

PREFACE

Introduction and apologia

Four men, who had been blind from birth, wanted to know what an elephant was like; so they asked an elephant-driver for information. He led them to an elephant, and invited them to examine it; so one man felt the elephant's leg, another its trunk, another its tail and the fourth its ear. Then they attempted to describe the elephant to one another. The first man said 'The elephant is like a tree'. 'No,' said the second, 'the elephant is like a snake'. 'Nonsense!' said the third, 'the elephant is like a broom'. 'You are all wrong,' said the fourth, 'the elephant is like a fan'. And so they went on arguing amongst themselves, while the elephant stood watching them quietly.¹

The notion of topos resembles an elephant in that it is possible to come up with very different descriptions of what topos theory is about, depending on the direction from which you approach the subject. It was André Joyal, in his lecture at the 1981 Cambridge Summer Meeting on Category Theory, who first drew attention to the parallel, and I am grateful to him for allowing me to use it as the basis for the title of the present work. He listed seven different descriptions of 'what a topos is like':

- (i) 'A topos is a category of sheaves on a site'
- (ii) 'A topos is a category with finite limits and power-objects'
- (iii) 'A topos is (the embodiment of) an intuitionistic higher-order theory'
- (iv) 'A topos is (the extensional essence of) a first-order (infinitary) geometric theory'
- (v) 'A topos is a totally cocomplete object in the meta-2-category \mathcal{CAT} of cartesian (i.e., finitely complete) categories'
- (vi) 'A topos is a generalized space'
- (vii) 'A topos is a semantics for intuitionistic formal systems'

In the 20 years that have passed since that lecture, the category-theory community has added a few more descriptions to the list (and the theoretical computer scientists have contributed yet others); for example,

- (viii) 'A topos is a Morita equivalence class of continuous groupoids'
- (ix) 'A topos is the category of maps of a power allegory'

¹ The story of the blind men and the elephant exists in numerous versions in Indian folklore. This version is freely adapted from E. J. Robinson, *Tales and Poems of South India* (London, 1885).

- (x) ‘A topos is a category whose canonical indexing over itself is complete and well-powered’
- (xi) ‘A topos is the spatial manifestation of a Giraud frame’
- (xii) ‘A topos is a setting for synthetic differential geometry’
- (xiii) ‘A topos is a setting for synthetic domain theory’,

and so on. But the important thing about the elephant is that ‘however you approach it, it is still the same animal’; this book is an attempt to demonstrate that the same is true of topos theory.

A brief bibliography of topos theory

This book resembles an elephant in another respect: namely its gestation period. To explain (and apologize for) this, I have to go back still further into the history of the subject. I must also crave the reader’s indulgence for what may seem an excessively egocentric view of that history in the paragraphs which follow.

The original notion of topos, as a ‘generalized space’ suitable for supporting the exotic cohomology theories required in algebraic geometry, sprang from the fertile brain of Alexandre Grothendieck in the early 1960s, and was developed in his ‘Séminaire de Géométrie Algébrique du Bois-Marie’ particularly during the academic year 1963–64. The duplicated notes of that seminar circulated widely among algebraic geometers and category-theorists over the next decade, until Springer-Verlag did the world a service by publishing a revised and expanded edition in three volumes of *Lecture Notes in Mathematics* in 1972 [36]. But by then the subject had already been ‘reborn’ in its second incarnation, as an elementary theory having links with higher-order intuitionistic logic, through the collaboration of Bill Lawvere and Myles Tierney during 1969–70 (which, in turn, built upon Lawvere’s work on providing a categorical foundation for mathematics, which had been developing since the early 1960s).

As a graduate student in the early 1970s, I was fortunate to ‘get in on the ground floor’ of elementary topos theory, by attending Tierney’s summer school lectures [1167] at Varenna in 1971, and later Saunders Mac Lane’s lectures during his Cambridge sabbatical in autumn 1972. At that time the only general accounts of the subject in print were the lecture notes [649] by Anders Kock and Gavin Wraith, which were not all that widely available, and Peter Freyd’s survey article [371], whose view of the subject was idiosyncratic to say the least – though they were joined over the next two years by the notes of Tierney’s lectures, already cited, and by Wraith’s second account [1236].

More or less immediately after completing my Ph.D. in 1974, I embarked on the writing of what became the first account in book form of the elementary theory of toposes, published by Academic Press in 1977 [504]. With the brash self-confidence of youth, I attempted to provide between one pair of covers a full account of as many aspects of the elephant as I could muster: the result was

a book which, though universally acknowledged as ‘far too hard to read, and not for the faint-hearted’,² remained in demand for a long time – and continues to be widely cited even today, despite being nearly 25 years out of date.

For some time [504] was the only available book on elementary topos theory. But by the time of Joyal’s 1981 lecture it had been joined by a second, Robert Goldblatt’s *Topoi: the categorical analysis of logic* [411], and over the next decade or so a good many other books appeared which either were devoted to topos theory or contained a good deal of topos-theoretic material: *Teoria toposurilor* [997] by Alexandru Radu in 1981–2, *Toposes, triples and theories* [87] by Michael Barr and Charles Wells in 1985, *Introduction to higher order categorical logic* [682] by Joachim Lambek and Philip Scott in 1986, *Toposes and local set theories* [95] by John Bell in 1988, *Categories, allegories* [381] by Peter Freyd and Andre Scedrov in 1990, *Models for smooth infinitesimal analysis* [856] by Ieke Moerdijk and Gonzalo Reyes in 1991, *Lecture notes on topoi and quasitopoi* [1248] by Oswald Wyler, also in 1991, *Sheaves in geometry and logic* [751] by Saunders Mac Lane and Ieke Moerdijk in 1992, *Elementary categories, elementary toposes* [760] by Colin McLarty, also in 1992, *Relative category theory and geometric morphisms* [238] by Jonathan Chapman and Frederick Rowbottom, yet again in 1992, and finally the third volume of Francis Borceux’ remarkable *magnum opus Handbook of categorical algebra* [147] in 1994.

Each of these books contributed a new view of the elephant, illuminating features not visible from other angles, and thus each of them was valuable. But I hope I shall not offend the authors listed above by saying that none of them, with the possible exception of the last, really attempted to give a picture of the whole animal: indeed, in several cases the concentration on one part of its anatomy is self-evident from the title.

Meanwhile, my own book had gone out of print in 1987, despite a corrected reprint (some time after the original stock had run out) in 1985. There was some suggestion at that time that Academic Press might be willing to publish a revised second edition of the book, if I were willing to write it; but their procrastination over, and lack of enthusiasm for, the 1985 reprint did not fill me with enthusiasm for that course of action. Another reason for my reluctance was the extent to which the subject had grown since the mid-1970s: since the principal merit of my book, in comparison with the others by then available or forthcoming, was its comprehensiveness, there would be no point in replacing it by anything other than a comprehensive account of the subject as it stood by then, and such an account would inevitably be very much longer than the 1977 book.

² The quotation is from the comments of an anonymous referee consulted by OUP about the present book.

The three blind men

By the mid-1980s Ieke Moerdijk had also become concerned (despite his involvement as co-author in two of the books mentioned above) about the lack of any single account of topos theory which could properly claim to give a comprehensive picture of the subject. However, like me he felt daunted by the prospect of writing such a book single-handed. Sometime early in 1986, he mentioned to me his idea that what was needed was a 'handbook' type publication in which a small team of experts would contribute extended essays on the different aspects of the subject, under the overall supervision of a co-ordinating editor – he felt that I was the ideal person to take on the latter rôle. I was not, at first, greatly taken with this idea: as I saw it, the chief problem with the 'handbook' style was that, although the elephant has many different aspects, their interconnectedness is such that it is impossible fully to appreciate any one of them without making some reference to most of the others. So I foresaw that the job of a co-ordinating editor, for such a volume, would be more than usually nightmarish.

Nevertheless, discussions between Ieke and myself about the idea continued to take place on the odd occasions when we met, and they soon came to involve my former student Andrew Pitts as another possible author. The 'handbook' idea gradually metamorphosed into something more like a 'Bourbaki' volume or the *Compendium of continuous lattices* [404], in which each author would have individual responsibility for the first drafts of particular sections, but the final text would have to be agreed by all the authors (and the individual authorship of the different sections would not be publicly identified). At a meeting between the three of us in December 1987, we formally agreed to go ahead with a three-author work on that basis, and a rough table of contents – including the division of the book into six parts, although they did not correspond exactly to the six parts in the present table of contents – was worked out, together with the responsibilities of the three of us for drafting the more introductory sections. The working title *Sketches of an elephant* was also agreed at this meeting.

By the late summer of 1988 we had enough material in draft to approach publishers with a definite proposal, and we signed an agreement with Oxford University Press in January 1989. At the time (I blush to recall it now) we hoped that we would be able to deliver a complete typescript to the publishers by the end of September 1991. We also estimated that the overall length of the book might be around 1000 pages; in the event, that estimate has proved more accurate than the temporal one.

Over the next couple of years, a substantial quantity of draft material was written, with P. T. J. working mainly on Parts A and B, I. M. on Part C and A. M. P. on Part D. Ieke produced large quantities of draft material, principally on topics from Chapters C2 and C3, although virtually none of it has survived into the version that exists today; on the other hand, stylistic experts may just be able to detect the origins of the present Sections D1.1–D1.3 in a draft written by Andrew, even though it has been heavily rewritten by me. But it had

become clear by early 1991 that the project was a much larger one than we had realized: in particular, the problems of co-ordination between the three of us, and of reaching agreement on a text that all three could be happy with, seemed overwhelming. It was at least partly these problems that led Ieke and Andrew independently to decide to withdraw from the project in late 1992.

I should make it clear that, despite their not seeing the project through to completion, I remain enormously indebted to both Ieke and Andrew. If it had not been for their support in the early stages, I should never have dared to embark on the writing of a work on this scale, and their encouragement kept me going until the book had progressed far enough to acquire its own momentum. Also, their constructive criticisms of my early drafts of Part A caught many infelicities that I should probably not have noticed myself. And on a more mundane level, Andrew taught me to use \LaTeX ,³ without which the project would have been not just foolhardy but downright impossible.

One blind man

I thus found myself in early 1993 faced with three options: to abandon the project myself, to try to find other co-authors, or to accept the burden of writing the whole book myself. The first of these was unattractive because it would mean abandoning the large amount of material (including a complete draft of Part A) that I had already written, and the second did not appeal since I could not think of any potential co-author who would be as easy to work with, or as productive, as I knew Ieke and Andrew to be. So I decided to go ahead with the third option, although I knew that that would involve a ‘long hard slog’ over several years – I found it impossible to estimate how many.

With my other commitments to teaching, research and administration, progress was indeed slow: often, several months would pass without my finding time to open the files where the text of the book was stored. In addition, I often felt like Heracles wrestling with the Hydra: every attempt to draft a new section seemed to entail the rewriting of at least two existing ones, so that I appeared to come no nearer to the goal of a complete text. (There was also, at times, the fear that the subject was advancing faster than I could write it down.)

Nevertheless, progress did sometimes get made. One thing which helped to spur the drafting, or revision, of certain sections was the chance to lecture about the material in them, to audiences around the world, on occasions which included the following:

- In June 1994 I was invited to lecture on classifying toposes at Shaanxi Teachers’ University, Xi’an, China, through the good offices of Wang Guo-Jun

³ Whilst I am on that subject, I should also record my indebtedness to another of my former students, Paul Taylor, whose \TeX diagram-drawing macros are used throughout the book.

and my former student Zhao Dongsheng. This involved the polishing of material from Chapters A1, D1 and D3.

- In May 1995, Francis Borceux invited me to lecture on locales at a summer school on ‘Topological Applications of Category Theory’ in Alle-sur-Semois, Belgium: this involved material from Chapters C1 and C3.
- During my sabbatical visit (arranged by Max Kelly) to Sydney, Australia, during February and March 1996, I gave a number of talks in the Sydney Category Seminar covering material from Chapters B1 and B3.
- In July 1997, as an invited speaker at the ‘Workshop on Logic, Language, Information and Computation’ (commonly known as WoLLIC) in Fortaleza, Brazil, I had the opportunity (thanks to Ruy de Queiroz) to give a tutorial on categorical logic before the main conference: once again, this provided an occasion to polish the material in Chapters D1 and D2.

I mention these meetings in order publicly to record my thanks to the organizers named: though they were not aware of it, they helped to bring the project closer to completion.

By the beginning of the third millennium, however, it had become clear to me that what was absolutely essential, if the writing of Parts A–D was ever to be finished, was to find an extended period of time in which I could set all other work aside and concentrate full-time on the book. (Although I had at various times begun work on several sections from Part F, I had by then taken the decision that the only sensible route to completion was to concentrate on the first four parts, which were too strongly interconnected to be written separately, and to leave the last two – which, dealing as they do with applications, are less strongly tied to the others – for completion at a later date and publication as a separate volume.) By great good fortune, such an opportunity arose with the invitation to participate in the ‘Mathematical Logic Year’ at the Institut Mittag-Leffler in Djursholm, Sweden, for a two-month period in January–March 2001. My gratitude to the organizers of that programme (Dag Normann, Viggo Stoltenberg-Hansen and Jouko Väänänen), for providing me with this opportunity, is impossible to exaggerate; and I am also deeply grateful to Kjell-Ove Widman and the staff of the Institut Mittag-Leffler for providing such a near-perfect working environment.

The speed with which everything fell into place, once I had the opportunity for uninterrupted work on the book, greatly surprised me. Although the proportion of the total text which was written in Sweden is not so great, the number of sections which reached (more or less) their final form during my time there is very large: about half of those in Part B, almost all in Part C apart from Chapter C1, and most of those in Chapters D4 and D5. (In addition, a good deal of material was added to the supposedly complete draft of Part A.) Within a month of my return to Cambridge, I was able to produce a complete text of the first four parts, except for three sections in Part D. Copies of this draft were circulated to a number of colleagues (including Ieke Moerdijk and Andrew Pitts,

as well as Martin Hyland, Anders Kock and Gavin Wraith) for their comments and suggestions; the latter (which were, in many cases, of considerable value) were taken into account in the final revision of the text which took place in August–September 2001, after the writing of the last three sections had been completed.

At the time of writing these words, Parts E and F of the original plan remain to be written. In view of my experience so far, I am making no promises about how long it will take to write them; but I hope that interested readers will not have to wait as long for the third volume as they have done for the first two.

How to read this book

‘Begin at the beginning,’ the King said gravely, ‘and go on till you come to the end; then stop.’⁴

Whilst the King of Hearts’s advice to the White Rabbit is generally sound, it is not particularly helpful in relation to this book. At the December 1987 meeting to which I have referred, the three of us took a number of decisions about the overall structure and content of the book: despite the many changes of detail which have taken place since then, the general principles which we agreed have by and large survived unchanged. Since some understanding of these principles is important for anyone who seeks to obtain information from this book, I shall describe them here.

The first and most important is that *the book is not a textbook*. It is not addressed to those who are trying to learn about topos theory for the first time,⁵ but rather to those who already have some acquaintance with the subject and who wish to deepen their understanding, or to learn about aspects of it which they have not previously encountered. In keeping with this, there are no ‘exercises for the reader’ (which are often the lazy author’s way of leaving out the proofs of results that he can’t be bothered to write out in full): instead, all results needed in the book (with the exception of a few results from other areas of mathematics, for which references to standard works are given) are proved in full, and all examples are fully worked out.⁶ This is not to say that the book *cannot* be used as a course text: it would, for example, be possible to base an introductory course on the purely category-theoretic side of topos theory on Chapters A1, A2 and A4, although the instructor would have to be fairly selective about the specialized material which has been inserted in those chapters for future reference, and would also probably want to leaven the presentation with some material from Parts C and/or D.

⁴ Lewis Carroll, *Alice’s Adventures in Wonderland* (London, 1865).

⁵ For those who do seek an introduction to the subject, the one by Mac Lane and Moerdijk [751] remains (in the present writer’s opinion) the most recommendable.

⁶ It should be understood that this policy does not debar the author from leaving routine ‘diagram-chasing’ and other calculations to the reader.

Secondly, *the book is not intended to be read sequentially*. The whole point of the six-part structure is that the book is seeking to present several different approaches to topos theory, so where one starts reading will depend on the line of approach one wishes to follow. Most readers, however experienced, will probably wish to start by at least skimming through Chapter A1, in order to familiarize themselves with the book's conventions on categorical notation and terminology (about which I shall have more to say below); but thereafter a reader who already knows the basic category theory of toposes may well decide to jump to Chapter B1, C1 or D1, depending on the approach he wishes to take. However, almost any reader will have to be prepared to refer backwards and/or forwards a good deal. Although it would have been theoretically possible to write entirely independent accounts of the geometric and logical aspects of topos theory in Parts C and D (and that was more or less our original intention), in fact the two are so thoroughly interconnected that such a course would have involved an immense amount of duplication (and of course increased the overall length of the book).

There is some duplication, which is quite deliberate: for example, the treatment of natural number objects in Section A2.5 overlaps with that in Section D5.1, the construction of classifying toposes is treated in two different ways in Sections B4.2 and D3.1, and exponentiability of toposes is treated in Sections B4.3 and C4.4. And there are numerous smaller instances, where an individual result is proved in two different ways, using material from different areas of the book. But in general, if there is an optimal proof of a particular theorem, and it clearly belongs (say) in Part D, there does not seem to be any point in writing it out a second time in Part C if it happens to be needed there. So the reader who wishes to see the full picture must be prepared to skip around from one part to another, when necessary. (On the other hand, a reader who is willing to 'take on trust' the results quoted from other parts of the book should be able to follow any given chapter as a (more or less) connected narrative.)

I have at least done my best to make the cross-references easy to use. All numbered references (definitions, lemmas, theorems, examples, etc.) within a given section are numbered in a single sequence: thus 'B3.4.5' denotes the fifth numbered reference (in fact a corollary) in the fourth section of the third chapter of Part B. For references within a given Part, the letter is omitted: thus if the corollary just mentioned is cited elsewhere in Part B, it appears as '3.4.5'. If a lemma, proposition, theorem, corollary or scholium⁷ is subdivided into parts, they are normally labelled (i), (ii), (iii) and so on, whereas definitions, examples and remarks are subdivided (a), (b), (c), The symbol \square is used to denote the end of a proof; if it appears at the end of the statement of a result, it means

⁷ I should perhaps explain that I use the term 'scholium' (literally, a marginal note) to denote something which follows directly from the proof of a preceding result, as opposed to a corollary which follows directly from the statement of the preceding result.

either that no proof is given, or that the proof is contained in the discussion which precedes the statement.

At the December 1987 meeting we adopted the principle that, for every major result to be included in the book, we should seek to find the optimal proof,⁸ although that should not preclude us from giving alternative proofs if they helped to shed additional light on the result. (I hope it is unnecessary to add that I am conscious of having not always succeeded in this aim.) Partly for this reason, we decided that we would not attempt to give detailed attributions for all the results in the book, since the optimal proof is often not the original one, but has been arrived at through successive polishing of the latter by several hands. Nevertheless, I have attempted to compile a comprehensive bibliography of topos theory, and to give attributions where credit is clearly due to a single source. (I considered attempting to classify the bibliography entries by subject area, as was done in [404], for example; but the task seemed too large and ultimately of doubtful value.) And in most sections I have provided a list of references at the end of the section for those who wish to pursue the topic(s) covered in the section in greater detail.

We also recognized that the choice of terminology and notation, in a book of this nature, was of immense importance. It is impossible to be entirely consistent about notation, in a subject of this size: there simply aren't enough different letters and typefaces for one to avoid re-using some of them in mutually contradictory ways. But I have done my best to ensure local consistency of notation, within each chapter. As regards terminology, I have consciously attempted to use this book to disparage some of the sloppier pieces of inappropriate usage which are common among category-theorists; whether it will have any influence in this respect of course remains to be seen, but readers who are familiar with the existing terminology may well find some of my choices rather surprising at first. I have, however, listed alternatives to my chosen terminology in the index, where they appear in italics, and in most cases my reasons for discarding them will be found in the text at the points to which the index entries refer.

A comprehensive index is also of great importance in a book of this nature. I have done my best to provide one: in addition to listing alternative terminology in italics, as already mentioned, it also contains entries for many of the principal theorems, particularly those which occur in several different versions in different places, and I hope that this will assist the specialist in finding his way quickly to the information which he wants. One point which should be mentioned about the index is that, where an index entry begins with a mathematical symbol, it is listed in the alphabetical order as if the name of the symbol had been spelled out: thus, for example, ' σ -pretopos' will be found between 'sifted coverage' and 'signature'. In addition to the general index, there is also an index of standard notation, ordered alphabetically on the same principle.

⁸ That is, the one which the late Paul Erdős would have described as belonging to 'The Book'.

Conclusion

I have already mentioned my indebtedness to several colleagues who have played particular rôles in the story which led up to the writing of this book. But there are many other workers in topos theory to whom I am indebted, for their ideas which have contributed to my own understanding of topos theory, and which have in many cases found their way into the text which follows without specific acknowledgement. It would be quite impossible to thank them all individually, so I hope that they will accept these few words of collective thanks.

Finally, I must thank the staff of Oxford University Press, and in particular Elizabeth Johnston and Alison Jones: not only for their enormous patience in waiting for this text ten years beyond the date when we originally promised to deliver it, but also for their enthusiastic and helpful response when I was finally able to let them know that its delivery was within sight. It has been a pleasure to work with them, and I look forward to doing so again when Parts E and F are ready for publication.

Cambridge
September 2001


P. T. J.

Sketches of an Elephant A Topos Theory Compendium

VOLUME 1

PETER T. JOHNSTONE

*Reader in the Foundations of Mathematics
University of Cambridge*



CLARENDON PRESS • OXFORD

2002