The Autism of Knowledge Management

By Patrick Lambe

There is a profound and dangerous autism in the way we describe knowledge management and e-learning. At its root is an obsessive fascination with the idea of knowledge as *content*, as *object*, and as manipulable *artefact*. It is accompanied by an almost psychotic blindness to the human *experiences* of knowing, learning, communicating, formulating, recognising, adapting, miscommunicating, forgetting, noticing, ignoring, choosing, liking, disliking, remembering and misremembering.

This autism is profound in its pervasiveness and in the intensity of effort devoted to its pursuit. It is dangerous in the way that it constructs unacknowledged power relationships into the use of technology, in the way that it resources depleted learning strategies and competencies to the exclusion of richer ones, and it is dangerous in its uncritical imperialism in relation to knowledge and learning.

Connectivity-oriented learning and knowledge management strategies, exemplified by current valuable work on communities of practice and social capital, exist in an uneasy alliance with content-oriented autism.[1] The two operate as distinct continua, represented by different technologies and practices, although the content approach often attempts "knowledge capture" strategies from within the web of interactions that successful social connectivity brings. The separation of content from connectivity further amplifies the failure of both management and technology to better understand and serve the knowledge and learning needs of workers in a hyper-connected, volatile and complex knowledge universe.

Autism

Autism, first described by American psychiatrist Leo Kanner in 1943, is thought to result from a brain disorder that takes place during the first two and a half years of childhood. It is characterised by social disconnectedness, failure to recognise and read the subtleties of human communication behaviours and interactions, an obsessive addiction to routines and repeatable behaviours, and what psychiatrists call *meaningless noncontextual echolalia*, the repetition of sentences and words without regard to their significance or the context in which they are spoken.

All of the elements of object oriented knowledge and learning management are present: disregard for the actual rituals and subtleties of human communication and interaction, disregard for the unique contexts within which the temporary construction of meaning takes place, disregard for the uniqueness of each learning situation, its history and its

[1] Jean Lave and Etienne Wenger *Situated Learning: Legitimate Peripheral Interaction* (Cambridge: CUP, 1991); Don Cohen and Laurence Prusak *In Good Company: How Social Capital Makes Organizations Work* (Cambridge, Mass: Harvard Business School Press, 2001); John Seely Brown and Paul Duguid *The Social Life of Information* (Cambridge, Mass: Harvard Business School Press, 2000)

direction, the addiction to repeatability and reuse, supported where necessary by a spurious economic argument of scalability and reach – spurious because the argument only works in the very few cases where the same knowledge artefacts actually do have widespread and repeated application.[2]

Where does this autism come from? Some say it is endemic to the environment of technology development. In a provocative December 2001 article, Steve Silberman of *Wired Magazine* called autism and its milder cousin Asperger's Syndrome "The Geek Syndrome" and suggested that its apparent increase in frequency in Silicon Valley was the result of a gene-mix that favoured maths and technology brilliance, but disfavoured social and communication skills.[3]

The suggestion is highly speculative, and while it may seem to find support in the occasional occurrence of *idiot savant* characteristics in autistic adults (extraordinary mathematical computation, visualization and memorization skills), it does not in any way reflect the true working environment of most technology development – intense team-based environments where disciplined communication and tight working relationships are absolutely fundamental to success – nor of technology entrepreneurship, a field of activity highly dependent on the ability to communicate, and form and maintain dense networks of supportive relationships. This is not a naturally autistic environment.

More promisingly, a portrait of autism also emerges from Max Weber's characterisation of the bureaucratic mode of organization as one that systematically "eliminates from official business… love, hatred, and all personal irrational elements."[4] The disjunction between the idea of working life as described in bureaucratic management protocols and management texts, and our real-world emotional and relational existence, says it all. It's almost too easy to point to bureaucracy's tendency to practice meaningless noncontextual echolalia (broadcast mantras and management visions and doctrines largely divorced from immediate reality), its mania for habit and routine, its intolerance towards change, and its general depersonalisation of communication.[5] They are all core features of autism.

And of course, such are also the features of the systematic management school that began to emerge in the United States from the 1870s onwards, as the complexity of managing consistency across large enterprises and over large distances sparked a crisis of control in American industry. "Only by replacing individual idiosyncrasy with system, individual memory with organizational memory, and personal skills with firm-specified skills at all levels did the systematizers feel that they could achieve the current and future efficiency they sought."[6] The result was a view of the organization as a machine made up of

[2] "When you think of it, building all these courses from scratch is like building kit cars in your garage. We should be mass-producing and distributing knowledge with higher-quality, predictable results with less cost." Richard Close and James Li, 'White Paper: The promise of eLearning and the practice of Knowledge System Design' (Irvine, California: Leadingway Corporation, 2000) p.5
[3] Steve Silberman, 'The Geek Syndrome' *Wired Magazine* December 2001
[4] Max Weber *Economy and Society* (New York: Bedminster Press, 1968) p.973
[5] Stephen Fineman, 'Organizing and Emotion' in J. Hassard and M. Parker eds *Towards a New Organization Theory* (London: Routledge, 1994)
[6] JoAnne Yates *Control Through Communication* (Baltimore: Johns Hopkins University Press, 1989) p.12

interchangeable employee-parts defined by task and job description, rather than the earlier, more organic analogies of family, nation or person.

The concerns that plagued large enterprises in the later nineteenth century are exactly the same as those that plague us today. They revolve around the problems of complexity, scale and distance, and the anxieties created by technologies that propagate opportunities for disorder at the same time as they promise tools for greater control.

In stark contrast to the apparent similarities between our condition and theirs, however, the rhetoric of the knowledge-based economy centres on supposed *differences* between the "new" economy and the "old" economy. It declaims the flattening of hierarchies, the devolution of control, the importance of relationships, the promise of mass-customization and personalization.[7] Yet scratch beneath the surface, and the old scientific management paradigms, based on the autism of the bureaucratic worldview, come swiftly to light. You don't have to look very far or very hard to find explicit reference to the icons of the industrialised, bureaucratic worldview, Frederick Taylor and Henry Ford, deep in the heart of the supposedly all-new, knowledge-based economy scripture.

Consider the second chapter of the influential book, *The Instructional Use of Learning Objects*, possibly the most coherent and consistent manifesto for object-oriented learning and knowledge management.

> The industry that focuses on the design, development, and delivery of computerized instruction is currently undergoing a period of standard setting focused on the distribution of instructional experiences over the Internet and World Wide Web. The instructional object – indexed by metadata – has great potential as a common building block for a diverse range of technology-based instructional products…
>
> It is hard to resist comparing these events with events in the history of steel-making technology. When Frederick Taylor showed in the opening years of the 20th century that reliable recipes for steel could be placed in the hands of relatively untrained furnace operators, an army of new and less-trained but fully competent furnace operators began to take over the mills. Greater quantities of steel could be produced at more precisely controlled levels of quality… Without these developments, steel quality would still be highly variable, steel products would have a much narrower range, and steel making would still be essentially an idiosyncratic craft practised by highly trained and apprenticed furnace operators.[8]

Further chapters in the book go on to make the case for creating similar, de-skilled access to rapid and consistent knowledge object creation, through the use of recipes of control, standards and scalability. That there is a qualitative difference between the process of steelmaking and *learning as a human experience* laden as it is with emotive colouring, and nested in an intricate, ever-changing web of relationships, is not noticed, or it is ignored.

---

[7] That many of these themes emerged quite early on in the 20th century as perennial issues in the wake of rapid industrialization, is conveniently forgotten. Cf Roy Jacques, *Manufacturing the Employee: Management Knowledge From the 19th to the 21st Centuries* (London: Sage, 1996)

[8] Andrew S. Gibbons, Jon Nelson, Robert Richards 'The Nature and Origin of Instructional Objects' in David A. Wiley ed., *The Instructional Use of Learning Objects* (Bloomington, Indiana: Agency for Instructional Technology, 2002)

Even the relatively humanistic approach of James Li's Leadingway Corporation cannot evade the autism of the mechanised view of knowledge: "In the same way Henry Ford figured out how to move from crafting cars to manufacturing and distributing cars, Dr Li figured out how to manufacture, manage, and distribute raw knowledge."[9]

One could posit arguments against the de-skilling and dehumanization of labour. There are plenty of precedents, from Thomas Carlyle's protest in 1829 that people were becoming "mechanical in head and heart, as well as in hand"[10] to Richard Sennett's observations on the transformation of a Boston bakery over a period of twenty-five years.[11]

Our point here is not a moral one, however, but a logical one. Approaching the use, acquisition, creation and adaptation of knowledge as if it is primarily a mechanical exercise that manipulates and processes *stuff*, akin to a blast furnace, is a *category mistake.* The analogies don't wash, because knowledge itself does not behave like physical stuff, and we only partially behave around knowledge as if it is stuff. The word "knowledge" is a noun, only because we make it so, not because it is a thing to be manhandled. If autism pays off in steel production, it most certainly does not in knowledge production.

Objects

We need to take a closer look at the tightly knit and over-influential family of ideas around reusable knowledge, learning and information objects to see why autism doesn't work.

There are several definitions of what reusable knowledge, information or learning objects are. M. David Merrill of the University of Utah is one of the more significant contributors to the school of object oriented knowledge. His work on the structures, classification and order of instructional knowledge objects mirrors the extraordinary influence and achievements of Melvil Dewey, creator of the Dewey Decimal Classification system for libraries in the nineteenth century.

Dewey reduced the universe of knowledge to 999 decimal categories, thereby imposing the illusion of correctness and the guilt of individual variance from the norm upon

[9] Richard Close and James Li, 'White Paper: The promise of eLearning and the practice of Knowledge System Design' (Irvine, California: Leadingway Corporation, 2000) p.5

[10] Thomas Carlyle 'Signs of the Times' in *The Collected Works of Thomas Carlyle vol.3* (London: Chapman and Hall, 1858)

[11] "Computerized baking has profoundly changed the balletic physical activities of the shop floor. Now the bakers make no physical contact with the materials or the loaves of bread, monitoring the entire process via on-screen icons which depict, for instance, images of bread color derived from data about the temperature and baking time of the ovens; few bakers actually see the loaves of bread they make… As a result of working in this way, the bakers no longer actually know how to bake bread… The work is no longer legible to them, in the sense of understanding what they are doing." Richard Sennett *The Corrosion of Character* (New York: W.W. Norton & Company, 1998) p.68

generations of librarians who followed him.[12] The value of his scheme was great, in giving intuitive access to large collections of books; but it is also limited, in its perennial lack of hospitality to newly emerging or newly related fields of knowledge.

Similarly, in its mechanistic and doctrinaire definition and classification of technology mediated knowledge and instruction, Merrill's work adds analytical value, but practical nightmares for devotees.

The focus of this approach is completely component-driven: "Learning cannot occur if the necessary knowledge components are missing. Learning will not be effective or efficient if the knowledge components are incomplete or inappropriate."[13] The implication, is that effective and efficient learning can be completely expressed by the deployment of necessary, complete and appropriate knowledge components.

Compare this detached view with Akira Kurosawa's recollection of his attempts to learn calligraphy as a child:

> When a student felt he had done a good piece of writing, he would carry it up to the teacher with great trepidation. The teacher would look at it and take a brush with red ink to correct the strokes he did not like. This procedure would be repeated over and over again. Finally, when the teacher approved the student's writing sample, he would take out a seal I couldn't read because it was carved in ancient seal script and stamp it in blue on the side of the student's work. Everyone called this the Blue Seal, and when you got the Blue Seal you could go home for the day. Since I wanted nothing except permission to leave quickly and go to Mr Tachikawa's house, I applied myself with fervor to copying the teacher's calligraphy. But you can't love what you don't like.[14]

At one level Kurosawa's learning process was itself a repetitive, mechanistic one, full of the notions of correctness, sequence and structure. Yet his knowledge of calligraphy is a social one, intimately bound up with feelings of constraint and limitation, situated forever in the historical and mythical world of fearful fellow students and the unexciting teacher with the mysterious Blue Seal. Nothing of the component theory of Merrill will convey the true essence of the learning experience, nor the character of the knowledge that endures – not even metadata will do that.

It is of the school that Merrill represents, that one sometimes hears this particularly unkind variation of the disrespectful adage:

> *Those who can, do.*
> *Those who can't do, teach.*
> *Those who can't teach, become instructional designers*[15]

All learning has context, and it has historicity. In both dimensions, in its context and its historicity, knowledge is imbued with meaning and emotion far beyond its informational

[12] Wayne A. Wiegand, *Irrepressible Reformer: A Biography of Melvil Dewey* (Chicago: American Library Association, 1996)
[13] M. David Merrill 'Knowledge Objects' *CBT Solutions* Mar/Apr 1998
[14] Akira Kurosawa *Something Like an Autobiography* (New York: Vintage Books, 1983) p.27
[15] Last heard at a workshop at Online Learning Asia, May 2002 in Singapore

content, and it is netted in a social understanding of the world. It is layered in time, overlaid, often obscured and sometimes revived and resurfaced, to take on fresh shades of significance. It has a past and a future. It means different things to different people. Knowledge as we use it is organic and contiguous to our existence as continuous, conscious and social identities. The snapshots and freeze frames of knowledge objects, and the elegant artifices of more complex knowledge artefacts such as books, movies and love letters, are static markers and extrusions of knowledge, not to be mistaken for the processes of learning and knowing themselves.

The clearest and most accessible description of information objects, learning objects, and how to compile them into learning events, was developed by a team of researchers at Cisco Systems.[16] The white paper they produced in 1999 is brilliant, simple and deceptively easy to comprehend and apply. Intended as an internal design document for themselves and third party content developers, its legibility has given it far wider reach and influence.

We should be cautious about deceptively easy things where processes in human heads are concerned. A colleague told me recently of a conversation he had had with a senior Cisco executive at a conference. "These reusable learning objects are fantastic" the executive stated enthusiastically, and paused. "The only trouble is, we haven't figured out how to reuse them yet."

Now this puzzle takes place in a relatively easy learning environment to serve: primarily technical, descriptive content to be served up to a limited and well-defined class of people – sales people and technicians. Cisco's learning objects evidently work well for the specific and limited objectives they serve. But they are not especially reusable. Their merits lie in their ease of production and deployment more than on their characteristics as components of larger wholes, or on their reusability as components. And if a recent shift of focus at Cisco is anything to go by, senior executives at Cisco recognise this. At the Online Learning Asia conference 2002, Cisco's Asia Pacific Solution Development Director Dominic Scott moved the focus off learning object theory and onto connectivity: the use of technology to help people to share and learn through *communicating*.[17]

Even the enthusiasts find it hard going. The Wisconsin Technical College System launched an Online Resource Center focused on learning objects creation and use in October 1999. One of the major challenges in the project was the tendency of developers to think in terms of whole learning experiences, rather than small stand-alone components that could combine and recombine with others. "[D]evelopers created highly innovative activities but did so by creating them thematically rather than looking at each

[16] Chuck Barritt, Deborah Lewis, Wayne Wieseler 'Cisco Systems Reusable Information Object Strategy' (San Jose, California: Cisco Systems, 1999)
[17] Dominic Scott 'Learning at Internet Speed: A Cisco Systems Case Study' keynote address at Online Learning Asia, Singapore May 16 2002

as an independent, self-contained chunk of learning; many in this collection needed to be re-adapted for use as learning objects that could stand alone."[18]

It was with some difficulty that the project team themselves created two sample learning objects to illustrate precisely what they should look like. Disengaging a piece of knowledge from its context is a remarkably difficult thing to do, even when you're trying to do it. The Cisco experience suggests that even if you manage to do that, the object's application to other contexts, its *reusability*, suffers. Context neutrality seems to disengage knowledge and learning from its immediate relevance, and makes it harder to ascribe significance to it.

Where does this failure spring from, and why can't the enthusiasts see the failure? Where people can't see problems, where enthusiasm and energy and rhetoric drive direction and investment of resources, you can usually smell a myth or two. In the case of learning and knowledge objects, there are five big myths.

The Myth of Reusability

The myth of reusability imagines a system of knowledge and learning components that can be combined and recombined in lots of different ways. Each "brick" has learning objectives and a test of understanding attached. All you do when you create a new programme is compile all the different pieces from the repository, and sequence them accordingly. The same objects can be used in lots of different courses, and in lots of different contexts, for lots of different types of people.

But even if knowledge did behave like normal physical objects (which it doesn't), reusability is often no more than skin deep even with relatively simple objects. Take a very humdrum object, an orange jumpsuit. To be a jumpsuit, it needs to fulfil only a very few simple criteria: it's a one-piece coverall, normally goes from neck to ankle, and usually fastens at the front. To be an orange jump suit it also has to be orange. What could be simpler than that? A better candidate for a reusable object in lots of contexts, you could not find. Your principal concern, surely, would be the simple one of providing a range of different sizes – and we could call that our personalization strategy.

Probe a little deeper, and context starts to intervene. Our jumpsuit might be needed by a repairman, mechanic or housepainter who frequently works outdoors in cold weather. He'll need a thicker, warmer, more durable material. And he'll want more pockets. Or it might be a jumpsuit worn in the cleanroom of a high-tech manufacturing company, where the properties of the material will include sophisticated requirements about how it reacts to moisture, static and dust.

[18] Kay Chitwood, Carol May, David Bunnow, Terri Langan 'Battle Stories from the Field: Wisconsin Online Resource Center Learning Objects Project' in David A. Wiley ed., *The Instructional Use of Learning Objects* (Bloomington, Indiana: Agency for Instructional Technology, 2002)

Thus far, we've stayed in the realm of cool, objective fact. The detailed properties of the jumpsuit can vary significantly just by virtue of the function their wearer fulfils, and the environment they work in.

But jumpsuits have other meanings as well, more socially constructed than functionally driven. Orange jumpsuits of a much more sophisticated and specialised kind are worn by NASA astronauts and by some Formula One drivers. In the latter case, padding against the cold and fire retardant material is critical. In both cases, however, the jumpsuits go beyond their functional use: they are also badges of honour. To wear them is a privilege, and invites respect. On the other hand, orange jumpsuits of a different kind are also standard issue for convicts in the United States, a fact we noticed suddenly when we saw hobbled and masked Al Qaeda prisoners shuffling into their prison compound at the US base in Guantanamo Bay. These jumpsuits do not have any pockets.

Orange jumpsuits, despite their apparent simplicity, are not reusable across categories, neither by function nor by meaning. We might be proud to possess a NASA or Formula One jumpsuit, indeed we might pay a lot of money for the privilege. We'd feel less enthusiastic about the prison garb of Guantanamo Bay, unless our affiliations placed Al Qaeda prisoners into the same heroic mould that NASA astronauts and Formula One drivers possess.

So if even a simple object like the jumpsuit cannot travel far out of its native context, why would we expect that an abstraction, an indeterminate knowledge object, would?

The Myth of Universality

The myth of universality is more of a blind assumption than a true myth. But it is the true death knell to spurious arguments about the economic value of mass producing standard knowledge objects for global distribution. The assumption, in simple terms, is that the same piece of knowledge can be applied universally. It's true, and relevant, everywhere.

There's some curious logic here. It's clearly not true when you look at all *times* and places. A piece of knowledge about the circumstances of John F. Kennedy's assassination is not applicable before he died, it is very important around the time of his death, and recedes in significance with every passing year after that. Knowing how to drive onto a freeway at speed and filter into heavy traffic using only your rearview mirror is an important (for us) piece of learning that would not have been much use in 1930, and is not much use now in freeway-free environments.[19] Conversely, knowing how to operate a telegraph and use Morse Code is now rarely needed, where once it provided a serious employment opportunity.

Even if you focus on the application of a piece of *currently true* knowledge to lots of places at the same time, the significance of *here and now* has a big impact on the relevance and importance of the knowledge. If somebody tells me that a nuclear bomb has fallen on London, I will react very differently depending on whether I am in New

[19] I owe this example to a conversation with Douglas Englebart.

York or in the suburbs of London. If I'm in New York, I will run for the phone to contact my friends and relatives. If I'm near London, I'll simply run.

Our knowledge needs differ depending on who, where and when we are. This gets clearer when you look at learning objects. Let's consider a nice simple topic: customer service. There could be nothing more straightforward, if the consistency of content in textbooks and training courses is anything to go by. But if you were designing an e-learning course for private bank employees in Wall Street, you'd be training your employees to judge the quality of cashmere in suits, and calculate customer lifetime value. If you were training bank tellers in Shanghai, you'd be teaching them not to throw the small change onto the counter with such a flourish. And in Dhaka, you might be training them to say good day in a pleasant fashion, and start small conversations.

Different circumstances and contexts have different knowledge and learning needs, even within the same corporation, despite the veneer of similarity cast by the long reach of Hollywood, Coca-Cola and Macdonald's. There is no such thing as universally applicable knowledge, and this is why the market for the localisation of instruction manuals, software and e-learning has blossomed in recent years.

But the term "localisation" itself is an impossibly light rendering of the real task. A truer model of knowledge transfer would be less like moving a piece of stuff and maybe sculpting it a bit to local needs, and much more like a local construction of knowledge at the source of need, achieved by checking our own experience and observations, asking other people's opinions and looking at other people's apparently relevant knowledge artefacts. The more localised these artefacts are, the easier it may be to accept them as apparently relevant, but the artefacts themselves do not accomplish the knowing.

The Myth of Interchangeability

The myth of interchangeability is a critical plank in the reusability hypothesis. It's not simply based on a translation of manufacturing methodology into knowledge, though that's certainly a theme. It's also based on the very successful programming technique developed by the pioneers of hypertext, where instead of continually rewriting the same lines of code for essentially the same sub-routines, you simply point the programme to "call" that sub-routine from the library whenever it's needed.

When it comes to knowledge and learning, different knowledge or learning objects can be called for different learners at different times and in different situations: "…each execution of the program can be unique since the examples for illustration, exploration or practice can be randomly selected from the knowledge base for each student."[20]

The building blocks slot into the same place and fulfil exactly the same function in a completely interchangeable way. It's an engineer's dream, and it really does inspire developers with a vision of more effective knowledge productivity, in the same way that

[20] M. David Merrill 'Knowledge Objects' *CBT Solutions* Mar/Apr 1998

the standardization of parts in machine manufacture enabled the huge productivity gains of mass production.

So let's look at a machine, made of interchangeable parts. Is interchangeability sufficient for its smooth operation? Let's look at a photocopier in an office with a service agreement that allows parts to be changed whenever there's a problem. Let's say that the photocopier suddenly starts a strange pattern of behaviour. It jams repeatedly every Monday morning. Technicians come and replace the rollers, but it still happens, and it happens every week. It's clear that we don't solve this problem simply by replacing parts. It's clear that something beyond the sum of parts is creating the problem.

The smart technician will start asking questions about habits of use. He might discover, for example, that because of cost cutting measures, the airconditioning is now switched off in the building over the weekend, the paper already in the machine expands from the higher humidity, and it jams easily because the rollers are calibrated for thinner, dryer sheets. He might also discover that the sales presentation meetings have been moved from Fridays to Monday afternoons, placing added stress on the machine on Monday mornings.

Understanding the science, and being able to label and replace components does little to resolve problems that arise from social and human initiatives and changes. And most knowledge and learning issues arise precisely from social and human initiatives and changes.

Thus far, everything we've looked at has asserted the overriding primacy of *local context* for the *applicability* of knowledge and learning. The assumptions underlying object based content strategies are deeply flawed when it comes to this, treating what ought to be a fundamental design strategy as a cosmetic localisation (Merrill might say *instantiation*) issue instead.

Beyond a basic informational level (and value added knowledge and learning need to go far beyond basic informational levels), when I have a specific working problem such as how to resolve a complex financial issue, the last thing I want is a necklace of evenly manufactured knowledge nuggets cross-indexed and compiled according to the key words I happen to have entered into the engine. Google can give me that, in many ways more interestingly, because it will give me different perspectives, different depths and different takes.

What really adds value to my problem-solving will be an answer that cuts to the chase, gives me deep insight on the core of my problem, and gives me light supporting information at the fringes of the problem, with the capability to probe deeper if I feel like it. Better still if the answer can be framed in relation to something I already know, so that I can call more of my own experience and perceptions into play. Evenness and interchangeability will not work for me, because life and the situations we create are neither even, nor made up of interchangeable parts.

We do have an evolved mechanism for achieving such deep knowledge results: this is the performance you can expect from a well-networked person who can sustain relatively close relationships with friends, colleagues and peers, and can perform as well as request deep knowledge services of this kind.

The Myth of Completeness

The myth of completeness expresses the content architects' inability to see beyond the knowledge and learning delivery. Out of the box and into the head, and hey presto the stuff is known. The evidence for this is in the almost complete lack of attention to what happens outside the computerised storage and delivery mechanism – specifically, what people *do* with knowledge, how it transitions into action and behaviour. How many people in knowledge management are talking about synapses, or the soft stuff that goes on in people's heads? Is it simply assumed, that once the knowledge is delivered, it has been successfully transferred?

My former landlady in Singapore wants to start her own small food stall. Her problem is that if she wants good traffic, the rentals are very high and it's hard to make a profit. If she goes for a low rental stall, the turnover will be low, and it's hard to make a profit. So she is casting around for a specialty dish common from her childhood, that isn't so well known now. In Singapore, food is a craze: if you hit the right nerve with the right dish, people will travel from far and wide to sample it, and they'll tell everyone they know. Low rental and high traffic, good profits.

My landlady, Mary, fixed on a type of riceflour cake dish called *wah kueh*. Her next task is to find out how to make it – to obtain "the knowledge". The whole idea is that the dish should be scarce, so finding somebody who still makes it is difficult. Her best bet is to find somebody from her network of friends, whose parents made and sold it. It will have to be somebody fairly close to her, because such knowledge is rarely circulated outside the family, even for commercial incentives, and even though the knowledge is quite likely to die with its bearer.

Let's say she finds the parent of a close friend who knows how to make *wah kueh*. They won't be able to show her how to make it, because they no longer have access to the range of equipment required to make it. They'll have to tell her. And they won't tell her precise measurements, they'll just say "add this to that, and do this to it until it's ready". So I asked Mary, "How will you know you are making it correctly?" "I'll have to spend a couple of months feeding my family *wah kueh*, until I get the taste right" she replied.

This story, in miniature, is how we actually normally acquire knowledge. Knowledge only has value if it is emerges into actions, decisions and behaviours – that much is generally conceded. But few content-oriented knowledge managers think through the entire lifecycle of the knowledge objects they deal in. Acquiring a knowledge artefact is only the first stage of what's interesting about knowledge. We don't truly *know* until we have internalised, integrated into larger maps of what we know, practised, repeated, made myriad variations of mistake, built up our own personalised patterns of perception and experience. Emotions too will play their part: Mary's family will gradually increase the

pressure on Mary to get her knowledge right quickly, so that they can return to a normal diet.

We might glibly call this process a conversion of explicit knowledge to tacit knowledge, but this simple phrase is too easily identified with reading, hearing and understanding the intended meanings within artefacts. Deep, performative knowing is a longer, more complex process, and it's rarely considered by knowledge managers, largely because autistic technology and bureaucratic process have very little to do with it.[21]

We might think that we can solve Mary's knowledge problem much more efficiently by setting up a *wah kueh* school, or at least reconstructing more complete recipes. In the real world of limited time and limited resources, however, does Mary's need justify such investment? And if such knowledge were so easily accessed, would Mary's enterprise continue to have value? It's the scarcity of the knowledge that will drive her anticipated profits.

And when it comes to complex knowledge problems, we really do want people to struggle over the internalisation and integration of knowledge artefacts. In complex, nuanced cases, we certainly want people to *refer* to relevant knowledge artefacts, but we don't want simple *reuse*. In August 2000 Ernst and Young of India came under serious attack when it was discovered that an Environmental Impact Assessment they had prepared for Murdeshwar Power on the $40 million Dandeli dam project in Karnataka, had been substantially plagiarised from a report on a different project 145 km away. If the first report had been created by Ernst and Young, it might have made a nice knowledge management case study about reuse of knowledge artefacts. The fact that the first report had been written by the Bangalore Environment Support Group made it plagiarism.[22]

Why do we get worked up about plagiarism, and how do we relate it to responsible reusability of knowledge artefacts? Part of the plagiarism debate is connected to the notion of property – this artefact is the product of *my work* – therefore you can't pass it off as your work, without attribution. There's a scent of fraud in the air if you do. But behind that, there's also the worry that Ernst and Young don't really understand the project they have been contracted to advise on.

"Two completely different rivers. Two completely different types of dams. Two completely different forest types to be submerged. Two different locations." Thus go the opening lines of the Bangalore Environment Support Group press release unveiling the scandal. The only expertise in this reuse of knowledge artefacts, is cut and paste expertise; and not especially complex cut and paste expertise – 60 pages of the 65 page report were exactly the same. Does Ernst and Young's advice on dam projects carry any weight, if reusability of knowledge artefacts is not supplemented by processing, integration, application to context, and reference to experience? To what extent is the

[21] Jeffrey Pfeffer and Robert I. Sutton *The Knowing-Doing Gap: How Smart Companies Turn Knowledge into Action* (Cambridge, Mass.: Harvard Business School Press, 2000)
[22] 'Ernst and Young rewriting dam report' *The Times of India* 3 September 2000

object-oriented knowledgebase designed for chunking and reusability simply authorised plagiarism, following the form rather than the substance of knowing and learning?

The gap between knowing and acting is a big gap, and a glaring one. The gap between knowledge and behaviour is a critical one to bridge: at its extreme, this gap lies behind any socially constructed catastrophe. It's becoming increasingly clear that there was sufficient knowledge in advance of the World Trade Center attack to have created at least the possibility of preventing or mitigating it. The knowledge was present, but *not actioned, or considered actionable, or it was inappropriately actioned*. The same is true of the Challenger space shuttle disaster, as of the BSE crisis in the United Kingdom.[23]

Knowledge management, e-learning and the technology used to support both, are not simply a delivery infrastructure for bits and bytes. The apportioning of resources and the rhetoric of both knowledge management and e-learning largely assume that they are.

There is another dimension to the myth of completeness. This perspective focuses almost completely on digitally mediated information, knowledge and learning. The challenge therefore is seen as a conversion challenge. Immense resources have been poured into conversion of different forms of knowledge representation into digital form, and on the capabilities of different media technologies so that different media can become digitally playable. Our drive to digitise everything compels us to broaden the range of types of knowledge representations we can capture, store and manipulate by this means.

There are lots of apparent advantages to this, not all of them valid on closer scrutiny. Practically cost-free reproduction and distribution is one of them – but are intellectual property rights and commercial interests at play? Preservation is another – but actually digital media are tethered to a complex technology and standards infrastructure whose persistence we cannot guarantee – who will be able to play mpeg3 files in 2050? Who can open WordStar files now? Paper-carried knowledge carries much simpler preservation issues.

Another argument for the digitizers loops back to the object-oriented approach. Many forms of knowledge are conveyed in linear format – the spoken word, the presentation or speech, the long message, the sustained reflection of a report on complex issues. Digitization allows for the possibility of rapid navigation via tagging and bookmarking, but even this is conceptually difficult unless you can easily map the structure of the artefact. Enter "chunking": the process of breaking down "long" pieces of knowledge into separable objects, carrying their own meanings independently of the whole. We've already looked at some of the difficulties posed by this approach in our critique of the myth of interchangeability.

[23] Diane Vaughan *The Challenger launch decision: risky technology, culture, and deviance at NASA* (Chicago: University of Chicago Press, 1996). The commission enquiring into the BSE crisis in the UK specifically criticised civil servants for the slowness with which they translated policy into action: *BSE Inquiry Report* (London: HMSO, 2000) volume 1 p.xviii

The higher difficulty is that as human beings we have amassed and are addicted to, a cornucopia of different types of *physical* knowledge artefacts. The electronic environment is a depleted one for such sensual creatures as we, and digital media, while giving us much greater *access* to knowledge, fare poorly in delivering up richness and emotional colour. I have in my possession a paperback edition of Thomas Mann's novel, *Buddenbrooks.* I took it on a cycling holiday to France in 1980, my first trip to that country. On the journey home, my bicycle fell over in the Gare du Nord railway station in Paris. The bottle of wine inside my bag, intended for my parents, smashed, liberally soaking the contents. I can still smell the fragrance of wine-soaked paper to this day, and the book carries far more than the typographical or literary significance it was intended for.

Why must we make the long journey to the Louvre to see the patina of age and the sheen of varnish on the *Mona Lisa*? Why doesn't the WebMuseum's jpeg file quite suffice? How is it that a ticket stub for a showing of Kurosawa's *Ran* in 1985 can immediately trigger an avalanche of perceptions, insights and comments to be shared with friends, where a database entry could not?[24] How is it that a three-minute story can communicate an insight to an audience when months of meetings with charts and figures and rational exhortations have failed?[25]

In fact, technology-mediated knowledge management is *incomplete,* unless it links us back into the physical, sensual, emotion-laden world we inhabit. Its pretence of separation, characterised by its autism and by its attempt to simulate tangible things and activities, is in fact a flight from the world, and dooms it to inefficacy. Unless knowledge management fuels our actions and behaviours in the physical world, it has no role to play. The people who design advertisements constantly worry about how their artefacts get translated into behaviours. Why don't knowledge managers? If they do, why do so few of them talk about it?

The obsession with digital objects that are pale representations of real things is an unhealthy one. It is a sad mistake to give the impression that you are recreating and preserving a heritage when you give depleted digital access to treasured works of knowledge and culture. In the absence of the real thing, or at a distance from it, providing depleted digital access is an invaluable service to perform. But it's not the same as the real thing, and does not constitute a liberation from the real thing.

The Myth of Liberation

So let's talk about liberation. You don't have to look far in the technical writing of knowledge-object enthusiasts to find almost mystical hyperbole about the potential of knowledge objects for humanity. Just look for the exclamation marks. Returning to the

[24] For a brilliant account of the roles, variety and richness of knowledge artefacts in our lives, see David Levy *Scrolling Forward: Making Sense of Documents in the Digital Age* (New York: Arcade, 2001)
[25] Stephen Denning *The Springboard: How Storytelling Ignites Action in Knowledge-Era Organizations* (Boston Mass.: Butterworth Heinemann, 2001)

manifesto produced by the Association for Instructional Technology, let's look at H. Wayne Hodgkins' vision for the future of learning objects.

> As we stand at the inflection point of a new learning economy, we realize that it will be shaped as we choose to shape it; it will be as rewarding and humane as we make it; the decisions we reach will determine what the world will be like for all of us.
>
> … As this vision of the future develops, it leads to one of the ultimate characteristics of the dream: tools and technology that truly have the ability to learn. It is all in the turn of a phrase, some say. Learning *about* technology is important. Technology *for* learning enables amazing results and advances. However, technology that *can* learn is going to provide the most revolutionary and significant change![26]

One piece of this espoused revolution is the creation of a new openness, where peer to peer knowledge sharing liberates individuals from traditional sources of power and control.

> Imagine what it means to have no concept of "servers" where everything is just a node on the net. A world where every person and every file can be connected directly, one to one. Think of the impact on learning, learners, and learning content. Think about every learning object connected to every other learning object, able to communicate, pass data, and manipulate the other. Think about a world where control of content is truly put in the hands of every individual, where everyone in need of a given piece of content can be connected directly with those who have it. What will it mean to have potentially billions of authors and publishers?[27]

We might point out that we have good reasons for the intricate social structures that validate, value and authenticate our knowledge artefacts. Pure knowledge sharing democracy without such structures is very hospitable to extremists who can propagate knowledge artefacts to their own narrow advantage and to the harm of others. But there's a deeper and more dangerous lie behind this apparent vision of democracy, and it's hinted at just prior to the last quotation.

> [The success of learning objects] will be dependant upon the masses being able to see their fundamentally high value and the ability of the masses to put them to use quickly and easily. This should NOT be confused with the underlying complexity that is required to make all this work and make it work transparently and easily. Indeed there is likely an inverse relationship between the external simplicity and ease of use of any technology and system, and the underlying complexity required to make it happen.[28]

The twin poles of simplicity for the masses and complexity designed by specialists are the two ingredients of a classic recipe for the definition (or redefinition) of power and control, whether conscious or not. Now we begin to sense a high priesthood behind our apparent knowledge democracy, powered by knowledge objects.

---

[26] H. Wayne Hodgkins 'The Future of Learning Objects' in David A. Wiley ed., *The Instructional Use of Learning Objects* (Bloomington, Indiana: Agency for Instructional Technology, 2002)
[27] H. Wayne Hodgkins 'The Future of Learning Objects' in David A. Wiley ed., *The Instructional Use of Learning Objects* (Bloomington, Indiana: Agency for Instructional Technology, 2002)
[28] H. Wayne Hodgkins 'The Future of Learning Objects' in David A. Wiley ed., *The Instructional Use of Learning Objects* (Bloomington, Indiana: Agency for Instructional Technology, 2002)

Let's pull apart this background complexity to see just how it gets constructed, and by whom. The ability to manipulate interchangeable, universal and reusable knowledge objects requires an underlying architecture, a structure of taxonomies and defined relationships, and a set of rules for how objects are created, defined, connected, moved and tagged. Hodgkins believes this is a matter for discovery, rather than creation – ie he believes such an architecture objectively exists, can be observed, and described.

> I believe we are on the verge of grand discovery in the areas of learning, content, knowledge and objects. Just as revolutionary as our discovery of the atomic and molecular models will be our discovery of the equivalent of the periodic table for all content or data. Mendeleyev's 1870s creation of the periodic table laid out the basic building blocks of all physical matter and revolutionized our view of that world. Similarly, an equivalent understanding of our data and informational world will give us a fundamental understanding and ability to manipulate, create, and build any substance possible.[29]

How does such an architecture get constructed? Hodgkins is less clear on that, but pioneering work at the US Defense Advanced Research Project Agency (DARPA) gives us some clues. Any knowledge base that learns must first begin with an ontology – a set of agreed definitions for what is true and valid in the subject area being considered.

> The ontology is the more general component of the knowledge base, characteristic to an entire domain, such as medicine or the military. A domain ontology specifies terms that are useful in a wide range of different applications in a domain. For example, a military ontology would include specifications of military units and military equipment that are likely to be included in the knowledge base of any agent developed for a particular military application. Moreover, there is wide agreement in any mature domain on the basic terms of this domain, allowing one to reuse ontological knowledge that was previously developed to build a new knowledge base.[30]

Once you have your ontology, you can then get your knowledge base to *learn* by interacting with a subject matter expert or experienced decision maker, and by abstracting and modelling their behaviours, their use of knowledge and their problem solving approaches.

The apparent objectivity of this approach collapses as soon as you look at how the ontology is derived. It must be widely agreed, and there must be common definitions with common meanings. Very little of real working life is run on agreed, common definitions. At best we run on approximations to that. Most of what we do is highly interpreted, time and place contingent, and constantly shifting. We disagree constantly on how to interpret the world.

The example of the military ontology is a good one. The ground of this ontology will be constantly shifting as military technology, organization, doctrine and geo-political circumstance change. Any working ontology will always be a compromise, a contingent snapshot taken for the purposes of doing something rather than nothing. Bear in mind that the ontology is fundamental to the knowledge base. Once a critical component of the

[29] H. Wayne Hodgkins 'The Future of Learning Objects' in David A. Wiley ed., *The Instructional Use of Learning Objects* (Bloomington, Indiana: Agency for Instructional Technology, 2002)
[30] Gheorge Tecuci, Mihai Boicu, Michael Bowman, Dorin Marcu, Murray Burke 'An Innovative Application from the DARPA Knowledge Bases Programs' *AI Magazine* Summer 2001

ontology diverges from the real world, the whole knowledge base becomes catastrophically unreliable.

> We now know that a knowledge-based system needs a great deal of knowledge to be truly useful. However, building a large and high-performance knowledge base is still a long, inefficient, and error-prone process.[31]

Hence the burning desire for reusable objects and ontologies. If you can reuse, you can build cheaply. But to be able to reuse, you have to go for the safest and slowest changing knowledge, and the most generic decision and problem solving strategies. Applicability to novel environments becomes a serious problem. At the level of modelling expert behaviours, the knowledge system is tied to the observation of historical behaviours of experts. Hence you have the question of how such a system can support new behaviours represented by innovation, rather than stifling them by constant recycling and reinforcement of past behaviours.

The essence of the knowledge object approach is to simulate a *predictable* environment so that it can predict what you need; and serve up pre-defined, appropriate objects to meet that predicted need. Far from being open, this is a *closed* system which assumes (falsely) that:

- the world can be objectively described in a stable way to everyone's satisfaction;
- subjective variations can be moderated using technology in a stable way to everyone's satisfaction;
- the world will not change faster than our capacity to describe it;
- some people are qualified to develop reliable architectures, ontologies and taxonomies that the rest of us will just trust, because we'll never understand how they really work.

So is this trust justified? The Y2K crisis gave us a vivid insight into the mythical objectivity of technology and software development. In April 1999, *Wired Magazine* carried a detailed expose of how software really gets written.

> It's almost a betrayal. After being told for years that technology is the path to a highly evolved future, it's come as something of a shock to discover that a computer system is not a shining city on a hill – perfect and ever new – but something more akin to an old farmhouse built bit by bit over decades by nonunion carpenters.
>
> … At its core, programming remains irrational – a time-consuming, painstaking, error-stalked process, out of which comes a functional but flawed piece of work… A programmer is presented with a task that a program must accomplish. But it is a task as a human sees it: full of unexpressed knowledge, implicit associations, allusions to allusions. Its coherence comes from knowledge structures deep in the body, from experience, memory. Somehow all this must be expressed in the constricted language of API, and all of the accumulated code must resolve into a set of instructions that can be performed by a machine that is, in essence, a giant calculator.[32]

[31] Gheorge Tecuci, Mihai Boicu, Michael Bowman, Dorin Marcu, Murray Burke 'An Innovative Application from the DARPA Knowledge Bases Programs' *AI Magazine* Summer 2001
[32] Ellen Ullman 'The Myth of Order' *Wired Magazine* April 1999

Irrationality and pragmatic compromise surround technology development and programming. They appear in how the computing industry is defined and supported and the stories it tells about itself. In how business practices and environments impose resource, functionality and schedule constraints. In the way that economic success or failure defines which solutions get to transmit their flaws to the next generation of applications.

Again, the argument for self-contained objects rears its head, closely followed by the argument for standards and certification. Interoperable, exchangeable, reusable programming objects allow for the possibility of cheaper development as well as the identification and replacement of flawed programming objects. But standards bring us back to that underlying architecture, and assume that we can construct that architecture in a pure, error-free and infinitely generative way – if we can't build software right, why would we expect that underlying, complex architectures can be made any better?

And standards are prone to infinitely more dangers than a single program. Standards burn bridges behind us, they deliberately limit the possibilities of what we can do in the future, they put us on a platform that isolates us from the affordances of the ground beneath. The brilliant scientist Nikola Tesla worked a century ago on the means to transmit wireless electricity free of charge. The dominance of standards for wire-borne electricity, supported by business cartels and political fears about the destructive possibilities for such technology, ensured the suppression of his patents, and closed the door for further research after his death.[33] While interoperability brings many benefits, to pursue a totalitarian concern with standards to the exclusion of variants in an immature field that is not understood well, is to flirt with disaster.

Let's return to the field of knowledge management and e-learning, where standards propaganda marches on, too infrequently challenged. In what way is the rigid application of standards for e-learning and knowledge management in our current state of awareness any different from a committee being formed in the Mainz of 1460 to legislate on the form, dimensions, specifications and design of the printed book for the purposes of proper housing and retrieval in libraries? The enterprise then would have been absurd, and it is absurd now. Standards, conventions and interoperability did emerge for the printed book, over decades, and through market forces, not by *diktat* of a self-appointed elite. A pragmatic concern for ensuring current interoperability of systems and content is one thing: the grander design of architecting an all-encompassing knowledge framework serving up knowledge objects wherever required, is foolish.

And we can see this absurdity in the inability of object-based standards to accommodate interesting and more complex forms of knowledge acquisition. Dr Geoffrey Ring of e-learning company ICUS understandably sees the strong advantages of standards in the field of online learning. But he also points out the difficulty of using more complex learning activities in the current, object-focussed environment.

[33] Marc J. Seifer *Wizard: The Life and Times of Nikola Tesla* (Toronto: Citadel, 1998)

> Learning design, the process of delivering learning experiences, is complex. The fast-growing simulation arena is one area beyond the reach of today's chunking and meta-tagging specifications. Learning is different from consuming content. Learning is active. The learning content is a commodity used by learners during a learning activity – the focus is the activity, not the content.[34]

In fact, anything with a world of organically connected elements, a scenario or a story in it, falters in the autistic world of objects. Stories and autism don't mix.[35]

Beyond Autism

And so we return to this notion of the life-cycle of knowledge – knowledge emerging into action and experience – in a contingent, continually changing world. The enterprise of constructing knowledge artefacts, whether it be as simple as a digital knowledge object, as intimate as a scented love letter or as grandiose as a knowledge ontology or periodic table, is similarly contingent, across its entire range. To suggest otherwise is dogmatic and fatally flawed.

To invest enormous energy in the myths of permanence, objectivity and universal applicability will be interpreted by our successors as either grand folly or dangerous folly, but folly nevertheless. The technology myth has been battered already by the investment collapses of the dot com bust, but it's also becoming apparent that our blind investments in this dream have been stupidly wasteful. In 2002 Gartner Group estimated that as much as 20% of global spending on technology is wasted.[36]

The least interesting thing about knowledge flowing down wires, printed on pages, painted on screens, or transmitted via communication signals, is how the signals are constituted and sequenced. Much more interesting is the complicity and sympathy of numerous brains acting in concert at the level of protocol (agreeing how to communicate, and knowing which genres to deploy when), while simultaneously pursuing and propagating variety and discord at the level of interpretation, constantly creating fresh, relevant, transient knowledge, out of the balletic dance and traffic of our more enduring but essentially static knowledge artefacts.

And even if, for the sake of content-oriented knowledge practitioners, we wanted to limit our purview to the play of knowledge artefacts in our lives, we would be foolish to forget that virtual knowledge artefacts represent only a small, unidimensional and impoverished selection of the full panoply available to us in our everyday lives, in the form of books, videos, magazines, letters, postcards, souvenirs, dog-eared photographs, stories, jokes and long-treasured gifts. Propaganda alone will not persuade us to migrate from our rich,

---

[34] Geoff Ring 'eLearning Standards and Reusable Learning Objects' (Singapore: ICUS, 2000) draft v.0.1

[35] Autistic children have difficulty in relating to stories as well. In narratives involving different characters, they can describe behaviours and sequences of actions, but fail to relate to the mental states of characters that influence their actions. This inability extends both to a depleted internal use of stories, as well as a failure to tell stories in ways that will engage the listener's attention and interest. Lisa Capps, Molly Losh, Christopher Thurber, ' "The Frog Ate the Bug and Made His Mouth Sad": Narrative Competence in Children with Autism' *Journal of Abnormal Child Psychology* April 2000

[36] Jim Hopkins and Michelle Kessler 'Companies Squander Billions on Tech' *USA Today* May 20 2002

enmeshed, emotional and storied existence into the depleted environment of the knowledge management system. Why would we want less than we have already? Our history has shown us that we always want more.

If knowledge management and e-learning truly want to serve our knowledge and learning needs in this rich, complex and shifting environment, it will have to start watching how we really do things.

There are six basic principles that seem to work consistently in our knowledge and learning habits; principles that knowledge management and e-learning technologies need to serve.

1. *Highly effective knowledge performers prefer knowledge fragments and lumps to highly engineered knowledge parts.*

Think about how we create new knowledge artefacts. If it's a Powerpoint presentation, we'll usually search for other stuff we have done in this field first. Then we might look for what other people have done. We'll lift slides and tinker with them or adapt them to fit, but we'll also lift text extracts, illustrations and figures. Each time we use this archived material, we'll be doing something different with it, and we want the freedom to use fragments. And then we'll put our own story, theme and wrap into it – integrating it with what we know to be true for this audience, and hopefully gearing it towards their concerns.

2. *Parts need to talk to their neighbours.*

In creating new knowledge artefacts, we do a lot of smoothing and linking, to make all the pieces fit together and make a coherent story. Much of what we do in knowledge work is about filling in the links between the chain of core nuggets of truth or interpretation that we want to deliver.

Even well-designed physical environments exhibit such smoothing, patches that help human beings interpret how the environment works and what it's for. The word "Push" on a swing door tells you which way it swings if it's not immediately apparent. If you go to the most well designed government department serving its citizens, you'll find handwritten notices propped up permanently where the design of the environment has not been clear enough to prevent frequent mistakes. If the "Push" sign on the swing door is handwritten, you can be sure that several people bumped their noses to warrant its presence. The association between neighbouring fragments of knowledge and between the fragments and their application often needs to be made apparent if your artefact is to be coherent. This smoothing will vary according to who is using the artefact, when, and where.

3. *The whole is more important than the parts.*

The story is much more critical than its components, the big picture determines how the small fragments are pieced together, sequenced or discarded. And the ability to grasp and describe the big picture quickly will expedite the ability to start gathering any available raw knowledge fragments and pulling them into the right shape for your purpose. Any

difficulty in discerning the big picture will mean a more laborious process of play, arranging and rearranging our knowledge fragments into different patterns to see which ones make sense. Coming at our knowledge artefact from the direction of the big picture allows us to vary our strategies for designing it: we can be minimalist or we can be lush and extravagant. The context of need for our knowledge artefact may also determine the degree of complexity and subtlety with which we render the story we want to convey.

4. *Knowledge artefacts provide just enough to allow the user to get started in the real world*

The creation and delivery of a knowledge artefact – whether physical or virtual – is simply an input to a larger knowledge process over which we have no control, but which we should seek to enable. Our artefact needs to be accessible to that process, and to assist it, meaning that interpretive barriers need to be as minimal as possible. And we need to be aware that only fragments of what we create might be used.

Knowledge is deployed and needed in many very different environments, not all of them tethered to personal computers and networks. The knowledge artefacts we create need to have affordances that work wherever they are needed. If it's a handwritten list of ingredients and it works, it works. There's no golden rule that it has to be digital. The genre and medium of the artefact are chosen according to how they can help people do knowledge work in their real world environments.

And when we resource knowledge management and learning, our resourcing apportionment needs to recognise the partial nature of content creation and delivery in the entire knowledge lifecycle. Other critical portions of the knowledge life cycle need resourcing just as much. Knowledge literacy – the competencies and skills that allow a person to recognise useful fragments and artefacts and apply them – is one portion. The affordances of the user's knowledge environment – the availability or lack of technology, the deployment of time, customary habits of knowledge use – is another. The way in which decisions are made and enacted socially, and habits and behaviours changed, is another. The way in which experience is built, enriched and utilised is another.

5. *Learning needs change faster than learning design*

Learning and knowledge acquisition are driven by the gap between what we know and can do, and what we need to know and do. If the world never changed, or changed very slowly, our learning and knowledge appetites would gradually reach a status quo. But the human race is an individuated, conflictual, constantly regenerating species. Our perspectives are always partial, we constantly change the rules of the games we play, and we interact with an environment whose capacity to surprise us exceeds our capacity to learn from it.

So learning and knowledge design are coping strategies that attempt to follow the world more often and less perfectly than they truly describe the status quo. Learning and knowledge will always be transitional, partial and contingent. It's better to try to make a halfway decent job of a contingent knowledge priority and move on to the next one, than to make a perfect job of an illusion that doesn't exist.

6. Variety is the spice of life

We have developed, and we maintain at great expense, an incredible variety of knowledge vehicles and artefacts, and the infrastructure that houses and sustains them. We started with alphabets and numbers, and moved to inscriptions and portable texts. We developed a multiplicity of highly nuanced genres, and experimented with visual representations in different media and on different scales. We concocted a chemistry of meanings around scents, sounds, architectural spaces, tastes and rituals. We developed technologies and intricate social systems to validate and authorise the products of those technologies.

Now we struggle with those very same issues in the new digital domain. Our engrossment with this domain is more a symptom of our anxiety about how to use it well, than it is an expression of intent to ditch our inherited knowledge systems. We like variety and disorder. We desire variety and disorder. We learn from variety and disorder. We create variety and disorder, because our constructed disorder is what heightens our ability to adapt, to create higher meanings, and higher senses of order, out of our too unpredictable world.

Jane Jacobs explores the relationship between order and disorder, variety and sameness in her brilliant book *The Death and Life of the Great American Cities*.

> Under the seeming disorder of the old city, wherever the old city is working successfully, is a marvelous order for maintaining the safety of the streets and the freedom of the city. It is a complex order. Its essence is the intimacy of sidewalk use, bringing with it a constant succession of eyes. This order is all composed of movement and change, and although it is life, not art, we may fancifully call it the art form of the city and liken it to the dance – not to a simple-minded precision dance with everyone kicking up at the same time, twirling in unison and bowing en masse, but to an intricate ballet in which the individual dancers and ensembles all have distinctive parts which miraculously reinforce each other and compose an orderly whole.[37]

Human beings learn "through a constant succession of eyes". We lay down pattern after pattern of perception based on samenesses and differences in our observed experience, and it is through the lens of this layered history of perceptions that we form judgements and manufacture prisms that work for us now. We have created variety and complexity in our knowledge artefacts because only this variety and complexity allows us to recognise, interpret and adapt to complex and different circumstances.

We have also created remarkable icons of simplicity in our knowledge artefacts, from the pure tone of a soprano singing a high C, to the elegance of a Japanese haiku. But such simplicity only works where it speaks against the complexity that surrounds it, and out of that dialogue, wrings for us a temporary note of recognition, resonance or truth.

Simplicity is a necessary and useful tool, but its value lies in recognising its limits. Autism on the other hand is a kind of pathological simplification of our environment.

---

[37] Jane Jacobs *The Death and Life of the Great American Cities* (New York: Vintage, 1961) p.50

When it comes to knowledge, that is a dangerous condition to be in. We don't need reusable knowledge objects. We need disposable ones.

Patrick Lambe is Founder of Straits Knowledge (www.straitsknowledge.com), a Singapore-based consulting and research firm specialising in knowledge management, e-learning and innovation. His articles can be found at www.greenchameleon.com and he is the author of The Blind Tour Guide: Surviving and Prospering in the New Economy. He is leading a two day KM workshop for the Singapore Institute of Management 26-27 August 2002.