Learning From "Relevant" and "Irrelevant" Information

(A Research Overview)

Leona F. Fass

Much of our research has focused on the formalization of inductive inference processes and on the mathematical and philosophical foundations of inductive learning. We began such work some years ago in connection with a specific problem of (formal language) learning and learnability. There we sought to develop a learning technique applicable to any member of a particular (linguistic) knowledge class. The class included both finite and infinite elements: successful learning in either case required characterizing the knowledge by finite means. The solution to that specific learning problem employed techniques of inductive inference to discover learnable models of (possibly infinite) knowledge from suitable, finite knowledge samples. The key to establishing learnability was establishing the ***existence*** of a finite information sample from which a model for an entire body of knowledge could be found. When we determined a suitable sample that would lead a learning system to discover a correct result, we dealt with the issue of ***relevance***.

From our perspective on learning and learnability we define ***relevant*** information as that which, once sufficiently available to a system, necessarily leads that system to achieve a specific goal or obtain a desired result. By contrast, ***irrelevant*** information is that which, if available to a system, neither effects nor facilitates its achievement of a learning goal.

For instance, learners of arithmetic acquire a model for multiplying ***any*** pair of integers, given a finite set of examples. In this "real life" case, examples conveyed by traditional multiplication tables, and some more sophisticated multi-digit examples that illustrate "add-and-carry" rules, would be considered ***relevant***. Most would consider subtraction examples to be ***irrelevant***, relative to a learn-how-to-multiply goal. In any event, "real life" teachers of successful students, and designers of successful expert or computational learning systems, would determine sufficient examples and practice problems, necessarily leading the learner to achieve the learn-how-to-multiply, or other, specified goal.

Our own views of necessary, sufficient and relevant information have evolved as we have refined and extended our original learnability results. When we first sought to establish the existence of learnable models of knowledge, we developed a paradigm for ***discovering*** "correctness" from samples of how a model should and should ***not*** behave. We then showed (in our specific problem domain) that a "bad" behavioral sample was "irrelevant": we could determine a correctly behaving model by ***constructive*** induction, from a positive information sample alone. Then, as we sought to extend our original work, we found the "irrelevant" negative information we'd eliminated from consideration became relevant from a different perspective. A specific sample of negative information ("bad" behavior) proved critical, as we determined an alternate method of discovering correctness by "adversarial" or "default" means. Seeking to generalize our original results (in formal language learning) to other problem domains, we developed some classification of the instances in which our theory and techniques for discovering and sampling "correctness" might be expected to apply. We now briefly describe our results in this area.

We originally solved the inductive inference problem for context-free languages, determining inferable syntactic models for *any* language in the context-free class. To do so we established that, from a ***suitably represented*** language ***sample***, a characterizing recognitive device for the entire language, could be inductively constructed. We then showed that a corresponding generative grammar could be inferred by similar means [7, 8]. Structural properties of the languages themselves (e.g., central recursion) made representation a critical factor in determining our results. Based on a suggestion of Levy and Joshi [2] we represented the language in a structural, skeletal fashion and, as recognitive processors, considered the class of skeletal automata that [2] first described. We were able to show (generalizing the Myhill-Nerode Theorem [4, 5, 8]) that each state of such a processor corresponded to a structural congruence class (reflecting the recursion). Elements of an inferable model were ***automatically*** determined from representatives of the congruence classes. Representatives of a selected set of congruence classes, corresponding to positive language constructs, proved to be all that was ***relevant*** to determining a recognitive result. A suitable finite sample and finite distinguishing experiments [3, 6] were sufficient for ***automatically*** discovering ***every*** congruence class and ***correct***, minimal (space-efficient) inferred results.

Adapting the process and theory, we showed ***incorrectness*** (complementary, relative to a containing domain) may be similarly, finitely defined [9]. With sufficient tests using samples of "incorrectness", a potential minimal model can be determined *not* to be incorrect. Thus it is ***automatically*** verified as correct by experimental, "default" means. The "bad behavior", irrelevant to inferring a model, becomes relevant in a test set (it corresponds to classes of bad linguistic constructs that would lead a recognitive device to a reject or "dump" state). If none of the selected "bad" behavior in the test set is exhibited by the potential model, we know it *cannot* be incorrect. It is "default verified" [9,11,12].

We have found this inference/testing paradigm, using positive/negative, relevant/irrelevant information is successful in modeling finitely-realizable behaviors with decidable membership queries. (Then we can finitely characterize the behavior and its complement, and distinguish members or representatives ***of*** the relevant class from those that are ***not***.) Thus learning or reasoning about many language processors, finite-state devices, and "minimal" compilers are important areas to which our theory can be successfully applied.

We have looked into problems of learning or reasoning about programs and software in general, and find our theory yields results no worse than those of other theoreticians. We investigated these problems at the suggestion of Cherniavsky [1], and have found the inference/testing, relevant/irrelevant information sampling paradigm to yield a reasonable approach to determining approximately correct (minimal) software. The finite experiments to ***automatically*** construct by induction (or default verify) ***approximations*** to correct software, produce results that improve upon actual software design processes often in use today. Verification by proof is not necessary. "Correctness" is assured from the sample information, using inductive constructive or "default" means [9-11].

In the case of learning behaviors, or bodies of knowledge, that cannot be described in terms so simple as might be exhibited by a program or a "finite-state device", we are only just beginning to look into applications of our theory. Our approach to relevant sampling, and approximations to models of general bodies of knowledge, is described in [11]. The critical factor in such problems may well be distinguishing what is relevant from what is not.

A SELECTION OF REFERENCES

[1] Cherniavsky, J. C., "Computer Systems as Scientific Theories: A Popperian Approach To Testing", *Proc. of the Fifth Pacific Northwest Software Quality Conference*, Portland (October, 1987), pp. 297-308.

[2] Levy, L. S., and A. K. Joshi, "Skeletal Structural Descriptions", *Inf. Contr.*, Vol. 39 (1978), pp. 192-211.

[3] Moore, E. F., "Gedanken-Experiments on Sequential Machines", in *Automata Studies*, C. E. Shannon and J. McCarthy, Eds., Princeton University Press, 1956, pp. 129-153.

[4] Myhill, J., "Finite Automata and the Representation of Events", *WADC Tech. Rept.*, 57-624, Wright-Patterson AFB, 1957.

[5] Nerode, A., "Linear Automaton Transformations", *Proc. AMS*, IX (1959), pp. 541-544.

[6] Thatcher, J. W., and J. B. Wright, "Generalized Finite Automata Theory with an Application to a Decision Problem of Second Order Logic", *Math. Sys. Th.*, Vol. 2 (1970), pp. 57-81.

Selected Pertinent Papers By The Author

[7] Fass, L. F., "Learnability of CFLs: Inferring Syntactic Models from Constituent Structure", *1987 Linguistic Institute, Meeting on the Theoretical Interactions of Linguistics and Logic*, Stanford, July 1987. Abstract, *J. Symbolic Logic*, Vol. 53, No. 4 (December, 1988), pp. 1277-1278. Research Note, *SIGART Special Issue on Knowledge Acquisition* (April 1989), pp. 175-176.

[8] Fass, L. F., "A Minimal Deterministic Acceptor for Any (Structured) Context-Free Language", preliminary version (1987). Extended version, *1990-91 Annual Meeting of the Linguistic Society of America*, Chicago (January, 1991); 35 pp, abstracted in *Meeting Handbook*, p. 17.

[9] Fass, L. F., "A Common Basis for Inductive Inference and Testing", *Proc. of the Seventh Pacific Northwest Software Quality Conference*, Portland (September, 1989), pp. 183-200.

[10] Fass, L. F., "Software Design as a Problem in Learning Theory (A Research Overview)", Notes of *AAAI-92, Workshop on Automating Software Design*, San Jose (July, 1992), pp. 48-49.

[11] Fass, L. F., "Inference and Testing: When 'Prior Knowledge' is Essential to Learning", in Notes of *AAAI-92 Workshop on Constraining Learning Through Prior Knowledge*, San Jose, CA (July, 1992), pp. 88-92.

[12] Fass, L.F., "On the Verification of a Grammar as 'Correct' ", presented at *Ninth Annual Conference on Computing and Philosophy*, Occidental College, August 1994.

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. Dr. Fass may be reached at

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