

1 INTRODUCTION

1.1 The social life of objects

One of the questions that we are most frequently asked is "What is Category Theory good for?" or "Why should I use Category Theory?" The question usually indicates a genuine and healthy reaction to the proposal of a new piece of mathematics that one is invited to learn, similar to the reaction we have each time we are asked to change our eating habits and, say, replace butter with olive oil when cooking... How is it going to make us happier? When is the change justified?

The question also indicates that the nature and role of Category Theory, namely in Computing, is not completely clear to most people. The way we like to present Category Theory is as a toolbox similar to Set Theory: a kind of mathematical *lingua franca* in the sense that it can be used for formalising concepts that arise in our day-to-day activity. It constitutes, however, a richer toolbox in the sense that the instruments that it provides are more sophisticated, and thus makes it easier to model situations that are more complex and involve structured objects. On the other hand, because these instruments are more sophisticated than those of Set Theory, it is necessary to learn how to use them. Briefly, in Category Theory one can do as much as in Set Theory, in an easier way when it comes to formalising and relating different notions of "structure", but at the cost of learning a few more concepts and techniques.

So, I would like to reformulate the original question as: "Having been brought up to think about the world in terms of sets, why should I now change and use another frame of reference?" The purpose of this book is to convince you that this change of reference is worth making, i.e. that this book will have been worth reading and that its ideas are worth using in everyday engineering practice!

For me, the key factor that justifies this change is related to the fact that, whereas concepts in Set Theory are typically formalised extensionally, in the sense that a set is defined by its elements, Category Theory provides a more implicit way of characterising objects. It does so in terms of the relationships that each object exhibits to the other objects in the universe of discourse. The best summary of the essence of Category Theory that I know is from the logician Jean-Yves Girard³. For him, Category Theory characterises objects in terms of their "social lives".

In my opinion, this focus on "social" aspects of object lives is exactly the reason for the applicability of Category Theory to Computing in general, and Software Engi-

³ What stronger evidence would one need of the close relationship between Logic and Category Theory...

neering in particular. To realise why this is so, one just needs to think that current software development methods, namely object-oriented ones, typically model the universe as a *society* of interacting objects. Agent-oriented methods are based on the same societal metaphor. This focus on interaction is not accidental; it is an attempt at tackling the increasing complexity of modern software systems. In this context, complexity does not necessarily arise from the computational or algorithmic nature of systems, but results from the fact that their behaviour can only be explained as emerging from the interconnections that are established between their components. By promoting interactions as a focal point in the definition of the structure of a system, one also brings software to the realm of natural, physical and social systems, something that seems to be essential for the development of well-integrated systems. Category Theory is advocated as a good mathematical structure for this integration precisely because it focuses on relationships and interactions! The work of Goguen on General Systems Theory [51,52,63], and recent books like [95], show exactly that.

This is why many of the examples that will be used throughout the book address what has become known in Software Engineering as "Software Architectures" [49], i.e. precisely the study of the gross modularisation principles that should allow us to design systems as, possibly standard, structures of smaller components and interconnections between them. The focus that Category Theory puts on morphisms, as structure-preserving mappings, is paramount for Software Architectures because it is the morphisms that determine the nature of the interconnections that can be established between objects (system components). Hence, the choice of a particular category can be seen to reflect, in some sense, the choice of a particular "architectural style". Moreover, Category Theory provides techniques for manipulating and reasoning about system configurations represented as diagrams. As a consequence, it becomes possible to establish hierarchies of system complexity, allowing systems to be used as components of even more complex systems (i.e. to use diagrams as objects), and for inferring properties of systems from their configurations.

The ultimate conclusion that we would like the reader to draw is that high-school education could well evolve in a way that equips our future generations with tools that are more adequate for the kind of systems that they are likely to have to develop and interact with. We believe that the teaching of mathematics could progressively shift from the set-theoretical approach that made it "modern" some decades ago (not to say in the last century) to one that is centred on interactions. This is, of course, a big challenge, mainly because it is not enough for the mathematical theory to be there: the right way needs to be found for it to be transmitted. We believe that recent books such as [74] are putting us in a path towards meeting this challenge, and we would hope this book to make a further contribution. However, we will be satisfied if the rationale for such a shift can, somehow, emerge from what we will present.

1.2 Categories versus sets

Let us consider an example in order to make clear what the difference of approach is between Set Theory and Category Theory. For this purpose, there is nothing better than showing how certain mundane set-theoretic constructions are modelled in Category Theory.

Consider, for instance, the characterisation of the empty set. In Set Theory, the empty set is characterised precisely by the property of not having any elements. There are several ways in which we can say this using a formal notation. Here is one:

$$(\forall x)x \notin \emptyset$$

The important point here is that any characterisation requires the use of the membership relation \in , i.e. the characterisation is made with respect to the elements that belong to the set.

Consider now the characterisation of the empty set in Category Theory. As discussed above, such a characterisation involves the definition of a "social life" of sets, i.e. it has to be made relative to the way sets interact with one another. Hence, the obvious question to ask first is "what is the social life of sets"?

The first important thing to understand is that Category Theory does not provide an answer to questions like this one; it is up to whoever is formalising a particular domain of discourse to come up with a definition of "social life" that is convenient. Convenience here has to be measured against the formalisation activity that is being undertaken, possibly as an abstraction of real-world phenomena. Hence, it is not subject to mathematical proof. What Category Theory does is require some basic properties of such a "social life" so that the whole mathematical machinery that we are about to describe can be applied successfully. Such basic properties are defined in section 2.2. In a nutshell, they prescribe ways in which one is related to oneself and the way one's relations' relations are our own relations.

Each time our audiences were first confronted with the need for defining a social life for sets, the immediate answer was:

"Set A is related to set B iff $A \in B$ "

Discussing the reason why this answer comes up spontaneously is well beyond the scope of this book⁴. The characterisation of the empty set based on this social life of sets is quite easy: the empty set is the only set that is related to every other set by this particular relationship.

Prompted for another definition of social life, the following answer was given several times:

"Set A is related to set B iff $A \notin B \neq \emptyset$ "

According to this definition, the empty set is the only set that is not related to any set; all non-empty sets relate at least to themselves. Incidentally, we shall see in section 2.2 that this definition of social life does not define a category, one of the reasons being that, in a category, every object is at least related to itself in a canonical way. Nevertheless, the example is useful for showing that changing the definition of social life may lead to quite different characterisations of the same objects.

We typically have to force the audience to come up with the "standard" definition of social life between sets:

⁴ Nevertheless, the author is willing to collaborate in any research project that aims at understanding the psychology of modern formalism.

1. INTRODUCTION

"A social relationship between a set A and a set B is given as a total function $f:A \rightarrow B$ "

The characterisation of the empty set in this case is very similar to the first one: the empty set is such that, given any other set A , there is one, and only one, function to A – the empty function. Notice that, although the concepts of "contains", "intersection" and "function" are defined via the membership relation, the characterisation of the empty set given in the three cases does not involve it directly.

For further evidence that the nature of objects changes according to the "social life" that is of interest, consider the characterisation of singletons. According to the first definition, a singleton is such that only the empty set and the singleton itself relate to it. In the second case, a singleton is such that it relates to the sets that contain the element: $\{a\}$ relates to B iff $a \in B$. This is a good evidence for the fact that the second definition is not very "categorical": it reduces to set membership and makes the identity of the elements visible. This is also the case with the first definition: in both cases, two singletons are socially equivalent, in the sense that they relate to any other set in the same way, iff they are equal. Incidentally, a possible set-theoretic characterisation of a singleton set $\{a\}$ is:

$$x \in \{a\} \text{ iff } x=a$$

The use of equality between elements makes clear the need to look inside the set.

The third definition characterises singletons as sets into which there is one, and only one, total function from any other set. Indeed, the function being total, a singleton offers no choice for the mapping to be established: everything is mapped to this single element. According to this definition, two singletons cannot be distinguished because they relate in exactly the same way with the other sets. This is quite intuitive because, not being able to use set membership, we cannot look inside the singletons and notice that they have different elements. Hence, this social life is a better abstraction from set membership than the other two.

The "social" way of characterising objects is, in fact, similar to the way, in object-oriented programming, the view of objects that matters for building systems is that of the methods that objects make available to the other objects. However, the view that matters for their implementation may be different, reflecting the fact that different uses reflect different structural views of the same concept (which, in Category Theory, means different categories).

Other examples could be given from branches of engineering or our day-to-day praxis. For instance, another example is the way we understand devices as mundane as hi-fi systems. When assembling a hi-fi system from separate components, what is important is to know how each component can be connected to the others. The way they are implemented in terms of microcircuits probably explains the restrictions on the connections that can be established, but it is something that a user would like to see abstracted away. The characterisation of human behaviour can also be used as an example. There is probably a neurophysiological justification for what we call "shy" or "expansive" characters, but the meaning that we normally attach to such features of human nature is derived from the way people interact with us.

In brief, like all mathematics, Category Theory is about providing us with abstraction mechanisms. In our opinion, the fact that these mechanisms relate to interaction make Category Theory particularly suited for Computing Science in general, and Software Engineering in particular.

1.3 Overview of typical application areas

We have already promised, perhaps too often, that we are about to put forward mathematical techniques that are useful for supporting software development. We have even mentioned a few specific application areas like Object-Oriented Programming and Software Architectures. Some of these applications will actually be developed, even if not entirely, throughout the book. In this section, we will summarise some of the main applications we know, with references to the literature, bearing in mind that our focus is Software Engineering and not Computer Science as a whole. We strongly suggest the reader to consult [26,56] for excellent introductions to (and overviews of) the field.

1.3.1 GENERAL SYSTEMS THEORY

When one thinks of it, this is really early work. Goguen started exploring Category Theory as a mathematical toolbox in the early 70s, applying it to General Systems Theory [51,52,63]. The famous *motto* [56] "*given a category of widgets, the operation of putting a system of widgets together to form a super-widget corresponds to taking a colimit of the diagram of widgets that shows how to interconnect them*" was first applied to mathematical models of system behaviour. We also find in this area the origins of latter applications to object-oriented modelling [57]. In fact, one of his early works is even called, very appropriately, "objects" [53].

But Goguen's early work is just one example of the many applications of Category Theory to the area of General Systems. We would like to encourage the non-specialist to get acquainted with the field through one of the recent books that have appeared in the wider market of "Science" like [70,95].

1.3.2 ALGEBRAIC DEVELOPMENT TECHNIQUES

This is one of the most "typical" areas of the applications of Category Theory to Computer Science. One of the earlier and most influential movements started in the second half of the 70s with the work of a group of researchers at IBM (Joseph Goguen, Jim Thatcher, Eric Wagner and Jesse Wright), very aptly named ADJ (for *adjunction*), around the initial semantics of abstract data type specification [64]. The pioneering work of Rod Burstall and Joseph Goguen around the language CLEAR [17,18] showed how simple universal constructions (e.g. colimits) could be used to give semantics to operations for structuring specifications (e.g. computing the sum of two specifications, parameterising a specification, etc). The area produced textbooks [e.g. 28,29,75,93] as well as surveys [8,26] that provide a good showcase for "categories at work", namely in what concerns elegance and economy of means.

One of the finest hours of this research programme was the development of the theory of Institutions, a categorical formalisation of the notion of Logic developed by Goguen and Burstall in the early 80s [61]. The aim of this effort was to provide a mathematical framework in which the specification building operators defined for the language CLEAR could be made independent of the underlying logic and thus available for other specification formalisms. The theory of Institutions is, in our opinion, one of the best examples of the usefulness of Category Theory as a "unifier" of concepts and techniques developed by different research teams in response to different

or similar needs. As a "science", Computing is still very young and, hence, fragmented in that it is often difficult to find the "universal laws" that we recognise as the objects of study in other sciences. Through the theory of Institutions, Category Theory showed to be a powerful tool for abstracting from individual and uncoordinated efforts some "universal laws" (or "dogmas" as they are called in [56]) that apply to specifications in general [e.g. 96]. This categorical framework has been subsequently extended and put to use, for example, in applying algebraic specification techniques to other computational paradigms (e.g. object-oriented), and defining ways for specifications to be mapped from one formalism to another.

1.3.3 CONCURRENCY AND OBJECT-ORIENTED SYSTEMS

Concurrency theory is another area to which categorical techniques have been applied in the "engineering" view that interests us in this book. This domain of application is dominated by the work of Glynn Winskel [109,110,111] who showed how models for concurrent system behaviour like transition systems, synchronisation trees and event structures can be formalised in Category Theory. The idea is that each process model is endowed with a notion of morphism that defines a category in which typical operations of process calculi are given as universal constructions. This work was continued in the 90s in collaboration with Mogens Nielsen and Vladimiro Sassone [98], exploring the use of adjunctions as a means of translating between different models. The chapter [112] provides an excellent overview of this application area.

This stream of work was extended for giving semantics to object-oriented systems by J.Félix Costa, Hans-Dieter Ehrich, and Amílcar and Cristina Sernadas [e.g. 20,21,25], also with contributions from Goguen [24]. In a nutshell, these authors showed how universal constructions can be used to express object-oriented features like encapsulation, inheritance and composition in concurrency models endowed with richer notions of state. The work of Goguen on sheaf-theoretic models [57] in particular establishes an interesting connection to the earlier work on General Systems Theory.

Applications of Category Theory to the specification of concurrent and object-oriented systems can also be found. These include our work in the late 80s and early 90s [38,39,40,43]. Basically, we showed how the techniques developed for the modularisation of equational and first-order logic specifications of abstract data types could be applied to concurrent and object-oriented systems by using the modal and temporal logics that had been developed for the specification of such systems (e.g. [13]).

Briefly, the idea was to show that the specification of such systems fitted in the "institutional" picture as set up by Goguen and Burstall [32,33]. We later showed that our work is directly related to the General Systems tradition [31], in the sense that the module structure defined by our categorical formalisation is directly compositional over the system structure. This lead to applications of the categorical techniques to parallel program design [44] in which the notion of morphism captures what in the literature is known as superposition or superimposition [19,47,71].

There is also a substantial amount of work in the application to concurrent and object-oriented systems of the algebraic approach to abstract data type specification [7], most notably by Egidio Astesiano and his colleagues in Genova. Basically, these approaches present different alternative ways of bringing states and transitions to the universe of discourse of abstract data types.

1.3.4 SOFTWARE ARCHITECTURES

Applications of categorical techniques to the semantics of interface description and module interconnection languages were developed by Goguen in the early 80s [54] and more recently recast in the context of the emerging interest in Software Architectures [58]. These applications are in the tradition of the algebraic approach to abstract data types, more specifically in what has become known as "parameterised programming" [55], in the sense that they capture functional dependencies between the modules that need to be linked to constitute a given program.

The view of architectures that is captured in this way is somewhat different from the one followed in the work of Dewayne Perry, David Garlan and other researchers who have launched Software Architectures as we know them today [14,49,100]. This more recent trend focuses instead on the organisation of the *behaviour* of systems as compositions of components ruled by protocols for communication and synchronisation. As explained in [2], this kind of organisation is founded on *interaction* in the behavioural sense, which explains why formalisms like the CHAM [16] are preferred to the functional flavour of equational logic for the specification of architectural components.

This why, in the early 90s, we proposed a categorical toolset for architectural description based on our work on the formalisation of parallel program design techniques [34,42]. The idea is that, contrarily to most other formalisations of architectural concepts that can be found in the literature, Category Theory is not another semantic domain in which to formalise the description of components and connectors. Instead, through its universal constructions, it provides the very semantics of "interconnection", "configuration", "instantiation" and "composition", i.e. that which is related to the gross modularisation of complex systems.

1.3.5 SERVICE-ORIENTED SOFTWARE DEVELOPMENT

But there is more to Category Theory than the ability to support such interaction-based views of system behaviour. By relying on "local naming" (as shown by the impossibility of distinguishing between singleton sets), Category Theory requires all such interactions to be modelled explicitly and outside the participating objects. This is exactly one of the aspects that makes the distinction between service and object-oriented system development.

Clientship, i.e. the ability to establish client/supplier relations between objects, leads to systems of components that are too tightly coupled and rigid to support the levels of agility that are required to operate in environments that are "business time critical", namely those that make use of Web Services, B2B, P2P, or otherwise operate in what is known as "internet-time". Because interactions in object-oriented approaches are based on *identities*, in the sense that, through clientship, objects interact by invoking specific methods of specific objects (instances) to get something specific done, the resulting systems are too rigid to support such levels of agility. Any change on the collaborations that an object maintains with other objects needs to be performed at the level of the code that implements that object and, possibly, of the objects with which the new collaborations are established. On the contrary, interactions in a service-oriented approach should be based only on an abstract description of what is required, thus decoupling the "what one wants to be done" from the "who does it and how". This is, precisely, the discipline of interconnection that Category Theory enforces.

Hence, our last claim, in support of the last paragraph of section 1.1, is that Category Theory is, definitely, the mathematics of the Internet-age (and beyond)!

1.4 What can be found in this book

This book emerged from tutorials and courses given in the past few years both in academia and more industry-oriented fora like OOPSLA, ObjectWorld, TOOLS, ECOOP, and FME – Formal Methods Europe. This experience showed that there was an audience for this particular way of exposing Category Theory.

The book is structured in three parts, leaving room for different reading/teaching paths to be followed. With respect to most other books in the market, this one will use examples of a different nature, focus on, and give more emphasis to, aspects that are less common in other fields of Computing, and adopt a different pace altogether.

- Part One, i.e. chapters 2-5, covers some of the "basics" of Category Theory, the "bare essentials" that are addressed in any book, but with a different emphasis and tone that are meant to be more appealing and accessible to software engineers. Mature readers may, hopefully, appreciate a different way of exposing and illustrating these familiar concepts and constructions.
- The material included in Part Two, chapters 6 and 7, is of a more advanced nature, not only because it is more challenging from a mathematical point of view, but because it makes appeal to an additional level of maturity in so far as Computing Science is concerned. There, the reader will find material that only a few other books cover as well, although not all at the same time. Again, examples are drawn from areas that are not normally found in other publications. The sections in this part should be accessible to anyone with basic knowledge of Category Theory, but a quick travel through Part One will help the reader become familiar with the notation and the examples that are used in Part Two.
- Part Three offers the chance of seeing Category Theory "at work" in a more ambitious project – giving semantics to CommUnity, a prototype language for architectural modelling. It can be ignored by readers who are not particularly interested in the applications to Software Engineering. On the other hand, it can be followed, to a large extent, without the material exposed in Part Two. This means that a novice to Category Theory but interested in Software Engineering, or anyone whose goal is to understand CommUnity or setup a course that aims at teaching (any subset of) the material in part three, can safely skip Part Two. It will be necessary to go back to Part Two in order to understand in full the mathematical structures that relate to Software Architecture, but this may be done later on, once the reader feels more at-ease with the mathematics.

Parts one and two use examples taken directly from Bertrand Meyer's book on Eiffel [88] to illustrate definitions and constructions. Readers who are not familiar with the particular notation of Eiffel, but are used to object-oriented modelling, will hopefully be able to understand the examples without much effort. The choice for Eiffel instead of a more "modern" language has to do with two facts. On the one hand, it *was* modern when the book started to be written... On the other hand, it is one of the object-oriented languages that has solid foundations, which allows it to be used to illustrate mathematical constructions such as the ones you are about to read.

1.3 Overview of typical application areas

The use of Eiffel allows us to illustrate applications of Category Theory to the more "static" aspects of system modelling, namely to what is related with classification (inheritance). Other examples are brought to bear from "Specification Theory" that involve Logics and Algebraic Models for system behaviour in order to show how the more "dynamic" and "evolutionary" aspects can also be handled. These different aspects come together in Part Three.

This means that different reading/teaching paths can be established that are based on the examples: one can follow the "Eiffel path" for a lighter approach, perhaps more suitable for "practitioners"; or one can follow the "Specification Theory path" that will enable the "scientists" to visit more challenging places and reach the end of Part Two.

