, 1997. Special Issue on Intelligent Systems on the Internet, Vol 18, No 2 (Summer 1997).

AI Magazine, 1998. Special Issue on Agents, Vol 19, No 2 (Summer 1998).

Communications of the ACM, 1994a. Special Issue on Intelligent Agents, Vol 37, No 7 (July 1994).

Communications of the ACM, 1994b. Special Issue on Internet Technology, Vol 37, No 8 (August 1994).

IEEE Expert, 1995. Special Issue on Intelligent Internet Services, Vol 10, No 4 (August 1995).

Jennings, N., Sycara, K, and Georgeff, M., eds, 1998. *Autonomous Agents and Multi-Agent Systems*, Premier Issue, Vol 1, No 1. Boston MA: 1998 Kluwer Academic Publishers.

Leona F. Fass received a B.S. in Mathematics and Science Education from Cornell University and an M.S.E. and Ph.D. in Computer and Information Science from the University of Pennsylvania. Prior to obtaining her Ph.D. she held research, administrative and/or teaching positions at Penn and Temple University. Since then she has been on the faculties of the University of California, Georgetown University and the Naval Postgraduate School. Her research primarily has focused on language structure and processing; knowledge acquisition; and the general interactions of logic, language and computation. She has had particular interest in inductive inference processes, and applications/adaptations of inference results to the practical domain. She may be reached at

Mailing address: P.O. Box 2914
Carmel CA 93921