The Knowledge Engineering Review

More comments on rule induction: A reply to Bratko and Michie

BRIAN P. BLOOMFIELD
Systems Group, The Open University, Milton Keynes, MK7 6AA

In their paper "Some Comments on Rule Induction" Bratko and Michie address only three of the seven concerns that I put forward in my original article. Significantly, they have chosen to ignore the deeper and more important philosophical issues which I raised, while seeking to deflect my concerns by proclaiming their own privileged position of "being in the know" as experienced rule inductionists. This is a strategy well-known to philosophers and sociologists of science, unfortunately, it tends to stifle debate and avoids the serious consideration of constructive criticism.[1] It is therefore important to respond to the comments made by Bratko and Michie and to re-focus attention on the central themes of my argument.

(1) In their introduction Bratko and Michie claim that I imply that the hand crafting of rules "somehow ensures more dependable or otherwise more satisfactory results" than does the machine induction of rules. However, I make no such implication. Indeed, elsewhere I have argued the poverty of conventional rule-based systems which are constructed solely on the basis of articulated knowledge while neglecting the visual/ostensive dimension of human expertise and the fact that knowledge is not an "object"—a "thing-in-itself" to be "mined" from the minds of experts—but is constitutively social in nature.[2] Throughout my paper I make reference to the difficulties inherent in interview/dialogue methods of knowledge elicitation (e.g. the problem that experts often cannot account for their own expertise) in order to show how in general rule induction does not circumvent them and in fact faces more severe problems.

(2) They claim that it is an "empirical finding" that an expert can more easily provide an adequate training set than a full explication of their expertise. This rather reminds me of the old adage that one man's terrorist is another man's freedom fighter; in this case the point is that one man's "empirical finding" is another's theoretical supposition. Now, this is not to deny the fact that rule induction might well work better with certain individual experts, or in certain particular domains; the crucial question here (which in fact was one of the major thrusts of my original piece) is *how do we go about the business of deciding appropriate domains for rule induction or indeed any other form of knowledge elicitation.*

(3) They argue that inductive tools help the expert to identify where attributes are missing in a knowledge base. In the case of Expert-Ease, missing attributes become evident, for instance, when two examples in a training set have the same respective values for existing attributes but are assigned different classifications—thus giving rise to a contradiction. The specification of a missing attribute helps to differentiate between the two conflicting examples because each will take on a different specific value in respect of that attribute. However, though this might "patch" the knowledge base we cannot assume *a priori* that the particular attribute that is added has any real epistemological weight. To know this requires the expert to reflect on his or her domain knowledge and provide reasons for the relevance of an attribute: reasons which go beyond the simple fact that the inclusion of the attribute allows a piecemeal or *ad hoc* patching of the knowledge base. Yet, such in-depth reasoning is the sort of time consuming activity that rule induction was originally meant to circumvent!

(4) In reply to my discussion of the decision trees induced by Expert-Ease, Bratko and Michie claim

that I simply repeat the familiar fact that a "rule cannot exceed the power of the description language provided". Further, they point out that this limitation applies equally to hand-crafted rules. In making this claim they ignore my central point—which is to illuminate the epistemological status of machine induced rules—and therefore the fact that while all rules are indeed limited by the language in which they are stated, some description languages may be more powerful than others. Moreover, there is another point here which stems from a consideration of what is involved in following or applying a rule. Wittgenstein showed us that on pain of infinite regress rules themselves cannot specify how they are to be applied; rather, rules are conventional in character, they are part of a shared culture. In the case of hand-crafted rules, the conventions involved are part of the background, unstated, or tacit knowledge in the given domain. With in-depth interviewing (though limited to knowledge that can be articulated) a knowledge engineer can begin to probe this background or tacit knowledge of the expert; in constructing rules he or she can seek to force the expert to justify every scrap of knowledge that is either volunteered or unearthed. (Certain elements of tacit knowledge may even become articulated knowledge and thereby transformed into facts or rules.)[3] But, in the case of rule induction the knowledge engineer constructs a decision tree or set of rules on the basis of knowledge which is restricted to mere examples of expertise. In short then, the description language for handcrafted rules can be more powerful than for induced rules (e.g. it can incorporate causal knowledge). Moreover, because of the greater social interaction between expert and knowledge engineer—where the latter can act as an apprentice, aiming to absorb some of the background knowledge and culture of the former—it rests on firmer epistemological ground. It would be interesting to know whether in building a hypothetical expert system for knowledge engineering, Bratko and Michie would employ rule induction as a means of transferring their own expertise, or whether they too would resort to hand-crafted rules.

(5) Bratko and Michie assert that I have omitted the criterion which matters most—namely, "experience with applying induction to practical problems". All I can say to this remark is that their point will be of little comfort to anyone not "in the know" who has bought Expert-Ease in the hope of avoiding difficult knowledge elicitation problems. The criteria I put forward arguably contribute toward the use of rule induction tools such as Expert-Ease, and at the very least go beyond the sort of simple admonitions found in the user manual of Expert-Ease—which, for example, states that the program should not be used in the case of problems where decisions are made at random.

(6) Finally, in their concluding remarks they assert that in many cases rule induction has achieved better results than dialogue techniques. This may or may not be so, but as it stands the assertion must be treated as a context-dependent interpretation rather than a concrete fact—it certainly does not merit the status of a general rule for knowledge engineering. Apart from my own experience with Expert-Ease I know of at least one instance where the need to supply domain examples proved "traumatic".[4] In any case, what is important to the future development of knowledge engineering is that we have a theory or explanation of why different elicitation techniques work (or not) in a range of domains. To move in that direction it is necessary to tackle the problems inherent in understanding the nature of knowledge and expertise—not to avoid them.

Notes

1 For examples from the history of mathematics see: Lakatos, I. [1976] *Proofs and Refutations* (Cambridge: Cambridge University Press). For an example from simulation modelling and systems analysis see: Bloomfield, B.P. (1986) *Modelling the World: The Social Constructions of Systems Analysts* (Oxford: Basil Blackwell), Chapter 6.

2 Bloomfield, B.P. (1986) "Epistemology for Knowledge Engineers", *Communication & Cognition—Artificial Intelligence*, vol. 3(4), 305–320.

3 Collins, H.M. et al. (1985) "Where's the expertise: expert systems as a medium of knowledge transfer", in M.J. Merry (Ed.) *Expert Systems 85* (Cambridge: Cambridge University Press), 325–334.

4 Roycroft, A.J., personal communication. Roycroft, a domain expert in chess endgames, underwent an 18 month secondment at the Turing Institute.