THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

A Science of Reasoning (Extended Abstract)

Citation for published version:

Bundy, A 1998, A Science of Reasoning (Extended Abstract). in H de Swart (ed.), *Automated Reasoning with Analytic Tableaux and Related Methods*. vol. 1397, Springer International Publishing, pp. 10-17.

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Published In:

Automated Reasoning with Analytic Tableaux and Related Methods

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

A Science of Reasoning: Extended Abstract *

Alan Bundy

Department of Artificial Intelligence
University of Edinburgh
Edinburgh, EH1 1HN, Scotland
Email: bundy@edinburgh.ac.uk, Tel: 44-31-225-7774

Abstract

How can we understand reasoning in general and mathematical proofs in particular? It is argued that a high-level understanding of proofs is needed to complement the low-level understanding provided by Logic. A role for computation is proposed to provide this high-level understanding, namely by the association of *proof plans* with proofs. Criteria are given for assessing the association of a proof plan with a proof.

1 Motivation: the understanding of mathematical proofs

We argue that Logic¹ is not enough to understand reasoning. It provides only a low-level, step by step understanding, whereas a high-level, strategic understanding is also required. Many commonly observed phenomena of reasoning cannot be explained without such a high-level understanding. Furthermore, automatic reasoning is impractical without a high-level understanding.

We propose a science of reasoning which provides both a low- and a high-level understanding of reasoning. It combines Logic with the concept of *proof plans*, [Bundy 88].

*The research reported in this paper was supported by SERC grant GR/E/44598 and an SERC Senior Fellowship to the author. I would like to thank other members of the mathematical reasoning group at Edinburgh for feedback, especially Frank van Harmelen, Colin Phillips, Mitch Harris and Toby Walsh. I am grateful for comments on the first draft of this paper from two anonymous referees. The full version of this extended abstract will appear in "Computational Logic: Essays in Honor of Alan Robinson", MIT Press, forthcoming.

¹We adopt the convention of using uncapitalised 'logic' for the various mathematical theories and capitalised 'Logic' for the discipline in which these logics are studied.

We illustrate this with examples from mathematical reasoning, but it is intended that the science should eventually apply to all kinds of reasoning.

2 The Need for Higher-Level Explanations

A proof in a logic is a sequence of formulae where each formula in the sequence is either an axiom or is derived from earlier formulae in the sequence by a rule of inference. Each mathematical theory defines what it means to be a formula, an axiom or a rule of inference. Thus Logic provides a low-level explanation of a mathematical proof. It explains the proof as a sequence of steps and shows how each step follows from previous ones by a set of rules. Its concerns are limited to the soundness of the proof, and to the truth of proposed conjectures in models of logical theories.

While Logic provides an explanation of how the steps of a proof fit together, it is inadequate to explain many common observations about mathematical proofs.

- Mathematicians distinguish between understanding each step of a proof and understanding the whole proof.
- Mathematicians recognise families of proofs which contain common structure.
- Mathematicians use their experience of previously encountered proofs to help them discover new proofs.
- Mathematicians distinguish between ‘interesting’ and ‘standard’ steps of a proof.
- Mathematicians often have an intuition that a conjecture is true, but this intuition is fallible.
- Students of mathematics, presented with the same proofs, learn from them with varying degrees of success.

3 Common Structure in Proofs

Several researchers in automatic theorem proving have identified common structure in families of proofs. For instance,

- [Bundy & Welham 81] describes the common structure in solutions to symbolic equations. This common structure was implemented in a process of *meta-level inference* which guided the search for solutions to equations.
- [Bundy *et al* 88] describes the common structure in inductive theorems about natural numbers, lists, *etc.* This common structure was implemented as an *inductive proof plan* which was used to guide the search for proofs of such theorems.

- [Bledsoe *et al* 72] describes the common structure in theorems about limits of functions in analysis. This common structure was implemented as the *limit heuristic* and used to guide the search for proofs of such theorems.
- [Wos & McCune 88] describes the common structure in attempts to find fixed-points combinators. This common structure was implemented as the *kernel method* and used to guide the search for such fixed-points.
- [Polya 65] describes the common structure in ruler and compass constructions. This common structure was implemented by [Funt 73] and used to guide the search for such constructions.

4 Proof Plans

Common structure in proofs can be captured in *proof plans*. A proof plan consists of two parts: a *tactic* and a *method*. A tactic is a procedure which constructs part of a proof by applying a sequence of rules of inference. High-level tactics are defined in terms of lower-level sub-tactics. The lowest level tactics will apply individual rules of inference. A method is a partial specification of a tactic. It consists of preconditions which must be satisfied before the tactic is executed and some effects which will be true provided the tactic application is successful.

Proof plans have been implemented within the OYSTER-CIAM system, [Bundy *et al* 88]. OYSTER is a theorem prover for Intuitionist Type Theory. CIAM is a plan formation program which has access to a number of general-purpose tactics and methods for inductive proofs. This system has been used to control the search for inductive proofs about natural numbers and lists. CIAM constructs a special-purpose proof plan for each conjecture out of its methods and tactics. The tactic of the proof plan is then executed. It instructs OYSTER to build a proof of the conjecture. The search for a proof plan at the meta-level is considerably cheaper than the search for a proof at the object-level. This makes proof plans a practical solution to the problems of search control in automatic theorem proving.

5 The High-Level Understanding of Proofs

Thus a high-level explanation of a proof of a conjecture is obtained by associating a proof plan with it. The tactic of this proof plan must construct the proof. The method of this proof plan must describe both the preconditions which made this tactic appropriate for proving this conjecture and the effects of this tactic's application on the conjecture. It must also describe the role of each sub-tactic in achieving the preconditions of later sub-tactics and the final effects of the whole tactic.

In fact, this association provides a multi-level explanation. The proof plan associated with the whole proof provides the top-level explanation. The immediate sub-tactics and

sub-methods of this proof plan provide a medium-level explanation of the major sub-proofs. The tactics and methods associated with individual rules of inference provide a bottom-level explanation, which is similar to that already provided by Logic.

The general-purpose tactics and methods which we will use to build proof plans, and the association of proof plans with proofs will constitute the theories of our science of reasoning. This extends the way in which logical theories and the association of logical proofs with real proofs and arguments, constitute the theories of Logic (especially Philosophical Logic). Just as Logic also has meta-theories about the properties of and relations between logical theories, we may also be able to develop such meta-theories about proof plans.

6 What is the Nature of our Science of Reasoning?

Before we can dignify this proposed study of the structure of proofs with the epithet *science* we must address a fundamental problem about the nature of such a science. Traditional sciences like Physics and Chemistry study physical objects and the way they interact. The subject of our proposed science is proof plans. But proof plans are not physical objects. If they can be said to exist at all it is in the minds of mathematicians proving theorems, teachers explaining proofs and students understanding them. Physicists assume that the electrons in the apple I am eating as I write are essentially the same as the electrons in some distant star. But proof plans will differ from mind to mind and from time to time. There will be billions of such proof plans. Are we doomed merely to catalogue them all? Given the difficulty of discovering the nature of even one such proof plan, what a difficult and ultimately pointless task this would be. We would prefer to narrow our focus on a few representative proof plans. But on what basis could these few be chosen?

Fortunately, this is not a new problem. It is one faced by all human sciences to some extent and it is one that has been solved before. Consider the science of Linguistics. In Linguistics the theories are grammars and the association of grammatical structure with utterances. Linguists do not try to form different grammars for each person, but try to form a grammar for each language, capturing the commonality between different users of that language. They try to make these grammars as parsimonious as possible, so that they capture the maximum amount of generality within and between languages. Linguists do not claim that everyone or anyone has these target grammars stored in their head — nor, indeed, that anyone has a grammar at all — only that they *specify* the grammatical sentences of the language.

Another example is Logic itself. Again judged by the arguments people produce, the logical laws differ between minds and vary over time. Logicians do not try to capture this variety, but confine themselves to a few logics which specify ‘correct’ arguments. As with grammatical sentences, correct arguments are identified by initial observation of arguments actually used and consultation with experts to decide which of these are correct.

I place our proposed science of reasoning between Linguistics and Logic. Proof plans are more universal than grammatical rules, but it is possible to associate different, equally appropriate proof plans with the same proof. The study of proof plans appeals both to an empirical study of the way in which mathematicians structure their proofs and to reflection on the use of logical laws to put together proofs out of parts.

Thus there are strong precedents for a science that takes mental objects as its domain of study and tames the wide diversity of exemplars by imposing a normative explanation informed by reflection and empirical study. It only remains to propose criteria for associating proof plans with proofs that will enable us to prefer one proof plan to another. This we can do by appealing to general scientific principles. Our proposals are given in the next section.

7 Criteria for Assessing Proof Plans

If there were no criteria for the association of proof plans with proofs, then we could carry out our programme by associating with each proof an *ad hoc* tactic consisting of the concatenation of the rules of inference required to reproduce it, and constructing an *ad hoc* method in a similar way. This would not go beyond the existing logical explanation.

The only assessment criterion we have proposed so far is *correctness*, i.e. that the tactic of the proof plan associated with a proof will construct that proof when executed. We now discuss some other possible criteria.

- *Intuitiveness*: the way in which the proof is structured by a proof plan accords with our intuitions about how we structure the proof.
- *Psychological Validity*: there is experimental evidence that all, most or some mathematicians producing or studying proofs also structured a proof in the way suggested by some proof plan.
- *Expectancy*: there must be a basis for predicting the successful outcome of a proof plan.
- *Generality*: a proof plan gets credit from the number of proofs or sub-proofs with which it is associated and for which it accounts.
- *Prescriptiveness*: a proof plan gets more credit the less search its tactic generates and the more it prescribes exactly what rules of inference to apply.
- *Simplicity*: a proof plan gets more credit for being succinctly stated.
- *Efficiency*: a proof plan gets more credit when its tactic is computationally efficient.

- *Parsimony*; the overall theory gets more credit the fewer general-purpose proof plans are required to account for some collection of proofs.

We might start designing proof plans using the criteria of intuitiveness and psychological validity as sources of inspiration, but then use the criteria of correctness, expectancy, generality, prescriptiveness, simplicity, efficiency and parsimony to revise them.

8 The Role of the Computer

So far we have not involved the computer in this methodological discussion. One might expect it to play a central role. In fact, computers have no role in the *theory*, but play an important *practical* role. *Computation* plays a central role in the theory, because the tactics are procedures and they are part of the theory of our science of reasoning. It is not, strictly speaking, necessary to implement these tactics on a computer, since they can be executed by hand. However, in practice, it is highly convenient. It makes the process of checking that the tactics meet the criteria of the §7 both more efficient and less error prone. Machine execution is convenient:

- for speeding up correctness testing, especially when the proof plans are long, or involve a lot of search, or when a large collection of conjectures is to be tested;
- to automate the gathering of statistics, *e.g.* on size of search space, execution time, *etc*;
- to ensure that a tactic has been accurately executed; and
- to demonstrate to other researchers that the checking has been done by a disinterested party.

In this way the computer can assist the rapid prototyping and checking of hypothesised proof plans. Furthermore, in its 'disinterested party' role, the computer acts as a sceptical colleague, providing a second opinion on the merits of hypothesised proof plans that can serve as a source of inspiration. Unexpected positive and negative results can cause one to revise ones current preconceptions.

9 The Relation to Automatic Theorem Proving

Although our science of reasoning might find application in the building of high performance, automatic theorem provers, the two activities are not co-extensive. They differ both in their motivation and their methodology.

I take the conventional motivation of automatic theorem proving to be the building of theorem provers which are empirically successful, without any necessity to understand why. The methodology is implied by this motivation. The theorem prover is applied to a random selection of theorems. Unsuccessful search spaces are studied in a shallow way and crude heuristics are added which will prune losing branches and prefer winning ones. This process is repeated until the law of diminishing returns makes further repetitions not worth pursuing. The result is fast progress in the short term, but eventual deadlock as different proofs pull the heuristics in different directions. This description is something of a caricature. No ATP researchers embody it in its pure form, but aspects of it can be found in the motivation and methodology of all of us, to a greater or lesser extent.

Automatic theorem provers based on proof plans make slower initial progress. Initial proof plans have poor generality, and so few theorems can be proved. The motivation of understanding proofs mitigates against crude, general heuristics with low prescriptiveness and no expectancy. The 'accidental' proof of a theorem is interpreted as a fault caused by low prescriptiveness, rather than a lucky break. However, there is no eventual deadlock to block the indefinite improvement of the theorem prover's performance. If two or more proof plans fit a theorem then either they represent legitimate alternatives both of which deserve attempting or they point to a lack of prescriptiveness in the preconditions which further proof analysis should correct.

Thus, we expect a science of reasoning will help us build better automatic theorem proving programs in the long term, although probably not in the short term.

10 Conclusion

In this paper we have proposed a methodology for reaching a multi-level understanding of mathematical proofs as part of a science of reasoning. The theories of this science consist of a collection of general-purpose proof plans, and the association of special-purpose proof plans with particular proofs. Each proof plan consists of a tactic and a method which partially specifies it. Special-purpose proof plans can be constructed by a process of plan formation which entails reasoning with the methods of the general-purpose proof plans. Ideas for new proof plans can be found by analysing mathematical proofs using our intuitions about their structure and, possibly, psychological experiments on third party mathematicians. Initial proof plans are then designed which capture this structure. These initial proof plans are then refined to improve their expectancy, generality, prescriptiveness, simplicity, efficiency and parsimony while retaining their correctness. Scientific judgement is used to find a balance between these sometimes opposing criteria. Computers can be used as a workhorse, as a disinterested party to check the criteria and as a source of inspiration.

The design of general-purpose proof plans and their association with particular proofs is an activity of scientific theory formation that can be judged by normal scientific criteria. It requires deep analysis of mathematical proofs, rigour in the design

of tactics and their methods, and judgement in the selection of those general-purpose proof plans with real staying power. Our science of reasoning is normative, empirical and reflective. In these respects it resembles other human sciences like Linguistics and Logic. Indeed it includes parts of Logic as a sub-science.

Personal Note

For many years I have regarded myself as a researcher in automatic theorem proving. However, by analysing the methodology I have pursued in practice, I now realise that my real motivation is the building of a science of reasoning in the form outlined above. Now that I have identified, explicitly, the science I have been implicitly engaged in for the last fifteen years, I intend to pursue it with renewed vigour. I invite you to join me.

References

- [Bledsoe *et al* 72] W.W. Bledsoe, R.S. Boyer, and W.H. Henneman. Computer proofs of limit theorems. *Artificial Intelligence*, 3:27–60, 1972.
- [Bundy & Welham 81] A. Bundy and B. Welham. Using meta-level inference for selective application of multiple rewrite rules in algebraic manipulation. *Artificial Intelligence*, 16(2):189–212, 1981. Also available from Edinburgh as Research Paper 121.
- [Bundy 88] A. Bundy. The use of explicit plans to guide inductive proofs. In *9th Conference on Automated Deduction*, pages 111–120, Springer-Verlag, 1988. Longer version available from Edinburgh as Research Paper No. 349.
- [Bundy *et al* 88] A. Bundy, F. van Harmelen, J. Hesketh, and A. Smaill. *Experiments with Proof Plans for Induction*. Research Paper 413, Dept. of Artificial Intelligence, Edinburgh, 1988. To appear in JAR.
- [Funt 73] B. V. Funt. *A procedural approach to constructions in Euclidean geometry*. Unpublished M.Sc. thesis, University of British Columbia, October 1973.
- [Polya 65] G. Polya. *Mathematical discovery*. John Wiley & Sons, Inc, 1965. Two volumes.
- [Wos & McCune 88] L. Wos and W. McCune. *Searching for fixed point combinators by using automated theorem proving: a preliminary report*. ANL-88-10, Argonne National Laboratory, 1988.