
Introduction

No classification system, any more than any representation, may specify completely the wildness and complexity of what is represented.

(Bowker and Star, 1999: 232)

Taxonomies are at the same time deceptively simple and fiendishly complicated. They are simple because they are absolutely basic to human consciousness, so everybody manipulates and creates them with great ease. Our world is populated with categories, of family, friends, social groups, things, concepts, activities, feelings, places, times and many other things besides.

Taxonomies are complicated because we use them for the most part unreflectingly – they are simply part of our mental and social background – and we use categories in a huge variety of ways, often in competing and inconsistent ways.

This makes it difficult to be consciously consistent and consistently conscious in our use of categories, and the implicit taxonomies of which they are a part. And yet to manage our knowledge we do need to be both conscious and consistent.

Consider the case of going for a drive. To achieve this relatively simple task we need to be able to manipulate a category system involving cars and parts of cars (at least the parts that we manipulate); categories of road sign; rules of the road and their relative importance; types of traffic conditions and appropriate responses; categories of other drivers and appropriate responses; categories of vehicles and their capabilities. Except among chauvinists who like to complain about women drivers, or among driving test candidates, very little of this taxonomic knowledge comes consciously to the surface. Our daily taxonomies remain largely tacit. As we'll see later on, most of the taxonomy work we do in organisations is also invisible and not consciously organised.

Moreover, human beings are programmed with the 'Babel Instinct'. If

we can organise things around us differently from other people, we will do so. Sometimes this is for pragmatic working needs – we like our things organised in such and such a way because it suits the tasks we have to do. Different people with different tasks will order their knowledge assets differently to suit the tasks they have to do. Sometimes we just organise them differently because we feel like it. We don't always know how to discriminate between pragmatic need and arbitrary inclination.

These few factors are what make taxonomy work difficult. Our clients can and do categorise continually. They are confident about knowing how to sort things out. It is easy for them to find fault with the taxonomies we design for them. They can be fluent in critique. But they are not skilled at conscious, strategic organisation of their knowledge assets to suit collective needs. In knowledge management, much of the work of a taxonomist is not in analysis, but in 'reading' the varying information perspectives of different groups in the client organisation, 'collecting' their languages and labels, and helping them reach a negotiated, well-structured compromise. It goes beyond that, to working within an organisation's information environment to make sure that the taxonomy is understood pragmatically, adopted consistently, applied productively and managed sustainably.

There are many contradictions in taxonomy work. It is intensely democratic, yet it is also a highly specialised art – everyone can do it, but few can do it well. Good taxonomies are simple; they become invisible and taken for granted, because they reflect so well the contours of their users' knowledge world – but only complex, difficult taxonomies are held to represent the true art of taxonomy building. Taxonomies work on principles of consistency and predictability, yet they must also accommodate inconsistency, contradiction and ambiguity, because so do the knowledge worlds that we are trying to navigate. Taxonomies are a losing battle, sandcastles shored up against the rising tide of change – but we fight nevertheless, because they give temporary respite from advancing chaos. Taxonomies are commissioned, constructed and managed as products, yet the most important part of taxonomies lies in the processes and environments that produce them, and the processes and environments where they are employed. Taxonomies make knowledge visible, but while they reveal, they also conceal – the 'stuff' that is not accommodated at all, and the attributes of our knowledge that our taxonomy builders considered of secondary importance. A taxonomy is a standard, and yet it is also highly contingent on current circumstance.

It's no wonder that such a confusing picture gives rise to popular assumptions about taxonomies that are only partially true, misleading, or just plain wrong. One of these, that it's a highly arcane domain inaccessible to the ordinary human being, we've already dealt with. Another, that only librarians and biologists understand taxonomies, is misleading – biological taxonomies are paragons of consistency and purity of principle, but they are totally unlike taxonomies for knowledge work. Our messy, confused world of knowledge and information artefacts does not follow the simpler laws of genetics. And while they are usually well-schooled in their own specific sets of classification principles, neither biologists nor librarians, for the most part, ever have to build taxonomies. In their professional roles, they will at most be passive users of existing taxonomic schemes. There is no ready reservoir of taxonomy construction experts – we are all muddling through.

In the first half of this book we'll challenge a number of assumptions about taxonomies and the work of taxonomy building, and relate this work to organisation effectiveness and knowledge management. Chapter 1 defines our terminology and introduces the basic concepts we'll be working with throughout the book. In Chapter 2, we tackle the assumption that a taxonomy has to look like a hierarchical tree. In Chapter 3 we show that taxonomy work is an integral part of information infrastructure development going far beyond information retrieval. In Chapter 4 we look in more detail at how taxonomy work influences the basic things that organisations do to be effective. Chapter 5 traces the history of taxonomies in knowledge management and challenges the assumption that taxonomy work is just a specialised area of work within knowledge management associated with content management and information retrieval. In that chapter too we look in more detail at the variety of contributions that taxonomy work can make to knowledge management initiatives.

In the second half of this book, we take a more practical approach and guide you through the steps involved in a 'typical' taxonomy project. Here we challenge the assumption that taxonomy development can be done in the abstract, by a consultant, sitting apart from the information and knowledge world of the organisation it is intended for. Very few taxonomies for knowledge management can be developed in that distant, unengaged way.

In Chapter 6 we look at the practical things that taxonomies can do for us in organisations, and how different taxonomies work towards different results. It is possible to do a lot of damage by applying taxonomies badly. Chapter 7 outlines the key steps that need to be

walked through in planning and preparing for a taxonomy project. Chapter 8 looks at the typical activities in designing and validating a taxonomy, while Chapter 9 looks at implementation and change management issues.

To close, in Chapter 10 we take a forward look at issues and challenges on the horizon for knowledge managers. What do the semantic web, folksonomies, ontologies and social tagging mean for taxonomy work? Will we need taxonomies at all? Here we challenge the assumption that taxonomies are the only 'true' way to organise and connect to information content.