AI Magazine Volume 16 Number 2 (1995) (© AAAI)

Behavioral Cloning:

A Correction

Rui Camacho and Donald Michie

We recently reported on the application of a machine-learning (ML) technique to automated flight control using a simulated F-16 combat plane (Michie and Camacho 1994). Subsequent tests of our data-induced flying model have broadly confirmed the reported results but have also identified a lack of robustness. We had underestimated the latter and now regard our report (Michie and Camacho 1994) as being, by omission, potentially misleading.

Successes and shortcomings of *behavioral cloning* (learning by imitation) have been reviewed by Urbancic and Bratko (1994). They discuss the following problem domains: pole-and-cart balancing (Michie, Bain, and Michie 1990; Chambers and Michie 1969), flight-simulator control (Sammut et al. 1992), telephone-line scheduling (Kibira 1993), and crane-simulator control (Urbancic and Bratko 1994). Conclusions are as follows:

First, successful clones have been induced using standard ML techniques in all four domains.

Second, the cleanup effect, whereby the clone surpasses its original, has been observed in all four domains.

Third, in all domains, the best clones were obtained when examples from a single human only were used.

Fourth, the present approach lacks robustness in that there is no guarantee of inducing with high probability a successful clone from given data.

Fifth, typically, the induced clones are not sufficiently robust with respect to changes in the control task.

Sixth, although the clones do provide some insight into the control strategy, they in general lack conceptual structure that would clearly capture the causal relations in the domain and the goal structure of the control strategy.

Pearson et al. (1993) worked with flight-simulator control (Cessna, on a Silicon Graphics workstation) at the University of Michigan. They constructed skill models using hand crafting rather than ML derivation. Robustness to varied starting conditions, although not systematically studied, appeared reasonable—attributable, we believe, to the goal directedness of their underlying architecture. Our own more systematic robustness trials of the earlier-described F-16 flying models, with respect to local changes in the control task, gave disappointing results. In one series, only 3 of the 81 starts saw successful completion of the entire mission, most of the remainder crashing either during landing or while lining up for descent, this in spite of minor factors of noncomparability between training and test conditions more likely to affect performance positively than negatively.

In light of these results, we reviewed the inductively generated decision trees that comprised the procedure bodies of the 72 individual agents flying the simulated plane (see the earlier-cited article for the full design architecture) and identified two major factors in the observed brittleness. In many cases, learned decision trees had failed to incorporate an attribute test that, from a perspective of goals and causality, was plainly mandatory. For example, the chances of achieving stage 6's postcondition (aircraft is lined up with respect to the runway) are clearly compromised if relevant agents remain blind to the craft's current bearing with respect to the runway! The other weakness that stood out was the cursory nature of the definitions provided for signaling accomplishment of the goals of individual stages ("achievement goals" in the terminology of Laird and coworkers). There is an evident connection between the two because it is precisely failure to relate sensed state variables explicitly to local goals that is lacking from pure learning by imitation. As a next step, one of us (R.C.) is refining the transition goals that signal completion of successive flight stages. The main tool for this purpose is a variant of inductive logic programming that supports the inductive learning of complex flight goals.

References

Chambers, R. A., and Michie, D. 1969. Man-Machine Co-Operation on a Learning Task. In *Computer Graphics—Techniques and Applications*, eds. R. D. Parslow, R. W. Prowse, and R. E. Green, 179–185. New York: Plenum.

Kibira, D. 1993. Developing an Expert Controller of a Black Box Simulation of a Telephone Line Using Machine Induction, AI Laboratory, University of New South Wales.

Michie, D.; Bain, M.; and Michie, J. E. 1990. Cognitive Models from Subcognitive Skills. In *Knowledge-Based Systems in Industrial Control*, eds. M. Grimble, J. McGhee, and P. Mowforth. London: Peregrinus.

Michie, D., and Camacho, R. 1994. Building Symbolic Representations of Intuitive Real-Time Skills from Performance Data. In *Machine Intelligence 13*, eds. K. Furukawa, D. Michie, and S. Muggleton, 385–418. Oxford, U.K.: Clarendon.

Pearson, D. J.; Huffman, S. B.; Willis, M. B.; Laird, J. E.; and Jones, R. M. 1993. A Symbolic Solution to Intelligent Real-Time Control. In *Robotics and Autonomous Systems 11*, 279–291. New York: Elsevier.

Sammut, C.; Hurst, S.; Kedzier, D.; and Michie, D. 1992. In Proceedings of the Ninth International Workshop on Machine Learning, 385–393. San Francisco, Calif.: Morgan Kaufmann.

Urbancic, T., and Bratko, I. 1994. Reconstructing Human Skill with Machine Learning. Presented at the 1994 European Conference on Artificial Intelligence (ECAI 94), 8–12 August, Amsterdam, The Netherlands.