Text-made Text

BILL COPE
Centre for Workplace Communuication and Culture, Melbourne, Australia
MARY KALANTZIS
RMIT University, Melbourne, Australia

ABSTRACT What is the nature of the change represented by digital communications technologies? How will the impact of the digital compare with the massive changes spawned in its time by print and books? These are the two key questions addressed in this article. The authors answer the first of these questions by comparing the emergence of the printed book with the emergence of the digital communications technologies. The next section of this article, 'Transformations in Ways of Meaning: the case of print', discusses the technological nature and textual consequences of the printing press. The following section, 'Transformations in Ways of Meaning: designing text digitally', does the same for digital connectivity. Several themes emerge.

Changing the Means of Production of Meaning

In December 1995, 16 million people were connected to the Internet. By December 1997, the figure had risen to 101 million; by December 1999, 201 million; by September 2001, 516 million; and by September 2002, 606 million. At the end of 2001, 29 million Chinese citizens were connected to the Internet; by June 2002, this figure had risen to 46 million. By the end of 2002, there were 275,000 Internet users in Uzbekistan, compared to 137,000 a year previously (Nua.com, 2003). Within a decade of its practical availability, 10% of the world's population had become connected to the Internet. The pace of growth continues, regardless of the burst of the 'dotcom' bubble in March 2000. Several years later, technology and telecommunications companies were still in the doldrums, a case of commercial realities bearing little direct resemblance to the social realities of use.

Instantaneous, global, free access to information – this is the revolutionary promise of the World Wide Web. And in copious quantities – as we were writing this paragraph, *Google* reported that its index covered

3,307,998,701 web pages. This massive corpus of published content takes a form which is radically new in a number of respects. Its platform is genuinely multimedia – a *mélange* of text, still and moving image and audio – all transmitted within the unifying framework of hypertext markup language and its successors. In some moments, the Web seems a bit like a book, in others like a movie, in still others a directory or index. It is all and none of these things.

The Web is also a new way of reading and writing. In direct contrast to the world of print, hypertext links create a radically non-linear reading environment. They require that readers take a newly active role in relation to text – to the extent that they might more accurately be called 'users'. This is a navigational role in which they actively build meanings from a range of hypertextual possibilities, rather than following the reading order as traditionally determined by the author. It is also highly decentralised, allowing communities of common interest to develop their own publishing places, and to be themselves in their own peculiar ways. Dematerialisation, the virtual, an economy that produces and circulates knowledge more than it does physical things ... these are just a few items in a growing list of startling social predictions (Mitchell, 1995; Gilster, 1997). Bleak views of these processes verge on the apocalyptic. They speak of the growing digital divide, wild zones, the bandwidth disadvantaged, electronic dazzlement, sedentary man, grey ecologies and a civilisation of forgetting (Lash & Urry, 1994; Harvey, 1996; Virilio, 1997).

And then, the pundits say, the book is dead.

Johann Gutenberg created Europe's first printing press in 1450. Within just 50 years, print shops were to be found in every major city and town in Europe, 1500 in all. Some eight million volumes had been printed, consisting of 23,000 titles (Eisenstein, 1979). Textual and mental devices came with this new apparatus: analytical contents, alphabetical ordering, standardised spelling, indexing, bibliography, citation of sources, identification of individual authorship, declaration of intellectual property rights and perhaps, most importantly, the key that holds so much of the apparatus together, pages numbered identically from copy to copy. Modern science, modern education and modern consciousness followed – empirical, analytical, rational – for better and for worse. Whereas the oral life-world of non-literate societies had been by temperament additive, aggregative, redundant, repetitive, conservative, situational, empathetic, participatory and concrete, the life-world of literacy was analytical, subordinative, linear, experimentative, abstracting, objective, conceptual and distancing (Ong, 1982). This was the beginning of the world of the book.

Half a millennium later, this world faces a challenge from the emerging order of digital connectivity. What is the nature of the change represented by digital communications technologies? How will the impact of the digital compare with the massive changes spawned in its time by print and books? These are the two key questions we address in this article.

We answer the first of these questions by comparing the emergence of the printed book with the emergence of the digital communications technologies. The next section of this article, 'Transformations in Ways of Meaning: the case of print', discusses the technological nature and textual consequences of the printing press. The following section, 'Transformations in Ways of Meaning: designing text digitally', does the same for digital connectivity. Several themes emerge.

First, a significant change is emerging in the mechanics of rendering, or the task of making mechanically reproduced textual meaning. In the world of print, text was 'marked up' for a single rendering. In the emerging world of the digital, text is increasingly marked up for structure and semantics (its 'meaning functions'), and this allows for alternative renderings (various 'meaning forms', such as print, web page or audio).

Second, linguistic, visual and audio meanings are constructed of the same stuff in the digital environment. This contrasts with the world of print, in which the physical production and rendering processes for the visual and the linguistic were more conveniently separated, and audio production and later reproduction unequivocally separated. The consequence in the digital era is a trend to multimodality, to the fabrication and distribution of texts which integrate linguistic, audio and visual modes of meaning.

Third, whereas the trend in the era of print was towards large, homogeneous speech communities and monolingual nationalism, the trend in the era of the digital may well be towards multilingualism and divergent speech communities which distinguish themselves by their peculiar manners of speech and writing – as defined, for instance, by technical domain, professional interest, cultural aspiration or sub-cultural fetish.

None of these changes is technologically driven, or at least not simplistically so. It is not until 30 years into the history of the digitisation of text that clear signs of shape and possible consequences of these changes begin to emerge.

The pivot point of this article is an account of the emergence of new textual practices several decades into the digital era. In contrast to the Gutenberg era, text forms are rendered on the basis of the semantic and structural markup of text functions. In this sense, text is used to make text. This peculiar shift in textual practice is essentially semiotic – in other words, a shift in the practices of representation and communication. This shift is not caused by the invention of the digital; rather, it exploits semiotic possibilities opened up by digitisation.

If this is what is happening with digital communications technologies, what then do we do about it? This is our second key question. And this is where our analysis of the phenomenon of text-made text turns into a program of action.

The section 'Digital Schemas for the Book' discusses a space in which a legacy information architecture finds a revealing new place in the era of digital connectivity. A new life is being breathed into books through new rendering

technologies, systems for digital resource discovery, electronic library cataloguing frameworks, e-learning systems, ecommerce protocols and digital rights management processes.

The following section introduces Common Ground Markup Language (CGML), a functional schema for authorship and publishing. CGML attempts to create a sound basis for interoperability between the schemas described in the previous section, incorporating their varied functions. CGML is an 'interlanguage'. Its concepts constitute a paradigm for authorship and publishing, drawing on an historically familiar semantics, but adapting this to the possibilities of the Internet. Its key devices are thesaurus (mapping against functional schemas) and dictionary (specifying a common-ground semantics). These are the semantic components for narrative structures of text creation, or the retrospective stories that can be told of the way in which authors, publishers, referees, reviewers, editors and the like construct and validate text. The purpose of this section of the article is both highly pragmatic (a description of an attempt to create a kind of functional grammar of the book), and highly theoretical (a theory of meaning function capable of assisting in the partially automated construction and publication of variable meaning forms).

The next section returns to the point with which we began, to the question of the nature and significance of the changes to practices of representation and communication in the digital era. Some changes, we suggest, may be less important than frequently assumed – hypertext, and the 'virtual', for instance, are founded on practices and phenomena entirely familiar to the half-millennium-long tradition of the printed book. Other changes are indeed significant, and perhaps less obviously so. Amongst these, the article has identified three as particularly significant: changes in the mechanics of rendering; an integration of the modalities of meaning; and an emerging polylingualism. These changes have the potential to breathe new life into an old medium, to reinvigorate and extend the reach of the book.

The book is dead. Long live the book.

Between the positive and negative scenarios of the digital age, we take a neutral stance. For history always remains open and contingent. Humans may make collective choices of meaning and practice which have all manner of potentials. Amongst these potentials, we have opted for a kind of strategic optimism, not daring to foretell the future, but nevertheless anticipating with some hope what we could possibly achieve as we refashion our means of production of signs, our tools for representation and communication.

By way of brief background, this article has developed at the confluence of three research endeavours. The first has been research over a 15-year period into the changing communications environment, and the consequences of these changes for literacy pedagogy (Cope & Kalantzis, 1993; New London Group, 1996). The second is a substantial new body of research undertaken by Common Ground in association with RMIT University for the Australian Department of Industry in 2000-2003, 'Creator to Consumer in a Digital Age: book production in transition'. This research has culminated in the production

of a series of 10 research reports in book format examining changing technologies, markets and human skills in the publishing supply chain (C-2-C Project, 2003). The third research endeavour is the Common Ground Markup Language project commenced by Common Ground in 2000 with the support of an AusIndustry Research and Development grant. If this article speaks strangely at times, it is because it speaks in three, sometimes incongruent, voices. And if its agenda seems peculiar, it is because it crosses backwards and forwards between historical reflection and a call to action.

Transformations in Ways of Meaning: the case of print

Of world defining events, historian Fernand Braudel says the invention of the compass may have been more important than the invention of the printing press. Here he is speaking of the global impact of a particular kind of European modernity, and some of the key inventions that helped Europeans establish an era of economic and perhaps also cultural dominance over the world. Interpreting this same European modernity, many say that the invention of the printing press was unequivocally the world defining moment. Huge claims are made for the significance of print – as the basis of mass literacy and modern education, as the foundation of modern knowledge systems, and even for creating a modern consciousness aware of distant places and inner processes not visible to the naked eye of everyday experience.

The consequence of print was a new way of representing the world. Contents pages and indexes ordered textual and visual content analytically. A tradition of bibliography and citation arose in which a distinction was made between the author's voice and ideas and the voice and ideas of other authors. Copyright and intellectual property were invented. And the widely used modern written languages we know today rose to dominance and stabilised, along with their standardised spellings and alphabetically ordered dictionaries, displacing a myriad of small spoken languages and local dialects.

The cultural impact was enormous: modern education and mass literacy; the rationalism of scientific knowledge; the idea that there could be factual knowledge of the social and historical world; the nation-state of interchangeable individuals; the persona of the creative individual author. All these are in part a consequence of the rise of book culture, and give modern consciousness much of its characteristic shape.

What was the defining moment? Was it Johann Gutenberg's invention of moveable type in 1450 in Mainz, Germany? This is what the Eurocentric version of the story tells us. Or is this giving too much credit to western modernity? By another reckoning, the defining moment may have been the invention of clay moveable type in China by Li Sheng in about 1040, or of wooden moveable type by Wang Zhen in 1297, or of bronze moveable type by officials of the Song Dynasty in 1341 (Luo, 1998). Or, if one goes back further, perhaps the decisive moments were the Chinese inventions of paper in the year

105, wood block printing in the late sixth century and book binding in about 1000.

Whoever was first, it is hard to deny that the Gutenberg invention had an immediate, world defining impact, whilst the Chinese invention of moveable type remained localised and even in China was little used. As mentioned earlier, within 50 years of the invention of the Gutenberg Press there were 1500 print shops, using the moveable type technology, located in every sizeable town in Europe; and 23,000 titles had been produced totalling eight million volumes (Eisenstein, 1979). The social, economic and cultural impact of such a transformation cannot be underestimated.

We will focus for a moment on just three aspects of this transformation, and these are the three themes of this article: the processes for the rendering of meaning form; the pre-eminence of the modality of language; and the rise of large communities of linguistic commonality. In this analysis, our focus will be on the tools for representation, on the means of production of meaning forms. For the printing press marks the beginning of the age of mechanical reproduction of meaning. This is new to human history and, however subtly, this development changes in some important respects the nature of the very process of signing.

The Mechanics of Text Rendering

Gutenberg had been a jeweller, and a key element of his invention was the application of his jeweller's skills to the development of cast alloy moveable type. Here is the essence of the invention: a single character is carved into the end of a punch; the punch is then hammered into a flat piece of softer metal, a matrix, leaving an impression of the letter; a hand-held mould is then clamped over the matrix, and a molten alloy poured into it. The result is a tiny block of metal, a 'type', on the end of which is a reproduction of the character that had been carved into the punch. Experienced type founders could make several hundred types per hour. The type was then set into 'formes' – blocks in which characters were lined up in rows, each block making up a page of text – assembled character by character, word by word, line by line. Finally, Gutenberg applied the technology of the wine press to the process of printing; the inked forme was clamped against a sheet of paper, the pressure making an impression on the page. That impression was a page of manufactured writing (Man, 2002).

The genius of the invention was that the type could be reused. After printing, the formes could be taken apart and the type used again. But the type was also ephemeral. When the relatively soft metal wore out, it could easily be melted down and new types moulded. One set of character punches was all that was needed to make hitherto unimaginable numbers of written words: the tiny number of punches producing as many matrices as were needed to mould a huge number of types in turn producing seemingly endless impressions on paper. In fact, Gutenberg invented in one stroke one of the fundamental

principles of modern manufacturing: modularisation. He had found a simple solution to the conundrum of mass and complexity. Scribing a handwritten book, or making a hand-carved woodblock (which, incidentally, remained the practice in China, well after Gutenberg and despite the earlier invention of moveable type by the Chinese), is an immensely complex, not to mention time-consuming, task. By reducing the key manufacturing process (making the types) to the modular unit of the character, Gutenberg introduced a process simplicity through which complex texts could easily and economically be manufactured.

But he also produced new means of meaning. Its significance can be evaluated on several measures: technology of the sign, epistemology of the sign and social construction of the sign.

The technology of the sign was centred around the new science and craft of 'typography'. The practices of typography came to rest on a language of textual design centred upon the modular unit of the character, not only describing its size (in points), visual design (fonts) and expressive character (bold, italic, book). It also described in detail the spatial arrangements of characters and words and pages (leading, kerning, justification, blank space) and the visual information architecture of the text (chapter and paragraph breaks, orders of heading, running heads, footnotes). Typography also established a newly sophisticated and increasingly complex and systematised science of punctuation marks – a series of visual markers of textual structure with no direct equivalent in speech. The discourse of typography described visual conventions for textual meaning in which meaning function was not only expressed through semantic and syntactic meaning forms but also through a new series of visual meaning forms designed to realise additional or complementary meaning functions.

The new regime of typography also added a layer of abstraction to the abstraction already inherent in language. Representation of the world with spoken words is more abstract than representation through images expressing resemblance, and this by virtue of the arbitrary conventions of language. Writing, in turn, is more abstract than oral language. Vygotsky says that writing is 'more abstract, more intellectualised, further removed from immediate needs' than oral language and requires a 'deliberate semantics', explicit about context and self-conscious of the conditions of its creation as meaning (Vygotsky, 1962). With printing, further distance was added between expression and rendering, already separated by writing itself. The Gutenberg invention added to the inherent abstractness of writing. This took the form of a technical discourse of text creation which was, in a practical sense, so abstracted from everyday reality that it became the domain of specialist professionals: the words on the page, behind which is the impression of the forme on the paper, behind which is the science of typography, behind which is the technology of type, a matrix and a punch. Only when one reaches the punch does one discover the atomic element upon which the whole edifice is

founded. The rest follows, based on the principles of modularisation and the manufacturing logic of recomposition and replication.

The social construction of the sign also occurs in new ways, and this is partly a consequence of these levels of scientific and manufacturing abstraction. With this comes the practical need to remove to the domain of specialists the processes, as well as the underlying understanding of the processes, for the manufacture of the sign. Speech and handwriting are, by comparison, relatively unmediated. To be well formed and to take on the aura of authoritativeness, printed text is constructed through a highly mediated chain of specialists, moving from author, to editor, to typesetter, to printer, and from there to bookstores and readers. Each mediation involves considerable backwards and forwards negotiation (such as drafting, refereeing, editing, proofing and markup) between author, editor, typesetter and printer. The author only appears to be the font of meaning; in fact texts with the aura of authority are more socially constructed than the apparently less authored texts of speech and handwriting. What happens in this process, however, is more than just a process of increasing social construction based on negotiation between professionals. A shift also occurs in the locus of control of those meanings that are ascribed social significance and power. Those who communicate authoritatively are those linked into the ownership and control of the means of production of meaning – those who control the plant that manufactures meaning and the social relations of production of meaning.

The Modality of Language

Meanings are made in a number of modes; there is more than one means of production of meaning. As well as linguistically, meanings are made in the visual, audio, gestural and spatial modes (Cope & Kalantzis, 2000a, b). In the Christian religion of medieval Europe, faith was acquired through the visual imagery of icons, the audio references of chant, the gestural presence of priests in sacred garb, and the spatial relationships of priest and supplicant within the architectonic frame of the church. The linguistic was deliberately backgrounded, an audio presence more than an linguistic one insofar as the language of liturgy was unintelligible to congregations (being in Latin rather than the vernacular), and the written forms of the sacred texts inaccessible in any language to an illiterate population (Wacquet, 2001). Gutenberg's first book was the Bible, and within a century the Reformation was in full swing. The Reformation sought to replace the Latin Bible with vernacular translations. Its agenda was to create congregations whose engagement with their faith was primarily linguistic – the sacred Word, for a people of the Book. The underlying assumptions about the nature of religious engagement were so radically transformed that it might as well have been a new religion.

So began a half-millennium-long (modern, western) obsession with the power and authoritativeness of language, and particularly its written form, over other modes of meaning. One of the roots of this obsession was a practical

consequence of the new means of production of meaning in which the elementary modular unit was the character. In a very pragmatic sense, until photoengraving and offset printing, it was about as hard to make images as it had been to hand-make whole blocks before the invention of moveable type. Furthermore, with the rise of moveable type letterpress it was difficult, although not impossible, to put images on the same page as text – to set an image in the same forme as a block of text so that the impression of both text and image was as clear and even as it would have been had they been set into different formes. Until offset printing if there were any 'plates' they were mostly printed in separate sections for the sake of convenience. If image was not removed entirely, text was separated from image (Cope, 2001a, b, c). And so began a radical shift from image culture to the word culture of western modernity (Kress, 2000). This was taken so far as to entail eventually the violent removal of images, the iconoclasm of Protestantism, which set out to remove the graven images of Catholicism in order to mark the transition to a religion based on personal encounter with the Word of God. This Word was now translated into vernacular languages, mass-produced as print and distributed to an increasingly literate populace. Gutenberg's modularisation of meaning to the written character was one of the things that made the western world a word-driven place, or at the very least, made our fetish for writing and word-centeredness practicable.

Enlarged Communities of Linguistic Commonality

If one radically new aspect of Gutenberg's invention was to position the character as the modular unit in the manufacture of text, the other was to create the manufacturing conditions for the mass production of texts. This is another respect in which his invention anticipated, and even helped to usher in, one of the fundamentals of the modern world. Indeed, the idea of modularisation and the idea of mass production were integrally linked. The component parts of text were mass-manufactured types assembled into formes. The atomic element of modularisation may have been reduced and rationalised; the process of assembly was nevertheless labour-intensive. One printed book was far more expensive to produce than one scribed book. The economy of the printed book was one of scale, in which the high cost of set-up is divided by the number of impressions. The longer the print run, the lower the per unit cost, the cheaper the final product could be sold and the better the margin that could be added to the sale price. Gutenberg printed only about 200 copies of his Gutenberg Bible, and for his trouble he went broke. He had worked out the fundamentals of the technology but not its commercial fundamentals. As it turned out, in the first centuries after the invention of the moveable type letterpress, books were printed in runs of about a thousand. We can assume that this was approximately the point where there was sufficient return on the cost of set-up, despite their lower sale price than scribed books.

Here begins a peculiarly modern manufacturing logic – the logic of mass, or the logic of increasing economies of scale.

Moreover, when culture and language are being manufactured, the assumption is one of scale in discourse communities, that there are enough people who can read and will purchase a particular text to justify its mass production. At first, economies of scale were achieved by publishing in the lingua franca of the early modern European cultural and religious intelligentsia, Latin. This meant that the market for a book was the whole of Europe, rather than the speakers of a local vernacular. By the seventeenth century, however, more and more material was being published in local vernaculars (Wacquet, 2001). With the expansion of vernacular literacy, local markets grew to the point where they were viable. Given the multilingual realities of Europe, however, not all markets were of a sufficiently large scale. There were many small languages that could not support a print literature; there were also significant dialect differences within languages that the manufacturing logic of print ignored. Driven by economies of scale, the phenomenon that Anderson calls 'print capitalism' set about a process of linguistic and cultural standardisation of vernaculars, mostly based on the emerging metrics of the nation-state – marginalising small languages to the point where they became unviable in the modern world and standardising written national languages to official or high forms (Anderson, 1983). This process of standardisation had to be rigorous and consistent, extending so far as the spelling of words – never such a large issue before – but essential in a world where text was shaped around analytical apparatuses such as alphabetical indexing. With its inexorable trend to linguistic homogenisation and standardisation, it was this print capitalism which ushered in the modern nation-state, premised as it was on cultural and linguistic commonality. And so 'correct' forms of national languages were taught in schools; newspapers and an emerging national literature spoke to a new civically defined public; and government communications were produced in 'official' or 'standard' forms. As a consequence, a trend to mass culture accompanied the rise of mass manufacture of printed text, and pressure towards linguistic homogenisation became integral to the modernising logic of the nation-state.

However rational from an economic and political point of view – realising mass literacy, providing access to a wider domain of knowledge, creating a modern democracy whose inner workings were 'readable' – there have also been substantial losses as a consequence of these peculiarly modern processes. Phillipson documents the process of linguistic imperialism, in which the teaching of literate forms of colonial and national languages does enormous damage to most of the ancestral and primarily oral languages of the world, as well as to their cultures (Phillipson, 1992). Mühlhäusler traces the destruction of language ecologies – not just languages but the conditions that make these languages viable – by what he calls 'killer languages' (Mühlhäusler, 1996).

And now, in the era of globalisation, it seems that English could have a similar effect on the whole globe to that which national languages had in their

day on small languages and dialects within the territorial domains of nation-states. By virtue in part of its massive dominance of the world of writing and international communications, English is becoming a world language, a lingua mundi, as well as a common language, a lingua franca, of global communications and commerce. With this comes a corresponding decline in the world's language diversity. At the current rate, between 60 and 90% of the world's 6000 languages will disappear by the end of this century (Cope & Gollings, 2001). To take one continent, of the estimated 250 existing in Australia in the late eighteenth century, two centuries later there are only 70 left possessing more than 50 speakers; perhaps only a dozen languages will survive another generation; and even those that survive will become more and more influenced by English and interconnected with Kriol (Dixon, 1980; Cope, 1998).

Transformations in Ways of Meaning: designing text digitally

In each of the three areas we have been discussing, the emergence of digital means of manufacture and distribution of meaning portends a significant shift away from the technical, commercial and cultural logic surrounding the Gutenberg press. The focal point of typography on visual rendering of meaning is being replaced by automated processes for rendering based on a description of meaning function. Meaning is increasingly multimodal, in sharp contrast with the separation of the printed word and image characteristic of the Gutenberg era. And there may be more space for polylingualism, possibly even to the extent of reversing the trend to linguistic homogenisation inherent to the technologies and economies of mass production, with their accompanying politics of nationalism and the pressure to be absorbed into mass culture.

Automated Rendering of Form based on Description of Function

It is not digitisation per se that created the shifts which are the focus of this article. Widespread use of digitisation in the manufacture of text dates from the mid-1970s, initially in the form of phototypesetting and soon after in domestic and broad commercial applications with the introduction of word processing and desktop publishing systems. However, it was not until the late 1990s that we witnessed the beginnings of a significant shift in the mechanics of digital representation. Until then, text structure was defined by the selfsame discourse of 'markup' that typesetters had used for 500 years. Historically, the role of the typesetter had been to devise appropriate visual renderings of meaning form on the basis of what they were able to impute from the meaning functions underlying the author's text. Markup was a system of manuscript annotation, or literally 'marking up' as a guide to the setting of the type. This is where the highly specialised discourse of typography developed.

The only transition of note in the initial decades of digitisation was the spread of this discourse into the everyday, non-professional domains of writing

(the word processor) and page layout (desktop publishing software). Never before had non-professionals needed to use terminology such as 'font', 'point size' or even 'leading' – the spaces between the lines which had formerly been blocked out with lead in the formes of letterpress printing. This terminology had been an exclusive part of the arcane, professional repertoire of typesetters. Now the author had to learn to be their own typesetter, and to do this they had to learn some of the language of typography. As significant as this shift was, the discourse of typography remained essentially unchanged, as well as the underlying relations of meaning function to meaning form.

The Gutenberg discourse of typography even survived the first iteration of Hypertext Markup Language or HTML, the engine of the World Wide Web (Berners-Lee, 1990). This was a vastly cut-down version of Standardised General Markup Language (SGML), which had originated in the IBM laboratories in the early 1970s as a framework for the documentation of technical text, such as computer manuals (Goldfarb, 1990). Particularly in its earliest versions in the first half of the 1990s, HTML was driven by a number of presentationally oriented 'tags', and in this regard was really not much more than a traditional markup language, albeit one designed specifically for web browsers. When markup tags such as , <italics> or <underlining> are placed before a word or a phrase, the tag effectively manufactures or renders that word or phrase in that visual form in the moment it is read by a web browser. Its genius was that it was based on just a few tags (about 100) drawn from SGML, and it was easy to use, accessible and free. Very soon it was to become the universal language of the World Wide Web.

The truly significant shift away from the world of Gutenberg commenced in the second half of the 1990s and is epitomised by the rapid emergence since 1998 of XML or Extensible Markup Language. This, too, is a vastly simplified version of SGML, albeit a simplification more rigorously true to the original SGML insight into the benefits of separating information architecture from rendering. In this respect, its conceptual bases are fundamentally different from those of HTML. Its main distinguishing feature is that it marks up for structure and semantics, instead of for presentation. This represents a truly revolutionary shift away from the practices of text manufacture that had predominated for the previous half-millennium. Within just a few years, XML has become pervasive, if mostly invisible – not only in the production of text but also increasingly in computer games, electronic commerce, mobile phones and the Internet. It may be decades before the depth of its impact is fully realised.

Unlike HTML, which is a markup language, XML is a framework for the construction of markup languages. It is a meta-markup language. The common 'M' in the two acronyms is deceptive, because marking up markup is an activity of a very different order to marking up. XML is a simplified syntax for the construction of any markup language. In this sense, XML is also a significant departure from SGML. Consequently, XML is a space where a plethora of markup languages is appearing, including markup languages that work for the

activities of authorship and publishing – the subject of the next section of this article.

Most significantly, however, XML brackets the logistics of presentation away from the abstract descriptive language of structure and semantics. Take the heading of this section of this article, for instance. In traditional markup, we might use the typographer's analytical tool to mark this as a heading, and that (for instance) might be to apply the visual definition or command such as <12pointTimesBold>. We could enact this command by a variety of means, and one of these is the screen-based tool of word processing or desktop publishing. The command is invisible to the reader; but somebody (the author or the typesetter) had to translate the meaning function 'heading' into a visual meaning form which adequately represents that function to a community capable of reading that visual rendering for its conventionally evolved meaning. If we were to mark up by structure and semantics, however, we might use the concept <heading>. Then, when required, and in a rigorously separated transformational space, this tag is translated through a 'stylesheet' into its final rendered form – and this can vary according to the context in which this information is to be rendered: as HTML by means of a web browser, as conventional print (where it may happen to come out as 12-point Times Bold, but could equally come out as 16-point Garamond Italic, depending on the stylesheet), to a mobile phone or electronic reading device, or as synthesised voice, or even as text or audio in translation to another language.

The semantic part of XML is this: a tag defines the semantic content of what it marks up. For instance, 'cope' is an old English word for a priest's cloak; it is a state of mind; and it is the surname of one of the authors of this text. Semantic ambiguity is reduced by marking up 'Cope' as <surname>. Combined with a tag that identifies this surname as that of an <author> of this text, a particular rendering effect is created depending on the transformation effected by the stylesheet that has been applied. As a consequence, in the manufacturing process this author's surname is rendered in the place where you would expect it to appear and in a way you could expect it to look – as a byline to the title.

And the structural part of XML is this: a framework of tags (a Document Type Definition or DTD) provides an account of how a particular domain of meaning hangs together, in other words how its core conceptual elements, its tags, fit together as a relatively precise and interconnected set of structural relations. A domain of meaning consists of a system of interrelated concepts. For instance, <Person>'s name may consist of <GivenNames> and <Surname>. These three concepts define each other quite precisely. In fact, these relations can be represented taxonomically – <GivenNames> and <Surname> are 'children' of the 'parent' concept <Person>.

Put simply, and applying now the terminological frame of reference being developed in this article, the effect of markup practices based in the first instance on semantics and structure rather than presentation, is that meaning

form is rigorously separated from meaning function. Digital technologies ('transformation stylesheets') automate the manufacture of form based on their peculiar framework for translating generalised function into the particularities of conventional understandings and readings of form. When 'Cope' is marked up as <Surname> and this is also a component of <author>, (the meaning function) then the stylesheet transformation will interpret this to mean that the word should be located in a particular point size at a particular place on the title page of a book (the meaning form). Meaning and information architecture are defined functionally. Then form follows function.

In the meantime, HTML is becoming more like XML. In recent versions of HTML, earlier presentational tags have been deprecated. And a fully XML-compliant version of HTML is gaining increasing acceptance in the form of XHTML. At the same time, word processing programs are moving in the general direction of XML. Not only does this include the capacity to 'save as' XML but the rudiments of the principles of structural and semantic markup are also to be found in the increasing use of 'styles' and 'templates' in order to define text structure and metadata tags attached to a document in order to define file content.

These trends are very much influenced by the emergence of multiple rendering options for textual (and visual and audio) meaning. The Gutenberg-inspired book was singular in its expression of meaning function in one rendering of meaning form – the printed book. The markup used to define its information architecture could serve one purpose only. It translated meaning function into just one meaning form. Today, the same text can also be read through a web browser as HTML, on various handheld reading devices using the Open eBook format, in the print-facsimile Portable Document Format, as Braille, or as synthesised speech by means of Digital Talking Book (Cope, 2001a, b, c). Important back-to-the future rendering formats are also in development, such as eInk and Smart Paper – flexible electronic substrates which can be read with reflected light and which attempt to replicate the traditional paper-based reading experience, but without the costs and environmental impacts of printing to paper (Coburn et al, 2001). The lesson of previous waves of innovation in other areas of the culture-media industry is that multiple formats (such as cinema, television, video/DVD) create new kinds of audience experience, provide broader access and extend market share. The same is destined to hold true for textual content, and XML and similar technologies designed for alternative renderings and multichannel distribution will be fundamental to the next phase of development of the business of authorship and publishing. Travellers will be able to travel guidebook in their bag, or to download a specific piece of local information to a Portable Digital Assistant (PDA) or a mobile telephone handset (about the monument they may be standing in front of, or the location of the nearest hotel) or they could listen to the text while driving their car or through a Walkman as they walk along the street.

From this highly synoptic account of shifts in digital text creation and rendering technologies, we would like to foreground one central idea, and that is the emergence of a new kind of means of production of meaning: text-made text.

Text-made text represents a new technique and new mechanics of the sign. Text is manufactured via an automated process in which text produces text. Text in the form of semantic and structural tags drives text rendering by means of stylesheet transformation. This is a technology for the automated manufacture of text from a self-reflective and abstracted running commentary on meaning function. It is a new mode of manufacture which, we contend, may well be destined to change the very dynamics of writing, and much more profoundly so than the accretion of typesetting language into everyday textual practices, which accompanied the first phase of digitisation in the last quarter of the twentieth century.

Accompanying this new mechanics is a new epistemology of the sign. It is to be expected that the emergence of new representational means – new relations of meaning form to meaning function – will also entail new ways of understanding meaning and new ways of thinking. Writing acquires yet another layer of 'conscious semantics', to use the Vygotskian phrase, and one which Vygotsky could not possibly have imagined when he was working in the first half of the twentieth century. More than putting pen to paper, or fingers to keyboard, this new mechanics requires the self-conscious application of systematically articulated information architectures. In effect, these information architectures are functional grammars, and the very manufacture of text will be driven by these functional grammars.

Critically, however, our historical evidence tells us that this epistemological shift is not technologically determined. It is not until a quarter of a century after the widespread digitisation of text construction, manufacture and distribution that the epistemological shift away from Gutenberg begins to take clear shape and gain wide acceptance. Digitisation itself provides new scope for meaning and broader meaning potentials. Take-up of these potentials is no more and no less than a purely semiotic process, and thus a simply human process, of self-transformation. The decisive shift is in the business of meaning making, a shift in the human practices of signing. As such, it is also a subtle extension of the ways we can be human.

The practical consequences of these developments could be enormous. For one thing, children in schools and people who aspire to work in the communications industries will need to do grammar again, or at least a grammar of sorts (Cope & Kalantzis, 2003). It will not be grammar in the traditional form of abstract word and sentence-level syntax – using concepts like 'noun', or 'verb', or 'adverbial clause of time' to describe written or verbal meaning. It will be a functional grammar that defines system and structure in text, and accounts for the complex and subtle ways these vary from text to text and genre to genre (Halliday, 1978; Hodge & Kress, 1988; Cope & Kalantzis, 1993). Literacy will no longer be the stuff of 'correct' expression involving

correctly spelled words and properly formed sentences adjudged to be 'good writing' by standardised tests. Instead, the focus of literacy learning will be the highly variable relationship of meaning functions to meaning forms.

The social relations of the sign are also destined to change. The most original and generative movements in contemporary technology are also politically motivated social movements (Williams, 2002). In the spirit of the Free Software and Open Source movements, XML is designed to prevent computerised content being locked up in proprietary code. Like HTML before it, but even more so, the effect of XML is to break the proprietary grip of closed commercial software systems, and thus to facilitate universal computer-to-computer communication independent of proprietary systems. It is free for anyone to use; its major tagging schemas are open source; and anybody can build a tagging schema which structures information relevant to their particular community of interest. XML also uses transparent natural language tags, so a stretch of computer code is literally readable. The distancing mystique of secret and obscurantist computer languages is regarded with disdain by the XML and open source coding communities. And when it comes to textual meaning, natural language structures natural language for the purpose of rendering its meaning functions in natural language. The temper of these new technologies is democratic and the motive of their proponents in the first instance is to serve human utility.

Even though modern economic fundamentals remain unchanged, and privately owned intellectual property is traded for the commercial benefit of those who can afford to take effective control of that intellectual property, the commercial effects of the emerging regime of the sign are against the grain of the economics of the Gutenberg era. The traditional print shop required expensive plant, a closed shop of waged tradespersons with specialist craft skills, and was premised on the economies of scale of mass production. The new technologies lower the capital and skills entry point, and economies of manufacturing scale are trending towards flat. In short, the new technologies provide cheap and transparent tools through which creators and communities of common interest can manage their own intellectual property. The era of Gutenberg favoured large and dominant languages, cultures and communities. The era of digitisation may in part reverse the hitherto distinctively modern trends to concentration of ownership, and the standardisation and homogenisation of knowledge and culture. This does not have to mean a loss of quality or the burgeoning of what is dismissively termed 'vanity publishing'. As we will argue later in this article, systematic social processes for the construction of text may be put in place which will enable communities of practice to publish for themselves, and these may produce publication decisions that are more reliable than those of mass production publishers and media conglomerates.

In fabricating textual meaning for multiple and divergent uses, this new era also provides broader access, as well as allowing more space for the reader to construct meanings in ways peculiar to their needs. Instead of a whole book,

for instance, text may be decomposed and recomposed according to an individual reader's needs. This, however, requires a reduction in the basic granular unit of text, from the book to information units as small as a paragraph or even a sentence, a subject to which we will return later in this article. Needless to say, even in the case of a digitally printed book, the reader will be able to make their own text (the poems, the recipes or the course pack they want), and in a format that suits their needs (such as regular or large print). Or they will be able to listen to the book as audio, or download the fragment they want to their own PDA, mobile phone or computer screen, or print it out on their own local printer. In practical terms, a more interventionist role for the reader provides greater access (from a social point of view) and opens a range of new niche markets (from a commercial point of view).

Multimodality

Reversing the language-centric tendencies of western meaning that emerged over the past half-millennium, a broad shift is occurring away from linguistic and towards increasingly visual modes of meaning (Kress & van Leeuwen, 1996; Kress, 2001). A key contributing factor to this has been the potential opened by digital technologies. And, once again, the technology itself does not determine the change; it merely opens human possibilities. The real shift is in human semiotic practices.

By contrast with the Gutenberg technology, it is remarkably easy to put the images and words together when digitally constructing meaning, and this is in part because text and images are built on the same elementary modular unit. The elementary unit of computer-rendered text is an abstraction in computer code made up of perhaps 8 (in the case of the Roman character set) or 16 bits (in the case of larger character sets, such as those of some Asian languages). This is then rendered visually through the mechanised arrangement of dots, or pixels (picture elements), the number of pixels varying according to resolution – a smallish number of dots rendering the particular design of the letter 'A' in 12-point Helvetica to a screen, and many more dots when rendering the same letter to a laser printer. Images are rendered in precisely the same way, as a series of dots in a particular differentiated range of halftones and colours. Whether they are text or images, the raw materials of digital design and rendering are bits and pixels.

One of the practical consequences of this change is that, amongst the text creation trades, typesetting is on the verge of disappearing. It has been replaced by desktop publishing, in which textual and visual design occur on the same screen, for rendering on the same page. Even typing tools, such as *Microsoft Word*, have sophisticated methods for creating (drawing) images, and also importing images from other sources, such as scanned images or photographic images whose initial source is digital.

In retrospect, it is ironic that in the first phase of digitisation, a simplified version of the discourse of typography is de-professionalised, and that this

discourse was a quintessential part of the language-centric semiotics of the Gutenberg era. Reading is actually a matter of seeing, and in this phase of digitisation, the attention of the text creator is directed to the visual aspect of textual design.

Digital technologies make it easy to relate and integrate text and images on the same page, and complex information architectures and multimodal grammars emerge around practices of labelling, captioning and the superimposition of text and image. The overwhelming evidence of this increasing co-location of text and image is to compare a newspaper or school textbook of today with its equivalent 50 or 100 years ago.

We are living in a world that is becoming less reliant on words, or more precisely, a world in which words stand simply and starkly in the linguistic mode. Sometimes the communication has become purely visual – it is possible to navigate an airport simply by using the international pictographs. At other times, the visual and the linguistic are powerfully interwoven in a common communicative framework (Cope & Kalantzis, 2000a, b). This is in part because written language and image, the symbolic and the iconic, are made of the same stuff and fabricated on the same plane.

If the meaning potential of changing technology is realised slowly by the relation of meaning function to meaning form, it occurs even more slowly in the shift away from language-centric to more visually oriented modes of meaning. Actually, this drift began earlier in the twentieth century and before digitisation. Offset lithography made it much easier to put images and text on the same plane, and this occurred through a series of photographically based darkroom and film composition practices. Digitisation merely opens the way for further developments in a long revolution, by making the process of assembling and rendering text with images simpler, cheaper and more accessible to non-tradespeople. And so, we begin a journey in which visual culture is revived, albeit in very new forms, and written-textual culture itself is more closely integrated with visual culture.

There remains, however, a paradox, and that is the increasing use of natural language tagging schemas for the identification, storage and rendering of digital media – visual and audio as well as textual. Prima facie, this seems to represent a setback in the long march of the visual. More profoundly, it may indicate the arrival of a truly integrated multimodality, with the deep inveiglement of the linguistic in other modes of meaning. Here, the linguistic is not just being itself, but also speaking of and for the visual.

One of the most striking manifestations of the effects of the interweaving of the linguistic and the visual is the challenge of multilingualism – or polylingualism as we shall call it – in a globalising communications and information network. For, as we will discuss in the next subsection, a revival of the visual is one significant element in the reversal of the tendencies to monolingualism and language standardisation inherent to textual practices in the era of print.

Polylingualism

Of readers and reading in the electronic age, Roger Chartier suggests two possibilities. One is the 'loss of common references ... the compartmentalisation of groups and the exacerbation of idiosyncrasies'. Yet, the opposite possibility equally presents itself. 'It could also bring about the hegemony of a single cultural model and destroy diversity' (Chartier, 2001).

The ensemble of digital technologies creates a potential in which these two possibilities are not polar opposites, nor even alternatives. Both occur, and simultaneously. Yet again, the change is not technologically determined, but the technologies make possible new kinds of semiosis. And as we take up these possibilities, we resolve the tension between cultural sameness and cultural difference in new ways. We will focus here on several aspects of these technological changes and changes in semiotic practice: new font rendering systems; an increasing reliance on the visual or the visually positioned textual; the emergence of social languages whose meaning functions have been signed at a level of abstraction above the meaning forms of natural language; machine translation assisted by semantic and structural markup; and a trend to customisable technologies which create the conditions for flat economies of scale, which in turn make small and divergent textual communities more viable. To denote the depth of this change, we have coined the word 'polylingual', foregrounding the polyvocal, polysemic potentials deeper than the simple language differences conventionally denoted by the word 'multilingual'.

The fundamental shift in the elementary modular unit of manufacture of textual meaning – from character-level to pixel-level representation – means that platforms for text construction are no longer bound by the character set of a particular national language. Every character is just a picture, and the picture elements can be combined and recombined to create an endless array of characters.

This opportunity, however, was not initially realised. In fact, quite the reverse. The first phase in the use of computers as text bearing machines placed the Roman character set at the centre of the new information and communication technologies. Above the 'bit' (an electric 'off' or 'on' representing a zero or a one), it was agreed that arbitrary combinations of zeros and ones would represent particular characters; and these characters became the foundation of computer languages and coding practices, as well as digitised written-linguistic content. The elementary unit above the 'bit' is the 'byte', using eight bits to represent a character. An eight-bit (one byte) encoding system, however, cannot represent more than a theoretical 256 characters, the maximum number of pattern variations when eight sets of zeros or ones are combined (27, plus an eighth stop bit). The international convention for Roman script character encoding was to became the American Standard Code for Information Interchange (ASCII), as accepted by the American Standards

Association in 1963. In its current form, ASCII consists of 94 characters in the upper and lower case and punctuation marks.

Although one-byte character encoding works well enough for Roman and other alphabetic scripts, it won't work for larger character sets such as the ideographic Asian languages. To represent languages with larger character sets, specialised two-byte systems were created. However, these remained for all intents and purposes separate and designed for localised country and language use. Extensions to the ASCII one-byte framework were also subsequently created to include characters and diacritica from languages other than English whose base character set was Roman. Non-Roman scripting systems remained in their own two-byte world. As the relationship between each character and the pattern of zeros and ones is arbitrary, and as the various systems were not necessarily created to talk to each other, different computer systems were to a large degree incompatible with each other.

However, to return to the fundamentals of digital text technology, pixels can just as easily be arranged in any font, from any language. Even in the case of ASCII, text fabrication seemed like just typing. Actually, it's drawing, or putting combinations of pixels together, and there's no reason why the pixels cannot be put together in any number of drawn combinations.

A new generation of digital technologies is now being built on the Unicode universal character set (Unicode, 2003). In Unicode, every character and symbol in every human language is represented in a consolidated two-byte system. The 94 ASCII Roman characters are now embedded in a new 16-bit character encoding. Here they pale into insignificance amongst the 95,221 characters of Unicode 3.2. These Unicode characters not only capture every character in every human language; they also capture archaic languages such as Linear B, the precursor to Ancient Greek found as inscriptions in Mycenaean ruins. Unicode also captures a panoply of mathematical and scientific symbols. It captures geometric shapes frequently used in typesetting (squares, circles, dots and the like), and it captures pictographs, ranging from arrows to international symbols such as the recycling symbol, to something so seemingly obscure as the set of 15 Japanese dentistry symbols. The potential with Unicode is for every computer and every printer in the world to render text in any and every language and symbol system, and perhaps most significantly for a multilingual world, to render different scripts and symbol systems on the same screen or the same page.

Unicode mixes ideographs and characters as though they were interchangeable. In fact, it blurs the boundaries between character, symbol and icon, and between writing and drawing. Ron Scollon speaks of an emerging 'visual holophrastic language'. He derives the term 'holophrastic' from research on young children's language in which an enormous load is put on a word such as 'some' which can only be interpreted by a caregiver in a context of visual, spatial and experiential association. In today's globalised world, brand logos and brand names (to what language does the word 'SONY' belong? he asks) form an internationalised visual language. A visual holophrastic sign brings

with it a coterie of visual, spatial and experiential associations, and these are designed to cross the barriers of natural language (Scollon, 1999).

These developments are but one aspect of the convergence of the visual-iconic and the linguistic-symbolic, discussed earlier in this article. And this, in turn, is one aspect of increasing polylingualism. The shift is in part a practical response to globalisation. Take the archetypical case of airport signage, where it is simply impossible to operate in the language of every traveller. And take technical manuals – meaning is expressed primarily via image and diagrams; if in the design of the manual text is kept to a minimum, it is a relatively inexpensive task to translate labels and text and insert this into the digitised pages. Now that text and image are fabricated or rendered on the same plane, narrative text, captions and labels are easily translated and one language is substituted for another in the source file. Communications, in other words, are increasingly built on visually structured templates, and the text is a secondary component.

Behind the multimodal semiotics of the visual, meaning functions are expressed via meaning forms in which natural language becomes genuinely arbitrary – or arbitrary in an additional sense to that originally proposed by founding linguist, Sassure, when he identified the arbitrary relation of the language signifier to its signified. Or, to use the terminology of this article, yet another layer of arbitrariness is added to the relation of forms of language to their meaning functions.

We will illustrate this additional layer of arbitrariness with a simple example: ecommerce has enabled banking which presents itself to account holders in the machine interfaces of the automatic teller machine, or Internet banking, or telephone banking. Banking involves what Scollon calls a 'chain of mediated action' (Scollon, 2001). The bank is keeping my money, and in this context I approach it and I ask for some of it back. What follows is a formal interchange, and, if my request proves valid and is approved, the bank says so and gives me some of my money. My sequence of action and communication is motivated by a 'funnel of commitment', in which I attempt to realise my purpose, my meaning function. What follows from my commitment is a kind of narrative structure. In the traditional shopfront bank my purposes are served as I negotiate a complex array of written documentation, supported by actual conversations which frame the details of the transaction and give context to the written documentation. In the world of banking before electronic commerce this was a heavily language-bound activity. You had to fill out a withdrawal slip that was almost invariably only available in the 'national' language of the bank and then speak to a teller in that language. Occasionally, in deference to multiculturalism and to break into niche markets, banks would make sure there were some bilingual tellers – that is, if they wanted to conduct transactions with international tourists, or to serve immigrant languages heavily represented in a local neighbourhood, for instance. But there were practical limits to this, the principal of which is the number of languages that can be serviced by a local branch.

Ecommerce-enabled banking – the ATM, online banking and automated phone banking – has the potential to change all of that. Various highly routine and predictable conversations, such as the 'I want some of my money' conversation, do not really (despite appearances) happen in English. They happen through a translation of the routine operation of withdrawing funds, or seeking an account balance, into a series of computer-generated prompts. The way these prompts are realised in a particular language is arbitrary. There is nothing peculiar or essential to the natural language of the banking conversation. Semantics and grammar, or meaning and information structure, are everything. The logic of the communicative exchange now operates below the level of language; it has been designed that way, and it works that way. Various 'banking conversations' are constructed as a universal, transnational, translinguistic code (actually, computer code, because the customer is 'talking' to the bank's computer), in which the manifestation of that code in natural language is, in a communicative sense, arbitrary. You can choose any language you like at the beginning of the online banking session and the visible 'tags' describing the effect of pressing alternative buttons will be translated into your language of choice. There is nothing to stop this being in any script; or the screen swapping its directionality if you were to choose Arabic; or in non-visual interfaces, such as Braille; or interfaces translating text to audio. The ATM and voice-synthesised telephone banking do the same thing, working off the same ecommerce-abstracted text. The rendering of the meaning form can vary radically; but the meaning function remains constant. The business of making the banking service available in another language is as simple as translating the labels which represent the tagged information to the bank customer – a few hundred words at the most.

Once, the grammar of language was the entry point into the grammar of banking. Unless the customer and the bank were able to operate competently in the same conversational, written and thus cultural world, there could be no transaction. Banking was a language-delimited game, and the prescribed language or languages were a non-negotiable precondition for playing the banking game. However, in the world of ecommerce, the functional grammar of banking is created first, and this grammar can be realised in any language. This functional grammar does not only speak; it also invites a number of physical actions, such as pressing a particular button and taking the money.

Coming back to the questions of polylingualism, this example captures a quite contradictory tendency. In so doing it also answers in part Chartier's rhetorical question. On the one hand, billions of people have been drawn into the culture of ATMs since they were introduced in the last quarter of the twentieth century. To use a term defined and developed by linguist Jim Gee, they have become proficient speakers of a 'social language' (Gee, 1996). In our example, we might want to call this social language 'global ATM' or 'electronic banking'. The particular natural language form in which this social language is realised in the instance of a single transaction is, measured in terms of human action and social meaning, an arbitrary and increasingly trivial accident of birth.

Yes, the culture of electronic commerce and modern banking is taking over the world, extending a certain kind of global sameness. Doing this sameness multilingually might be seen as a kind of ploy. But the facility also supports another kind of linguistic and cultural diversity, because it has reduced the need for speakers of smaller or locally marginal languages to move over into a dominant language, at least for the purposes of banking. Now you can play the global banking game, but you do not have to leave your culture and language behind to do it. You can be in a country where your language is not spoken in banks, and it does not matter because you can go to an ATM or ring telephone banking and deal with synthesised audio, or if required, be directed to a live operator in a call centre somewhere in the world who speaks your language. This is just one small and symptomatic example of the way in which new communications technologies may support language diversity, and make it less important in many settings to know a lingua franca such as English (Cope, 2001, a, b, c).

Behind this shift lie the 'tagging' technologies described earlier in this article. Tagging frameworks operate structurally and semantically. So, two of the tags behind our banking conversation may be:

<customername>Mary Kalantzis</customer name>

and

<withdrawalamount>50.00</withdrawalamount>.

The data contained by tags, in this case, can be any language or scripting system – and, for that matter, can be digitally recorded images or sound. However, a layer of arbitrariness is added to the tags themselves. The tags have a meaning function and, even in the computer code, they can be in any language and any script and still work as representation and communication within a particular social language. This is because meaning has been designed as a kind of functional map of that social language, the typical ways in which chains of mediated action play themselves out – the various conversational alternatives in the 'I want some of my money' conversation, for instance.

The underlying design technique is based on the conceptualisation of meaning function. The practical solution, which this article will explore in greater depths in subsequent sections, is to stabilise each tag schema as a controlled vocabulary. This is supplemented by tag dictionaries that spell out in translation, across however many natural languages that may be required, the precise referents and the ontologically given structural and conceptual relations between the meaning functions to which the tags refer. Schemas are used to represent tags paradigmatically, typically represented in taxonomies. Tag relations can also be represented as narrative, as activity sequences of a syntagmatic variety, and these alternative conversational or narrative sequences may be represented in flow diagrams. These two, essentially visual, devices represent meaning function at a level of abstraction beyond the level of natural language. They are tools for the construction of a relatively stable

semantic ground below the level of natural language. Now, the primary basis for the design of meaning is not the instantiation of meaning in the meaning forms of language (although this is the equally important but now secondary concern of stylesheet transformations). The basis, rather, is the activity and conceptual structures of human intention and experience, or meaning functions.

Tag dictionaries and the linguistic elements of the façade of user interfaces for digital text may be created through human translation – a relatively easy and inexpensive process when the language elements of an interface, or printed text for that matter, have been designed to be substitutable. Marking up for meaning function, however, is an important basis for the increasingly sophisticated technologies of machine translation (Gerber, 2001). It makes meaning functions less dependent on contextual markers and shared understandings between communicants. Mention of 'cope' would normally need to be contextualised in order to distinguish its particular meaning amongst its various meaning possibilities. But if we are working in a publishing schema in which the word is marked up semantically as <surname> and in a structural context where that surname refers to an <author> role, the markup will assist accurate translation.

The possibility created by these technologies is to reduce the relevance of language differences. It will be increasingly possible to participate in the all-encompassing world of global modernity without having to submit to one of its domineering language forms. And this is the response to Chartier's dilemma – we will all be able to speak with each other, and that capacity to communicate will be without prejudice to diversity.

There is still another profound way in which the post-Gutenberg technologies fundamentally shift the means of production of meaning and the ground of culture, and that is to reverse the logic of mass production and economies of scale. This is exemplified by the myriad websites (compare this to the number of 'publishers' in the world of print media), and the shift from the logic of broadcasting and mass media to the 'narrowcasting' or even 'pointcasting' of emerging electronic media systems. However, although it has none of the aura of newness of the digital electronic media, we will focus our attention on the less immediately obvious case of the printed book.

In the Gutenberg era, letterpress, and later gravure and lithographic offset printing, involved the creation of plates, the cost of which have to be amortised over the length of a run (Dunn et al, 2001). Typically, the set-up cost of creating a plate using traditional print technologies is between two hundred and a thousand times the cost of a single print in a viable print run. This means that the core commercial and cultural logic of traditional print manufacture centres around a process of reproduction. In its very nature it involves the replication of many copies of an identical original image. The effect is to favour markets and cultures of large scale. Jobs will only be printed if the run length can justify the cost per unit, and the longer the run, the lower the cost per unit. The bigger the culture, the more likely it is to get supported by traditional print.

This fundamental commercial and cultural logic is at work in all technologies of reproduction.

Although some aspects of the printing process have been digitised since the 1970s, it was not until the 1990s that digitisation was applied directly to the manufacturing process. This occurred through the application of laser technology to the older electromagnetic technology of Xeroxography (Dunn et al, 2001). This fully digital print constructs a final image directly on a substrate, or an impression medium, by arranging the elements of the image (pixels) dot-by-dot from computer code.

The most revolutionary feature of this technology is its variability. Rapidly printed consecutive pages can differ from each other with no fluctuation in speed and printing functionality. In other words, every new impression can be different from the previous impression as easily as it can be the same. A number of terms are frequently used to describe this technological shift including 'print on demand' and 'digital printing', but none of them capture the shift as appropriately as the notion of 'variable print'. In some respects, the term 'digital print' is a misnomer, for the reasons already discussed – almost all print was already digital in some respects. Only fully digitised print manufacture captures the dramatic potential of variability. Similarly, 'print on demand' is not strictly accurate or a particularly useful concept. Every commercial printer will tell you they print on demand in terms of meeting their customers' expectations as quickly and effectively as possible. And even fully digital presses need to go through a number of business systems and process steps (ordering, production, dispatch), which never means that printing is precisely on demand.

By contrast with the ensemble of Gutenberg technologies, the commercial and cultural logic of fully digital or variable print begins with the fact that every print is an original. In the case of digitally rendered images, every one has been constructed, not from a reproduced original, but within a source file of ephemeral and unreadable computer code and the final print is the first and only rendering. There is no reason why it should be rendered in this particular form ever again, and if it is, technically speaking it is another original. This, in turn, engenders flat economies of scale. As every pixel is formed afresh in the rendering of each impression, there is no difference between the cost of rendering identical successive impressions and rendering different impressions. As there are no plates or set-up costs, one impression shares the same cost per unit as one thousand impressions. As a consequence, there are no economies of manufacturing scale. In the domain of manufacturing, long runs have no particular advantage over short runs; niche markets are no less viable than mass markets; small languages and cultures can be serviced as easily as large ones; and even the 'digital divide' can be bridged with a few computers combined with digital print.

This ensemble of changes opens the possibility (although by no means does it preordain the inevitability) of richer polylingualism. Returning to the spirit of our definition of this term, this not only applies to natural languages

(conventional understandings of multilingualism), it also applies to the social languages of discourse communities, and these may be defined by profession, ethnicity, sub-cultures, fashions, style, fad or fetish. The remarkable paradox of globalisation is that, even when these social languages are expressed in the more and more extensive lingua francas of natural language (Chinese or English, for instance), social languages are diverging. They are, in fact, becoming less mutually intelligible, and that reducing intelligibility is manifest in the likes of technicality and dialect.

Meaning Form and Meaning Function

In this article, we have been developing the concepts of 'meaning form' and 'meaning function' as the foundations for a discussion of some of the changes in the means of production of meaning brought about by digital technologies. Our argument weaves between an historical and theoretical interpretation of the significance of aspects of these changes in the preceding two sections of this article and a discussion in the following sections of our own practical attempt to create a functional grammar of textual meaning focused on the practices of authorship and publishing in general, and on the information architecture of the book in particular.

Signs are the elementary components of meaning. And 'signs', say Kress & Leeuwen (1996), are 'motivated conjunctions of signifiers (forms) and signifieds (meanings)'. Rephrasing the terms of this definition, we would call motivated meanings – the impulse to represent the world and communicate those representations, 'meaning functions'. The business of signing, motivated as it is by representation and communication, entails an amalgam of function (a reason to mean) and form (the use of representational resources which might adequately convey that meaning). The meaning function may be a flower in a garden upon which we have fixed our focus for a moment through our faculties of perception and imagination. For that moment, this particular flower captures our attention and its features stand out from its surroundings. The meaning function is our motivation to represent this meaning and to communicate about it. How we represent this meaning function is a matter of meaning form. The meaning form we choose might be iconic – we could draw a sketch of the flower, and in this case, the act of signing (form meets function) is realised through a process of resemblance. Meaning form – the drawing of the flower – looks like meaning function, or what we mean to represent: the flower. Or the relation between meaning form and function may be, as is the case of language, arbitrary. The word 'flower', a symbolic form, has no intrinsic connection with the meaning function it represents. In writing or in speech the word 'flower' conventionally represents this particular meaning function in English. We can represent the object to ourselves using this word in a way which fits with a whole cultural domain of experience (encounters with other flowers in our life and our lifetime's experience of speaking about and hearing about flowers). On the basis of this conventional understanding of meaning

function, we can communicate our experience of this flower or any aspect of its flower-ness to other English speakers.

This, in essence, is the stuff of signing, the focal interest of the discipline of semiotics. It is an ordinary, everyday business, and the fundamental ends do not change when employing new technological means. It is the stuff of our human natures. The way we mean is one of the distinctive things that make us human.

We have been attempting to describe in this article the key dimensions of the 'digital revolution', in which the mechanics of conjoining meaning functions into meaning forms is changing in some subtle but nevertheless profound ways. Much has been said about this revolution. Some of it is optimistic hyperbole – about, for instance, the liberating power of cheap and easy information access and instantaneous hyperlinks which break apart linear structures of meaning and transform the role of the reader. Some of it is bleakly denigratory – about the sterility of machine-mediated communication and the centralisation of ownership and control of intellectual property in the hands of a class of wealthy individuals and affluent countries and by the exclusion of the bulk of the world's population from powerful domains of knowledge. At this moment, it is hard to say whether the optimists or the pessimists will be proved correct.

A less noticed aspect of the digital revolution is the series of interconnected changes in the means of production of signs which we have described thus far in this article. We have developed our account of these changes through an analysis of *the sign as the motivated conjunction of meaning form and meaning function*. To summarise, the changes are threefold:

- *The mechanics of rendering*. The immediate focus of the business of signing – more broadly conceived as the design of meaning – shifts from configuring meaning form (the specifics of the audible forms of speaking and the visual form of written text) to 'marking up' for meaning function in such a way that alternative meaning forms, such as audio or visual (written) forms of language, can be rendered by means of automated processes from a common digital source.
- *Modality*. In the digital era, written, imaged and audio forms of meaning are manufactured of the same stuff (the zeros and ones of computing bits that can equally generate symbolic character representations as bytes, iconic representations as pixels and audio representations as ticks). This new means of production of signs puts the visual, the written-textual and the audio on the same plane; and the visual and the written-textual even more so, as bytes are rendered through the configuration of pixels that constitute a particular character. As these varied digital meaning forms are made of the same raw materials, it is simply more practical then ever before to realise meaning functions multimodally.
- *Language Specificity*. Meanings are increasingly designed on the basis of underlying social languages, generated from non-language-specific structures of meaning (information architectures or activity sequences), with

various techniques used to facilitate translatability into a variety of natural languages – such as the techniques of machine translation, holophrastic visual representation or the shift from symbolic (language delimited) to iconic (increasingly universal, internationalist) representation. On top of the abstract arbitrariness of language itself is added another layer of arbitrariness – the genuine and practical arbitrariness of a particular language to the realisation of a meaning function. And once again the shift in the mechanics of semiosis as away from the logistics of language-bound meaning form to an analysis, description and record of meaning function.

On one scale – the scale which is the domain of semiotics, on which we find some of the key characteristics of what makes us human – these changes are tiny. However, on another scale of reference – when we take the modern tradition of writing and publishing as our measure, for instance – the changes are probably quite significant. Put simply, the fundamental tools we use for the creation of signs, or the means of production of meaning, are changing. And as we change our tools, our tools change us. Then we may discover that we are starting to be human in ever-so-subtly new ways.

The following sections of this article will investigate the specifics of one social language, the social language of the book, and the new ways in which this is being expressed in the world of digital meaning.

Digital Schemas for the Book

The current mix of digital technologies utilised in the process of making a book – digital print, the Internet, and integrated digital content, process and financial management systems – represents a moment of change that will be perhaps as significant as Gutenberg's invention of the printing press. One of the defining features of this change is the emergence of electronic standards which are free, open source, and which cross different proprietary technologies and different parts of the production process. Just as other digital information technologies have stabilised to agreed standards (the Internet, CD, DVD, etc.), the same is destined to be the case in book production.

It is increasingly proving to be the case that these standards perform a wide-ranging, fundamental and integrated set of functions. They *contain* the content – the electronic files that provide structural and semantic shape for the data which will be rendered as a book. They *describe* the content – for the purposes of data transfer, warehousing and retrieval. They *manage* the content – providing a place where job process instructions and production data are stored. And they *transact* the content – managing ecommerce transactions (B-2-B or B-2-C).

A number of digital tagging schemas have emerged which provide a functional account of these processes of containing, describing, managing and transacting books. More broadly, they provide a functional account of the world of textual content in general. Each tagging schema has its own functional purpose, or 'funnel of commitment', to use Scollon's terminology. We will

briefly describe a few of these below, categorising them into domains of professional and craft interest: typesetting and content capture, electronic rendering, print rendering, resource discovery, cataloguing, educational resource creation, ecommerce and digital rights management. The ones we will describe also happen to be those we have mapped into the most recent iteration of Common Ground Markup Language, as discussed in detail earlier in this article.

Typesetting and Content Capture

Unicode (www.unicode.org) appears destined to become the new universal, character encoding standard, covering all major language and scripts (Unicode, 2003), and replacing ASCII, which was based solely on Roman script.

A number of tagging schemas have been created for the purpose of describing the structure of text, and to facilitate its rendering to alternative formats. These schemas are mostly derivatives of SGML. HTML4 (World Wide Web Consortium, 1999) and XHTML (World Wide Web Consortium, 2002) are designed primarily for rendering transformations through web browsers. The DocBook standard, sanctioned by the Organization for the Advancement of Structured Information Standards and the United Nations Educational, Scientific, and Cultural Organization, is for structuring book text, which can subsequently be rendered electronically or to print (DocBook Technical Committee, 2003). The Text Encoding Initiative is 'an international and interdisciplinary standard that helps libraries, museums, publishers, and individual scholars represent all kinds of literary and linguistic texts for online research and teaching' (Text Encoding Initiative, 2003). And LaTeX is a system for structuring text for typesetting, 'with features designed for the production of technical and scientific documentation'. It has become, in the words of the LaTeX website, 'the de facto standard for the communication and publication of scientific documents' (LaTeX, 2003).

Although the primary purpose of each schema may be a particular form of rendering, this belies the rigorous separation of semantics and structure from presentation. Alternative stylesheet transformations could be applied to render the marked-up text in a variety of ways. Using different stylesheets, it is possible, for instance, to render DocBook either as typesetting for print or as HTML.

Electronic Rendering

Electronic rendering can occur in a variety of ways – as print facsimiles in the form of PDF (Portable Document Format), or as HTML readable by means of a web browser. Other channel alternatives present themselves as variants or derivatives of HTML: the Open eBook Standard for handheld electronic reading devices (International Trade Standards Organization for the eBook Industry, 2003) and Digital Talking Book (American National Standards

Institute/National Information Standards Organization, 2002), facilitating the automated transition of textual material into audio form – for the visually impaired, or the convenience of listening to a text rather than reading it.

Print Rendering

The Job Definition Format (JDF) appears destined to become universal across the printing industry (CIP4 Organization, 2003). Specifically for variable print, Personalised Print Markup Language has also emerged (PODi, 2003).

Created by a cross-industry international body, the Association for International Cooperation for the Integration of Processes in Pre-Press, Press and Post-Press, the JDF standard has been embraced and supported by all major supply-side industry participants (equipment and business systems suppliers). It means that the one electronic file contains all data related to a particular job. It is free (in the sense that there is no charge for the use of the format) and open (in the sense that its tags are transparently presented in natural language; it is unencrypted, its coding can be exposed and it can be freely modified, adapted and extended by innovators – in sharp distinction to proprietary software).

The Job Definition Format functions as a digital addendum to offset print, and as the driver of digital print. Interoperability of JDF with other standards will mean, for instance, that a book order triggered through an online bookstore (the ONIX space, as described below) could generate a JDF wrapper around a content file as an automated instruction to print and dispatch a single copy.

The Job Definition Format serves the following functions:

- Pre-press: full job specification, integrating pre-press, press and post-press (e.g. binding) elements, in such a way that these harmonise (the imposition matches the binding requirements, for example). These data are electronically 'tagged' to the file itself, and in this sense, they actually 'make' the 'printing plate'.
- Press: the job can then go onto any press from any manufacturer supporting the JDF standard (and most major manufacturers now do). This means that the press already 'knows' the specification developed at the pre-press stage.
- Post-press: once again, any finishing is determined by the specifications already included in the JDF file, and issues such as page format and paper size are harmonised across all stages in the manufacturing process.

The effects of wide adoption of this standard by the printing industry include:

- Automation: there is no need to enter the job specification data from machine to machine, and from one step in the production process to the next. This reduces the time and thus the cost involved in handling a job.
- Human error reduction: as each element of a job specification is entered only once, this reduces waste and unnecessary cost.

- Audit trail: responsibility for entering specification data is pushed further back down the supply chain, ultimately even to the point where a customer will fill out the 'job bag' simply by placing an order through an online B-2-B interface. This shifts the burden of responsibility for specification, to some degree, to the initiator of an order, and records by whom and when a particular specification was entered. This leads to an improvement in ordering and specification procedures.
- Equipment variations: the standard reduces the practical difficulties previously experienced using different equipment supplied by different manufacturers. This creates a great deal of flexibility in the use of plant.

Resource Discovery

Resource discovery can be assisted by metadata schemas that use tagging mechanisms to provide an account of the form and content of documents. In the case of documents locatable on the Internet, Dublin Core is one of the principal standards, and is typical of others (Dublin Core Metadata Initiative, 2003). It contains a number of key broadly descriptive tags: <title>, <creator>, <subject>, <description>, <publisher>, <contributor>, <date>, <resource type>, <format>, <resource identifier>, <source>, <language>, <relation>, <coverage> and <rights>. The schema is designed to function as a kind of electronic 'catalogue card' to digital files, so that it becomes possible, for instance, to search for Benjamin Disraeli as an author <creator> because you want to locate one of his novels, as opposed to writings about Benjamin Disraeli as a British Prime Minister <subject> because you have an interest in British parliamentary history. The intention of Dublin Core is to develop more sophisticated resource discovery tools than the current web-based search tools, which, however fancy their algorithms, do little more than search indiscriminately for words and combinations of words.

A number of other schemas build upon Dublin Core, such as the Australian standard for government information (Australian Government Locator Service, 2003), and the EdNA and United Kingdom (UK) National Curriculum standards for electronic learning resources. Other schemas offer the option of embedding Dublin Core, as is the case with the Open eBook standard.

Cataloguing

The MARC (Machine Readable Catalog) format was initially developed in the 1960s by the US Library of Congress (Mason, 2001; MARC Standards Office, 2003a, b, c). Behind MARC is centuries of cataloguing practice, and its field and coding alternatives run to many thousands. Not only does MARC capture core information such as author, publisher or page extent, it also links into elaborate traditions and schemas for the classification of content such as the Dewey Decimal Classification system or the Library of Congress Subject Headings.

MARC is based upon ISO 2709 'Format for Information Exchange'. MARC has recently been converted into an open XML standard.

The original markup framework for MARC was based on non-intuitive alphanumeric tags. Recent related initiatives have included a simplified and more user-friendly version of MARC: the Metadata Object Description Schema (MARC Standards Office, 2003c) and a standard specifically for the identification, archiving and location of electronic content, the Metadata Encoding and Transmission Standard (MARC Standards Office, 2003b).

Various 'crosswalks' have also been mapped against other tagging schemas, notably MARC to Dublin Core (MARC Standards Office, 2001) and the MARC to the ONIX ecommerce standard (MARC Standards Office, 2000). In similar territory, although taking somewhat different approaches to MARC, are Biblink (United Kingdom Office for Library and Information Networking, 2001) and Encoded Archival Description Language (Encoded Archival Description Working Group, 2002).

Educational Texts

Cutting across a number of areas – particularly rendering and resource discovery – are tagging schemas designed specifically for educational purposes. EdNA Online (2000) and the UK National Curriculum Metadata Standard (National Curriculum Online, 2002) are both variants of Dublin Core.

Rapidly rising to broader international acceptance, however, is the Instructional Management Systems Standard (IMS Global Learning Consortium, 2003) and the related Shareable Content Object Reference Model (ADL/SCORM, 2003). Not only do these standards specify metadata to assist in resource discovery, they also build and record conversations around interactive learning, manage automated assessment tasks, track learner progress and maintain administrative systems for teachers and learners. The genesis of IMS was in the area of metadata and resource discovery, and not the structure of learning texts. One of the pioneers in the area of structuring and rendering learning content (building textual information architectures specific to learning and rendering these through stylesheet transformations for web browsers) was Educational Modelling Language (OUL/EML, 2003). More recently, EML has been grafted into the IMS suite of schemas and renamed the IMS Learning Design Specification (IMS Global Learning Consortium, 2002).

Ecommerce

One tagging schema has emerged as the dominant standard for B-2-B ecommerce in the publishing supply chain – the ONIX, or the Online Information Exchange standard, initiated in 1999 by the Association of American Publishers, and subsequently developed in association with British publishing and bookselling associations (EDItEUR, 2001; Mason & Tsembas, 2001). The purpose of ONIX is to capture data about a work in sufficient detail

to be able automatically to upload new bookdata to online bookstores such as Amazon.com, and to communicate comprehensive information about the nature and availability of any work of textual content. ONIX sits within the broader context of interoperability with ebXML, an initiative of the United Nations Centre for Trade Facilitation and Electronic Business.

Digital Rights Management

Perhaps the most contentious area in the world of tagging is that of digital rights management (Cope & Freeman, 2001). Not only does this involve the identification of copyright owners and legal purchasers of creative content; it can also involve systems of encryption by means of which content is only accessible to legitimate purchasers; and systems by means of which content can be decomposed into fragments and recomposed by readers to suit their specific needs. The <indecs>, or Interoperability of Data in E-Commerce Systems framework, was first published in 2000, the result of a two-year project by the European Union to develop a framework for the electronic exchange of intellectual property (<indecs>, 2000). The conceptual basis of <indecs> has more recently been applied in the development of the Rights Data Dictionary for the Moving Pictures Expert Group's MPEG-21 framework for distribution of electronic content (Multimedia Description Schemes Group, 2002). From these developments and discussions, a comprehensive framework is expected to emerge, capable of providing markup tools for all manner of electronic content (International DOI Foundation, 2002; Paskin, 2003).

Amongst the other tagging schemas marking up digital rights, Open Digital Rights Language is an Australian initiative which has gained wide international acceptance and acknowledgement (ODRL, 2002). And XrML, or Extensible Rights Markup Language, was created in Xerox's PARC laboratories in Paulo Alto. Its particular strengths are in the areas of licensing and authentication (XrML, 2003).

What Tagging Schemas Do

The tagging schemas we have mentioned here do almost everything conceivable in the world of the written word. They can describe that world comprehensively, and to a significant degree, they can make it happen. The typesetting and content capture schemas provide a systematic account of structure in written text, and through stylesheet transformations they can literally print text to paper, or render it electronically to screen or manufacture synthesised audio. Digital resource discovery and electronic library cataloguing schemas provide a comprehensive account of the form and content of non-digital as well as digital texts. Educational schemas attempt to operationalise the peculiar textual structures of traditional learning materials and learning conversations, where a learner's relation to text is configured into an interchange not unlike the ATM conversation we described earlier in this

article. Ecommerce and digital rights management schemas move texts around in a world where intellectual property rights regulate their flow and availability.

They are part of a larger movement known as the 'semantic web', that promises to take networked electronic data one step beyond the jumble of the World Wide Web (Berners-Lee et al, 2001; W3C, 2002). The semantic web builds ontologies – each a systematic, coherent and consistent set of labels to describe the meaning elements of a particular domain, as well as the systematic relationships of these concepts to each other. This may be represented as paradigm (using syntagmatic devices such as taxonomy) or as narrative (an account of the funnel of commitment and the alternative activity sequences or navigation paths in the negotiation of that commitment). Ontologies are like theories, except, unlike theories, they do not purport to be hypothetical or amenable to testing; they purport to tell of the world, or at the very least a part of the world, like it is – in our case that part of the world inhabited by authors, publishers, librarians, bookstore workers and readers. The next generation of ontology-based markup brings with it the promise of more accurate discovery, machine translation and, eventually, artificial intelligence. A computer really will be able to interpret the difference between Cope, cope and cope. Even in the case of the <author> with the seemingly unambiguous <surname> Kalantzis, there is semantic ambiguity that markup can eliminate or at least reduce, by collocating structurally related data (such as date of birth) to distinguish this Kalantzis from others and by knowing to avoid association with the transliteration of the common noun in Greek, which means 'tinker'.

In the world of XML, tags such as <author> and <surname> are known as 'elements', which may well have specified 'attributes'; and the ontologies are variously known as, or are represented in, 'schemas', 'application profiles' or the 'namespaces' defined by 'document type definitions' or DTDs. As our interest in this article is essentially semantic, we will use the concepts of 'tag' and 'schema'. In any event, 'ontology' seems the wrong concept insofar as tag schemas are not realities; they are specific constructions of reality within the frame of reference of highly particularised social languages. Their reality is a social reality. They are no more and no less than a 'take' on reality which reflects and represents a particular set of human interests. These interests are fundamentally to get things done (funnels of commitment) more than they are mere reflections of objectified, inert being. Schemas, in other words, have none of the immutable certainty implied by the word 'ontology'. Reality does not present itself in an unmediated way. Tagging schemas are mediated means rather than static pictures of reality.

Most of the tagging frameworks relating to authorship and publishing introduced above were either created in XML or have now been expressed in XML. That being the case, you might expect that an era of rapid and flexible transmission of content would quickly dawn. But this has not occurred, or at least not yet, and for two reasons. The first is the fact that, although almost all content created over the past quarter of a century has been digitised, the

formats are varied and incompatible. Digital content is everywhere, but most of it has been created, and continues to be created, using legacy design and markup frameworks. These frameworks are embedded in software packages that provide tools for working with text which mimic the various trades of the Gutenberg universe: an author may use *Word*; a desktop publisher or latter-day typesetter may use *Quark*; and a printer will use a PDF file as if it were a virtual forme or plate. The result is sticky file flow and intrinsic difficulties in version control and digital repository maintenance (Cope, 2001a, b, c). How and where is a small correction made to a book that has already been published? Everything about this relatively simple problem, as it transpires, becomes complex, slow and expensive. However, in a fully comprehensive, integrated XML-founded file flow, things that are slow and expensive today should become easier and cheaper – a small change by an author to the source text could be approved by a publisher so that the very next copy of that book purchased online and printed on demand could include that change. Moreover, even though just about everything available today has been digitised somewhere, in the case of books and other written texts, the digital content remains locked away for fear that it might get out and about without all users paying for it when they should. Not only does this limit access but what happens, for instance, when all you want is a few pages of a text and you do not want to pay for the whole of the printed version? And what about access for people who are visually impaired? It also puts a dampener on commercial possibilities for multichannel publishing, such as the student or researcher who really has to have a particular text tonight, and will pay for it if they can get it right away in an electronic format – particularly if the cost of immediate access is less than the cost of travelling to the library specially.

The second reason that a new era of text creation and transmission has not arrived is semantic. Even though XML is spreading quickly as a universal electronic lingua franca, each of its tagging schemas describes their worlds in their own peculiar ways. Tags may well be expressed in natural languages – this level of simplicity, openness and transparency is the hallmark of the XML world. But herein lies a trap. There is no particular problem when there is no semantic overlap between schemas. However, as most XML application profiles ground themselves in some ontological basics (such as people, place and time), there is nearly always semantic overlap between schemas. The problem is that, in everyday speech, the same word can mean many things, and XML tags express meaning functions in natural language.

The problem looms larger in the case of specialised social languages. These often develop a high level of technical specificity, and this attaches itself with a particular precision to key words. The more immersed you are in that particular social language – the more critical it is to your livelihood or identity in the world, for instance – the more important these subtle distinctions of meaning are likely to be. Communities of practice identify themselves by the rigorous singularity of purpose and intent within their particular domain of practice, and this is reflected in the relative lack of terminological ambiguity

within the social language that characterises that domain. As any social language builds on natural language, there will be massive ambiguities if the looser and more varied world of everyday language is assumed to he homologous with a social language which happens to use some of the same terminology.

The semantic differences between two social languages in substantially overlapping domains are likely to be absolutely critical. Even though they are all talking about text and can with equal facility talk about books, it is the finely differentiated ways of talking about books that make authors, publishers, printers, booksellers and librarians different from each other. Their social language is one of the ways you can tell the difference between one type of person and another. These kinds of difference in social language are often keenly felt and defended. Indeed, they often become the very basis of professional identity.

This problem of semantics is the key dilemma addressed by this article, and the focal point of the research endeavour which has produced Common Ground Markup Language. Our focus in this research has been the means of creation and communication of textual meaning, of which the book is an archetypical instance. Each of the schemas we have briefly described above channels a peculiar set of 'funnels of motivation' in relation to books – variously that of the author, typesetter, printer, publisher, bookseller, librarian and consumer. And although they are all talking about the same stuff – textual meaning in the form of books – they talk about it in slightly different ways, and the differences are important. The differences distinguish the one funnel of commitment employing its own peculiar social language to realise that commitment, from another. It is precisely these differences which give shape and form to the tagging schemas which have been the subject of our investigations.

The schemas we have identified range in size from a few dozen tags to a few thousand, and the total number of tags across just these schemas would be in the order of tens of thousands. This extent alone would indicate that the full set of tags provides the basis for a near-definitive account of textual meaning. And although it seems as if these schemas were written almost yesterday, they merely rearticulate social languages that have developed through 500 years of working with the characteristic information architectures of mechanically reproduced writing, of bibliography and librarianship, of the book trade, and of readership. Given that they are all talking about authorship and publishing, the amount of overlap (the number of tags that represent a common semantic ground across all or most schemas) is unremarkable. What is remarkable is the subtle variations in semantics depending on the particular tagging schema or social language; and these variations can be accounted for in terms of the subtly divergent yet nevertheless all-important funnels of commitment.

So, after half a century of computing and a quarter of a century of the mass digitisation of text, nothing is really changing in terms of the core business of representing the world using the electrical on/off switches of digitisation.

The technology is all there, and has been for a while. The half-millennial shift is in the underlying logic behind the design of textual meaning. This shift throws up problems which are not at root technical; rather they are semantic. Interoperability of tagging schemas is not a technical problem, or at least it is a problem for which there are relatively straightforward technical solutions. The problem, and its solution, is semantic.

The commercial implications of the emergence and stabilisation of electronic standards are also enormous. These include:

- Efficiencies and cost reduction: electronic standards facilitate cross-enterprise automation of file flow, including the process and commercial aspects of that flow – from the creator to the consumer. Efficiencies will also created by B-2-B and B-2-C relationships based on standards, including error reduction, single entry of data and moves towards progressive automation of the production process.
- Supply chain integration: electronic standards also mean that closer relationships can and should be built between the links of the publishing supply chain. For instance, a publisher ordering a print run of books can enter data into the printer's Job Definition Format via a web interface. It is also possible to transfer ONIX data automatically into this format, thus creating publisher-printer-bookseller supply chain integration. The key here is the creation of trusted 'most favoured supplier' relationships and the development of a sense that the destinies of closely related enterprises are intertwined, rather than antithetical to each other's interests.
- New business opportunities: as new, hybrid enterprises emerge which create links across the supply chain, offering services such as the multipurposing of content (for instance, to the Web, to handheld reading devices and to digital talking books) and data warehousing. This will be supported particularly by the emergence of supply chain-wide product identification protocols such as the Digital Object Identifier.

The key issue is the flow of business information and content between the various players in the book production supply chain. Addressing this issue can produce efficiencies and competitive advantage for individual enterprises, and the whole industry. Many players are now arguing, in fact, that addressing this issue is not a choice – it is a necessity given the fact that standards are rapidly emerging, stabilising and gaining wide acceptance.

Common Ground Markup Language is an attempt to address these semantic and commercial challenges.

Common Ground Markup Language

The aim of the Common Ground Markup Language (CGML) research and development endeavour has been to develop software that enables digital text production (electronic renderings and print renderings) using a markup language which offers stable and reliable interoperability across different standards. These standards include typesetting and text capture, electronic

rendering, print rendering, B-2-B ecommerce, e-learning, digital rights management, Internet resource discovery and library cataloguing.

CGML addresses one of the fundamental issues of the semantic web – the problem of interoperability between different but overlapping and related electronic standards. Commercially and functionally, the intended result is a software environment in which texts render simultaneously to electronic and print formats (for instance, a bound book, computer screen, handheld reading device or synthesised voice) from a common source file. Metadata generated by this software is simultaneously able to create a library cataloguing record, an ecommerce record (automated entry to Amazon, international bookdata databases, etc.) and make a published text an e-learning object and conform to the current and emerging digital rights management protocols. The foundation of this software is Common Ground Markup Language.

At the time of writing, CGML consists of nearly a thousand tags, interpolated into an XML Document Type Definition (DTD). These tags are defined in the *Common Ground Dictionary of Authorship and Publishing*, which currently runs to some 25,000 words. CGML and the Dictionary are published dynamically (without 'editions' or 'versions'), with tags being constantly added and definitions refined as the Common Ground research endeavour proceeds. The Common Ground Research and Development team has created its own ontology building software, *CommonGroundLEXICOGRAPHER*, which houses the *Common Ground Dictionary of Authorship and Publishing*, as well as providing a foundation for the export of data into a range of XML text and publishing schemas.

The purpose of this section of our article is to outline the design principles underlying CGML, and the practical ways in which it attempts to address the semantic problem which we have identified as a key challenge in the development of the digital communications environment.

CGML takes the textual artefact of the book as its point of departure. The book is the archetypical, most established and first characteristically modern communications medium. In practical terms, the book is a carrier of text and still images, and as a consequence, the focus of CGML is on a certain range of linguistic and visual modes of meaning. However, insofar as the digital medium also serves as a construction tool, repository and distributional means for audio, moving image, software, databases and the like, CGML in its current iteration incorporates rudimentary reference to these representational forms. CGML is, nevertheless, designed to be fully extensible into all domains of creative media and cultural artefact.

From a technology point of view, CGML sets out to tackle one of the fundamental challenges of the semantic web – the problem of interoperability between overlapping and related electronic standards. To this end, Common Ground researchers developed the interlanguage mechanism described later in this article, and currently the subject of an international patent application (Common Ground, 2003a, b). This mechanism has the potential to extend the useability of content across multiple standards, XML schemas, ontologies or

database structures. Developed by Common Ground's open source software development team, the approach taken by CGML may begin to address the enormous problem of interoperability in general, not just in publishing but in other areas of the semantic web. Stated simply, electronic files do not flow well along production and distribution supply chains, and this is not only because file formats vary, but so also do the metadata which define file content format and uses. In the case of published material, there are enormous inefficiencies in file flow from author to publisher to printer to electronic rendering formats, as well as the ecommerce mechanisms which deliver content from enterprise to enterprise and finally to consumers.

Even though each electronic standard or XML DTD has its own functional purposes, there is a remarkable amount of overlap between these standards. The overlap, however, often involves the use of tags in mutually incompatible ways. Our extensive preliminary mapping of 17 standards in various text-related and publishing-related fields shows that, on average, each standard shares 70% of its semantic range with neighbouring standards. Despite this, it is simply not possible to transfer data generated in one standard to a database or XML schema using another. Each standard has been designed as its own independent, standalone DTD. This, in fact, points to one of the key deficiencies of XML as a meta-markup framework: it does not in itself suggest a way for DTDs to relate to each other. In fact, the very openness of XML invites a proliferation of DTDs, and consequently, the problem of interoperability compounds itself.

This produces practical and commercial problems. In the book publishing and manufacturing supply chain, different links in the chain use different standards: typesetters, publishers, booksellers, printers, manufacturers of electronic rendering devices and librarians. This disrupts the digital file flow, hindering supply chain integration and the possibilities of automating key aspects of supply chain, manufacturing and distribution processes. Precisely the same practical problems of interoperability are now arising in other areas of the electronic commerce environment.

Although our main interest is the world of authorship and publishing, the longer-term possibilities of technologies of interoperability such as CGML are in the areas in which the semantic web has so much – as yet unfulfilled – promise. This includes: indexing, cataloguing and metadata systems; product identification systems; systems for the production, manufacture and distribution of copyright digital content; knowledge and content management systems; systems for multichannelling content and also providing for disability access; machine translation from one natural language to another; and artificial intelligence.

More practically, the challenge of interoperability is this: in a scenario where there are many more than two parties, where the information is not covered by a single standard, where the resources and skills of the parties cannot facilitate costly and time-consuming integration, an approach is needed which caters for the complexity of the messages, while providing tools which

simplify the provision and extraction of data and metadata. This is the crux of semantic interoperability. Such an approach involves providing a systematic mapping of associated XML standards to a common XML 'mesh', which tracks semantic overlays and gaps, schema versioning, namespace resolution, language and encoding variances, and which provides a comprehensive set of rules covering the data transfer – including security, transactional and messaging issues.

The idea of a 'meta-schema' – a schema to cover other related schemas – was initially considered to be sufficient. Research has demonstrated, however, that this is not enough, being subject to many of the same problems as the individual schemas being mapped – versioning, terminological differences and so on.

The core operational principles of CGML are outlined in the following subsections: meaning form or rendering is rigorously separated from, yet reliably follows, markup tags expressing meaning function; interoperability of tagging schemas can be achieved by mapping though an interlanguage governed by a set of semantic and structural rules; a tag schema expresses paradigmatic relations; a tag thesaurus expresses relations between tagging schemas; a tag dictionary expresses semantics; interoperability mechanisms are automated or semi-automated; and tag narratives anticipate a range of activity sequences driven by funnels of commitment and realised through alternative navigation paths.

As a terminological aside, CGML deliberately recruits some quite ordinary concepts from the world of textual meaning, such as the ideas of a dictionary, thesaurus and narrative. If we are going to have text-made text, we might as well use these historically familiar devices – albeit with some refinement and levels of precision required by the logistics of digital meaning.

Form Follows Function

CGML provides a framework for marking up and storing text as structured data. The storage medium can be XML files. Or it can be a database in which fields are named by tags, and from which exports produce XML files marked up for structure and semantics, ready for rendering through available stylesheet transformations. The result is text that is more easily located by virtue of the clarity and detail of metadata markup, and capable of a range of alternative renderings. CGML structures and stores data on the basis of a functional grammar of text, not just as object but as a process of collaborative construction. The focal point of CGML is a functional grammar of the book as text, as well as a kind of grammar (in the metaphorical sense of generalised reflection) of the social context of book publishing. However, with CGML, 'functional' takes on a newly active meaning. The markup actually manufactures the text in the moment of rendering, through the medium of stylesheet transformation.

What does it mean to take a functional approach to meaning? CGML's perspective on meaning is characteristic of modern functional linguistics, as distinct from the approach to meaning taken by the tradition of formal linguistics. The most common practice in western understandings of language and meaning has been to take as a point of departure the forms of meaning (words and how these words are arranged in syntax, and how these are realised as speech or writing). Formal linguistics traces elaborate structures and patterns amongst these forms. In this endeavour 'grammar', or the arrangement of words, is the primary interpretative device. Only after the structure of forms has been established is the question posed, 'what do these forms mean?' In contrast, functional linguistics turns the question of meaning around the other way: 'how are meanings expressed?' Language is conceived as a system of meanings; its role is to realise or express these meanings. It is not an end in itself; it is a means to an end (Halliday, 1994). Meaning function underlies meaning form. An account of meaning form must be based on a functional interpretation of the structures of meaning. Meaning form of a linguistic variety comprises words and their syntactical arrangement, as well as the expressive or presentational processes of phonology (sounding out words or speaking) and graphology (writing). Meaning form needs to be accounted for in terms of meaning function.

In CGML, as is the case for any digital markup framework that separates structure and semantics from presentation, the elementary unit of meaning function is marked by the tag. The tag specifies the meaning function for the most basic 'chunk' of represented content. Tags, in other words, describe the meaning function of a unit of content. For instance, a word or phrase may be tagged as <Emphasis>, <KeywordTerm> or <OtherLangaugeTerm>. These describe the peculiar meaning function of a piece of content. In this sense, a system of tags works like a partial functional grammar; they mark up key features of the information architecture of a text. Tags delineate critical aspects of meaning function, and they do this explicitly by means of a relatively consistent and semantically unambiguous meta-language. This meta-language acts as a kind of running commentary on meaning functions which are otherwise embedded, implicit or to be inferred from context.

Meaning form follows mechanically from the delineation of meaning function, and this occurs in a separate stylesheet transformation space. Depending on the stylesheet, each of the three functional tags <Emphasis>, <KeywordTerm> and <OtherLangaugeTerm> may be rendered to screen or print either as boldface or italics, or as a particular intonation in the case of rendering as synthesised voice. Stylesheets, incidentally, are the exception to the XML rule strictly to avoid matters of presentation; meaning form is their exclusive interest.

A set of tags constitutes a controlled vocabulary of meaning functions for a particular field. The semantics of each tag is defined with as little ambiguity as possible in relation to the other tags in a tag schema. Insofar as the tags relate to each other – they are indeed a language – they can be represented by means

of a tag schema making structural connections (a <Person> is named by <GivenNames> and <Surname>) and counter distinctions against each other (the <City> Sydney as distinct from the <Surname> of the late eighteenth-century British Colonial Secretary after which the city was named). Schemas define tags paradigmatically. In today's computer parlance, this paradigmatic domain is increasingly being called 'ontology'. This is the basis of the futuristic projects of the semantic web, machine translation and artificial intelligence. Finally, tags are put to practical representational and communicative use in activity sequences, or narratives.

Returning to the question of ontology briefly addressed earlier in this article, the relation of the tag to tagged content requires further clarification. The form is the represented or communicated content with all its presentational nuances. The tag represents the meaning function of which that form is an expression; it is itself invisible, but, via the inference rules built into stylesheets it influences the manner of presentation of the content. To explore the relation of tag to reality, we will take Kant's famous example of the willow and the linden tree, and express it the way an XML tagging schema might. We could mark up these words semantically as <tree>willow</tree> and <tree>linden tree</tree>. The tagging may have a presentational effect if these terms need highlighting, if they appear as keywords in a scientific text, for instance; or it may assist in searching. This markup tells us some things about the structure of reality, and with its assistance we would able to infer that a <tree>beech</tree> falls into the same category of being (ontology). Our controlled markup vocabulary comes from somewhere in the field of commonsense biology. In that field, a <tree> is but one instance of a <plant>, and a plant in turn is but one instance of a <lifeform>. We could represent these structural connections visually by means of a taxonomy. However, <tree> is not an unmediated element of being; rather, it is a category of being. How do we create this tag category? How do we come to name the world in this way?

Here is Kant's answer:

> 'I see, for example, a willow and a linden tree. By comparing these objects, first of all, I note they are different from each other with regard to the trunk, branches, leaves etc.; but then, on reflecting only upon what they have in common: the trunk, branches and the leaves themselves, and by abstracting from their size, their shape, etc., I obtain the concept of a tree.' What follows is a process Kant calls 'the legislative activity of the intellect'. From the intuition of trees, the intellect creates the concept of tree. '[T]o form concepts from representations it is ... necessary to be able to compare, reflect, and abstract; these three logical operations of the intellect, in fact, are the essential and universal conditions for the production of any concept in general.' (Quoted in Eco, 1999)

Trees exist in the world. This is unexceptionable. We know they exist because we see them, we name them, we talk about them. We do not talk about trees

because they are a figment of pure imagination, the result of a purely capricious act of naming. There is no doubt that there is something happening, ontologically speaking. However, we appropriate trees to thought, meaning, representation and communication through mental processes which take the raw material of sensations and from these construct abstractions in the form of concepts and systems of concepts or schemas. These concepts do not transparently represent the world; they represent how we figure the world to be.

And how do we do this figuring? When we use the concept 'tree' to indicate what is common to willows, linden trees and beeches, it is because our attention has been fixed on specific aspects of apprehended reality – what is similar (though not the same) between the two trees, and what is different from other contiguous realities, such as the soil and the sky. But equally, we could have fixed our attention on another quality, such as the quality of shade, in which respect a tree and a built shelter share similar qualities.

Tags and tag schemas build an account of meaning function through mental processes of abstraction. This is by no means an ordinary, natural or universal use of words. Vygotsky and Luria make a critical distinction between complex thinking and conceptual thinking. Complex thinking collocates things that might typically be expected to be found together: a tree, a swing, grass, flower beds, a child playing and another tree – put together, the young child learns to call these a playground. From the point of view of consciousness and language, the world hangs together through syncretic processes of agglomeration. A playground is so named because it is this particular combination of things. The young child associates the word 'playground' with a concrete reference point. Conceptual thinking also uses a word, and it is often the same word as complex thinking. However, its underlying cognitive processes are different. Playground is defined functionally, and the word is used 'as a means of actively centring attention, of abstracting certain traits, and symbolising them by the sign' (Vygotsky, 1962; Luria, 1981; Cope & Kalantzis, 1993). Then, beyond the level of the word-concept, a syntax of abstraction is developed in which concept relates to concept. This is the basis of theoretical thinking, and the mental construction of accounts of a reality underlying what is immediately apprehended, and not even immediately visible (Cope & Kalantzis, 1993). And the way we construct the world mentally is not just a product of individual minds; it is mediated by the acquired structures of language with all its conceptual and theoretical baggage – the stuff of socialised world-views and learned cultures.

Conceptual thinking represents a kind of 'reflective consciousness' or meta-consciousness. Markup tags are concepts in this sense and tag schemas are theories which capture the underlying or essential character of a field. When applied to the particularities of a specific piece of content, they work as a kind of abstracting meta-commentary, relating the specifics of a piece of content to the generalised nature of the field.

Markup tags do not reflect reality. Nor do they represent it comprehensively. Rather, they highlight focal points of attention relevant to a particular expressive domain or social language. In this sense, they are world-views or cultural artefacts. A tag does not exhaustively define the meaning function of the particular piece of content it marks up. Rather, it focuses on a domain-specific aspect of that content, as relevant to the representational or communicative purposes of a particular social language. In this sense, to reiterate, 'schema' is a more accurate and useful concept than 'ontology'.

Notwithstanding these reservations, there is a pervasive underlying reality, an ontological grounding which means that schemas will not work if they are mere flights of fancy. Eco (1999) characterises the relationship between conceptualisation and the reality to which it refers as a kind of tension. On the one hand, 'being can be nothing other than what is said in many ways ... every proposition regarding that which is, and that which could be, implies a choice, a perspective, a point of view. ... [O]ur descriptions of the world are always perspectival, bound up with the way we are biologically, ethnically, psychologically, and culturally rooted in the horizon of being'. But, contrary to the assumptions of the prevailing philosophies of postmodernism, this does not mean that 'anything goes'. 'We learn by experience that nature seems to manifest stable tendencies. ... *[S]omething* resistant has driven us to invent general terms (whose extension we can always review and correct). The world can never be simply a figment of our concept-driven imaginations. Even granting that the schema is a construct, we can never assume that the segmentation of which it is the effect is completely arbitrary, because ... it tries to make sense of something that *is there*, of forces that act externally on our sensor apparatus by exhibiting, at the least, some resistances' (Eco, 1999).

Interlanguage

Markup schemas or software tagging systems use a variety of encoding formats, including Extensible Markup Language (XML) and Resource Definition Framework (RDF). They promise to overcome two of the most serious limitations of the World Wide Web: the fact, first, that searching is simply for semantically undifferentiated strings of characters; and second, the fact that rendering alternatives are mostly limited by data entry methods – printed web pages do not live up to the historical standards of design and readability of printed text, and alternative non-visual renderings, such as digital talking books, are at best poor.

Specific schemas are designed to provide more accurate search results than is the case with computer or web-based search engines. Examples include the Dublin Core Metadata Framework and MARC electronic library cataloguing system. However, metadata harvested in one scheme cannot be readily or effectively be used in another.

Specific schemas are also designed for a particular rendering option. For instance, amongst schemas describing the structure of textual content, HTML

is designed for use in web browsers, DocBook for the production of printed books, Open eBook for rendering to handheld reading devices and Digital Talking Book for voice synthesis. Very limited interoperability is available between these different schemas for the structure of textual data, and only then if it has been designed into the schema and its associated presentational stylesheets. Furthermore, it is not practically possible to harvest accurate metadata from data, as data structuring schemas and schemas for metadata are mutually exclusive.

Figure 1. CGML as an interlanguage.

The field of the semantic web attempts to improve the inherent deficiencies in current digital technologies both in the area of resource discovery (metadata-based search functions) and rendering (defining structure and semantics in order to be able to support, via stylesheet transformations, alternative rendering options).

CGML attempts to interrelate the principal extant in the tag schemas for the world of authorship and of publishing. However, unlike other tag schemas in this domain, it does not purport to be ontologically grounded. It does not attempt to name or rename the world. Rather, CGML builds a common ground between contiguous and overlapping tag schemas which already

purport to name the world of authorship and publishing. It is not a language. It is an interlanguage.

The challenge of interoperability of tagging schemas (standards, or application profiles, or namespaces) has typically been addressed through schema-to-schema 'crosswalks'. A crosswalk is a listing of tag-to-tag translations not dissimilar to a language-to-language dictionary. For instance, as mentioned earlier, crosswalks have been created between MARC and ONIX (MARC Standards Office, 2000) and between MARC and Dublin Core (MARC Standards Office, 2001). As Paskin notes, when there are N schemas, (N/2)(N–1) mappings are required (Paskin, 2003). For instance, as of writing, CGML maps to 17 schemas. For full interoperability, 136 crosswalk mappings would be required. Or, to take a natural language analogy, if there are 60 languages in Europe, translation between all 60 languages can be achieved with 1770 language-to-language dictionaries – Italian-Gaelic, Gaelic-Vlach, Vlach-Italian, etc.

In fact, things are more complicated even than this. Each dictionary is, in fact, two dictionaries. Italian-Gaelic and Gaelic-Italian are not mirror inversions of each other because each language frames the world in its own semantically peculiar way. Similarly, the MARC to ONIX exercise (MARC Standards Office, 2000) is quite a different one to the ONIX to MARC exercise (EDItEUR, 2003). MARC to ONIX translates a library cataloguer's understanding of the nature and content of the book into a form intelligible to a publisher or a bookseller; and ONIX to MARC translates a publisher's or bookseller's understanding of the book into a form intelligible to a library cataloguer. In each case, the frame of reference or the starting point is defined in terms of a subtly distinctive social language. Each crosswalk is a quite separate intellectual and discursive exercise. So, we need to modify Paskin's crosswalk formula as follows: the number of mappings to achieve interoperability between N tagging schemas is $2\{(N/2)(N–1)\}$. In a terrain encompassed by the current version of CGML, 272 crosswalks would be required; Europe needs 3540 dictionaries for comprehensive cross-translation of all its languages. (And, while we're on this train of thought and although it is tangential to our point, cross-translation of all the world's estimated 6000 languages would require a practically impossible 17,997,000 dictionaries.)

Creating a single crosswalk is a large and complex task. As a consequence, the sheer number of significant overlapping tagging schemas in a domain such as authorship and publishing presents a barrier to achieving interoperability – and this without taking into account the fact that the schemas are all in a state of continuous development. Every related crosswalk needs to be reworked with each new version of a single tagging schema. Moreover, new tagging schemas are regularly emerging and every new schema increases the scale of the problem exponentially. Five cross-translations require 10 crosswalks; 10 cross-translations require 90 crosswalks.

Figure 2. Language pairs – full interoperability of 17 schemas requires 272 crosswalks.

Paskin suggests that this level of complexity can be eased by mapping 'through a central point or dictionary' (Paskin, 2003). This is precisely the objective of CGML. CGML is an intermediating language, or an interlanguage through which a full set of translations can be achieved. Tag by tag, it represents a common ground between tagging schemas. Tag <x> in the tagging schema A translates into tag <q> in CGML, and this in turn may be represented by <y> in tagging schema B and <z> in tagging schema C. The 'common-ground' tag <q> tells us that <x>, <y> and <z> are synonyms. A theoretical 272 crosswalks are replaced by 17 thesauri of tag synonyms. If, by analogy, all European languages were to be translated through Esperanto, a language deliberately fabricated as a common-ground language, 60 dictionaries would be needed to perform all possible translation functions instead of a theoretical 3540. Even simpler, in theory just one dictionary would suffice, translated 60 times with 60 language-to-Esperanto thesauri. This is precisely what CGML does. It attempts to solve the semantic aspect of the interoperability problem by creating one dictionary and 17 thesauri of tag synonyms. (And, incidentally, returning to natural language for a moment, this technique can be used as a semantic basis for machine translation, bringing the inter-translatability of all human languages at least into the realm of possibility.)

CGML has a number of distinguishing features which mean that it is constitutionally a very different kind of tagging schema to all the others against which it maps. It is this constitutional character that defines it as interlanguage, as distinct from a language.

An interlanguage has no life of its own, no independent existence, no relation to reality other than a mediated relationship through other languages. We will outline the operational principles for the construction of such an interlanguage through the subsequent subsections of this article.

Figure 3. The interlanguage approach – full interoperability between 17 schemas requires a thesaurus with just 17 sets of tag synonyms.

Before this, however, we want to mention some of the unique characteristics of an interlanguage such as CGML. As an interlanguage, CGML is designed to be open to the possibility of mapping of new schemas that may emerge within or substantially overlap its general domain. It is also designed to be able to absorb tagging which finely distinguishes conceptual subsets of certain of its core interests. In the case of authorship and publishing this might include, for instance, geospatial tags to define precise location, or tags representing controlled subject vocabularies in specific field-domains. By comparison with the crosswalk alternative, this mapping is achieved with relative ease.

Full subsumption and overlap are both cases of vertical integration of tagging schemas into CGML. However, CGML is also designed to be amenable to horizontal integration of schemas designed to define contiguous or complementary domains, such as the integration of other digital media or museum objects with the world of books. After all, books are routinely made

into movies, bookstores sell DVDs and printed books and libraries store individual copies of rare and unique books as artefacts.

As an interlanguage, CGML is infinitely extensible, absorbing subsidiary, overlapping and contiguous schemas to the extent that seems necessary and useful. A the time of writing, CGML consists of nearly a thousand tags – and these happen to be the tags for which there is the greatest degree of semantic common ground identifiable as synonyms across the interlanguage. The tags which represent the greatest degree of overlap also happen to be the most fundamental to the representational and communicative activities of authorship and publishing. However, there is no reason why CGML should not extend to ten thousand or a hundred thousand tags as it describes progressively more arcane bywaters in each tagging domain (vertical integration) or as it spreads its range of reference into contiguous domains of meaning (horizontal integration).

To reiterate, CGML is an interlanguage which maps against any other related schema (or, as they are variously termed, standards, namespaces or application profiles) in the domain of authorship and publishing. It works through tag-to-tag translation links between schemas – be they competing within a substantially overlapping domain or serving varied functions in divergent but still overlapping domains. The CGML term is an intermediary or interlanguage term. CGML is not a schema in and for itself. Rather, it is a way of talking to schemas.

The conventional approach to evaluating the efficacy of alternative tag schemas with a particular semantic domain is to undertake a process of comparison and contrast, the purpose of which is to select the one which would, it seems for the moment at least, be most appropriate to one's expressive needs, or the one that appears to be the most internally coherent and robust, or the one that happens to be most widely used amongst the players within a particular community of practice.

As an interlanguage, however, CGML is entirely agnostic about the ontological validity of the schemas to which it maps. If they move a community of practice, or have the potential to move a community of practice, they are worth the trouble of mapping. New standards may emerge, and if they appear to be sufficiently cogent and practically useful, they are also worth the trouble.

CGML itself does not set out to be a competing or alternative standard. Rather, CGML takes the approach that the prevailing uncertainty about which standards will predominate and the likelihood of the emergence of new standards, is to a significant degree a diversion. In the interlanguage approach, standards are everything – CGML needs to talk with the main existing and emerging publishing standards from the pragmatic point of view of interoperability. Yet, in another sense, standards are nothing – it is immaterial if some standards fall into desuetude or if new standards emerge. Dogmatic debate about the value or lack of value of a particular schema or standard is of little value. Shoehorning social practices into ill-fitting received standards is also

a fraught exercise. CGML cares about standards but eschews standardisation, or making things the same for the sake of sameness.

Our decision to take the interlanguage approach, paradoxically in the light of our scepticism about the ontological pretensions of tag schemas, is based on the stability inherent to the semantic ground, or a kind of ontological pragmatism. Behind the varied 'takes' on reality reflected by tag schemas, there is still a relatively stable and thus predictable material and social reality. The 'resistances' of which Eco speaks are frequently insistent. Although we conceptualise the world paradigmatically through tag schemas and operationalise these schemas through activity narratives, these paradigms and narratives do have a reference point, and this reference point is essentially a matter of ontology. As we have said, ontology does not simply present itself; it is mediated by paradigms and narratives. However, ontology practically grounds paradigm and narrative. In fact, through language, paradigm and narrative make themselves integral to the ontological reality of society or culture.

This grounding provides stability and thus a certain predictability of paradigm and narrative within a particular semantic domain. If authorship and publishing is our domain of interest, for instance, this represents a set of social practices – practices of representation and communication – that have deep and only gradually changing roots. There are authors who write; these authors have names; their writings have titles; and these writings have characteristic generic structures and fields of representation or subjects. Any new tagging schema that turns up – no matter how fancy it is or how innovative its intentions and methodology (e-learning, digital rights management, variable rendering and the like) – is still going to have to name these insistent realities.

The basis of CGML, in other words, is in the semantic ground of publishing, and there is an essential stability in the everyday life-world of authorship and publishing. The technologies may be changing, but there are still creators (writers, editors, illustrators) creating works (books in print and electronic formats, chapters, articles and other written, visual and audio texts) which are subject to copyright agreements, which are then used by consumers (readers, learners). Schemas do no more than represent that life-world from a particular perspective – be that the perspective of the library, digital resource discovery, rights, commerce, education or rendering/production. Schemas may come and go, but the life-world they purport to represent and facilitate remains relatively stable. At most, it changes incrementally.

The interlanguage approach of CGML also provides a tool for literature, science and curricula to be built in small languages and endangered languages, including, with the aid of Unicode, publication in any script. CGML can run in any language and any script, and this is achieved simply by translating the tags and tag definitions. This may seem a relatively small move in a practical sense. Conceptually, however, it is a huge move. In fact, it turns a linguistically expressed term into a mere 'token' of a core concept that exists above and beyond any particular language. And an indirect effect of this move is to add

multilingual functionality to markup schemas which currently exist only in English. In addition, by virtue of its structural and semantic approach to markup, CGML could serve as an aid to effective and accurate human and machine translation. In other words, by these various means, CGML could literally find itself in the space of an interlanguage between various human languages.

In a globalised and multilingual world, Ron Scollon argues, social languages or discourses are more similar across languages than within languages (Scollon, 1999). The way academics write for their particular discipline, for instance, whether it is in English or Japanese, is similar in terms of the structure of their texts and the ways those texts describe the world. A structural and semantic framework for structuring text such as CGML, which includes elaborate structural and semantic markup linked to controlled keyword vocabularies, will work across languages once the tags and the specialist vocabularies are translated, and this is because the most important thing about the discourse does not sit inside a particular language. Text structured and rendered in this way may become the platform for multilingual, multi-script publishing in communities more and more defined by their social language (what they want to do in the world, as expressed in peculiar ways of communicating about the world) than by the accident of mother tongue.

Paradigm

These, then, are the core concepts and principles of CGML: tags fit into schemas and these schemas function as paradigms. Tags mark up the narrative flow of activity sequences around the construction of meanings, and the architectures of meaning characteristic of specific social languages. Tagged narratives represent meaning functions and, in the rendering process, form follows function.

CGML's field is the ongoing and now relatively stable historical tradition of the book. It provides an account of the internal information architecture of the book as well the world of books in the plural – the systematic ordering of books in the world of libraries and bookselling. It is a theory of text structure and the social world of creators, their creations, the relation of their creations to other creations in the world, and the referents in the world to which their creations refer. CGML has two primary forms of expression: a paradigmatic expression in the form of the Taxonomy of Authorship and Publishing (supported by a Dictionary and a Thesaurus) and an open framework for the construction of Authorship and Publishing Activity Narratives which link the CGML tag-concepts into activity sequences focused on products (the lifecycle of a work, for instance) or roles (the activity structures of authoring, publishing or browsing for a book, for instance) (Common Ground, 2003a, b, c).

In terms of current computer science terminology (and even though we might question the use of the term) CGML is an ontology (Denny, 2002). 'In philosophy,' say Berners-Lee et al (2001), 'an ontology is a theory about the

nature of existence, of what types of things exist; ontology as a discipline studies such theories. Artificial-intelligence and web researchers have co-opted the term for their own jargon, and for them an ontology is a document or file that formally defines the relations among terms. The most typical kind of ontology for the Web has a taxonomy and a set of inference rules.' In this specialised sense, computer science sense, CGML is an ontology – even though we would question the application of the word to computer science in the light of its philosophical connotations.

Represented as a taxonomy, CGML relates its thousand-odd tags into eight orders of concept, or eight levels linked by branch or parent-child relationship – whichever metaphor one might use to choose to describe taxonomy. As is required by XML expression languages, there is a single first-order concept or 'root element' (Harold & Means, 2002). This root element is <Meaning>. <Meaning> has two children: <Function> and <Form>. As CGML has little interest in <Form>, no children are noted, although children could be added if and when there appeared to be a need to develop a new account of the realm of presentation and stylesheet transformation. This realm is taken as given within the realm of <Form>. In CGML, this is a space where various existing stylesheet transformations can be applied as designed for the various structural and semantic tagging schemas with which CGML interoperates. We nevertheless include <Form> as one of our two second-order concepts because it is of fundamental importance. From a representational or communicative point of view, <Function> remains unexpressed without a material realisation as <Form>. <Function> has no practical existence without <Form>.

At a taxonomic third level <Function> splits into three: a <SemanticGround>, a process of <Creation>, and the means of <Distribution>. The <SemanticGround> consists at a fourth level of the activities of a <Party> (a <Person> or <Organisation> at the fifth level), in a specifiable <Location>, at or during a point of <DateAndTime> and a <Subject> indicting the material, social or metaphysical referent of the creative work, to which a reader or user's attention may be directed. The process of <Creation> consists at a fourth level of primary <Creator>s, ancillary <Contributor>s, whose creative efforts have an inherent <Design> (which at a fifth level becomes a <Work> and a sixth level becomes a <Product> such as, at a seventh level, a <Book> or a <Map> for instance). The third-level process of <Creation> may also involve ascribing a fourth level <Status> (such as <Proposal>, <Draft> or <Edition> at fifth level), providing a <Description>, noting the form of linguistic presentation in a natural <Language>, indicating <Relations> to encompassing or subsidiary <Works> or <Products>, naming a <Publisher>, defining <Rights>, ascribing a unique <Identifier> such as a product number or Digital Object Identifier, and describing <Format>. Still at a fourth level, the products of the <Creation> process have an inherent <Structure> or information architecture (covering everything from <MacroStructure> such as <Chapter> and <Index> and

<LocalTextStructures> down to the level of <Paragraph> or <Emphasis> for words or phrases). These are supplemented by <Externals> which refer to the <Work> in question, such as a <Review> or <RefereeReport>. The final third-level concept <Distribution> provides a framework for the tagging of <Audience> (who a <Work> is meant for), Availability (where and how it can be found), <Consumer> (who reads or uses it), <Item> (an individual manifestation of a <Product>), <Transaction> (the legal basis of a particular <Consumer> use), <Delivery> (how the <Item> reaches the <Consumer>) and <Provenance> (where the <Item> has been during its life). This is the beginning of a paradigm which currently runs to a thousand <Function>s within the field of <Meaning>, and whose main focus at this stage is the creative process of authorship and the publication of books.

Figure 4. CGML Taxonomy of Authorship and Publishing: first to fourth level concepts. The remaining of the approximately thousand tags add detail at the fifth level and beyond.

Within CGML, there are two types of tags: open tags and closed tags. Open tags mark up any content which they happen to enclose, for instance: <MainTitle>Any Conceivable Title</Title>. In the XML expression format, these are called 'elements'. Closed tags specify a strictly defined range of content alternatives, and these take the form of a predetermined list of secondary tags. For example, in CGML as it currently stands, <MeaningMode> can only be defined amongst the alternatives <LinguisticMode>, <VisualMode>, <AudioMode>, <GesturalMode>, <SpatialMode> and <Multimodal>. In the XML expression format, these are called 'attributes'.

Paradigm is constructed in CGML by means of a number of taxonomic construction rules. Although CGML tags are written in natural language, this

belies a level of precision not found in natural language. Natural language involves considerable semantic ambiguity, whereas a tagging schema needs to attempt to reduce this as much as practicable. It does this by rigorously applying two semantic logics that exist somewhat less rigorously in natural language: the logic of distinction-exclusion and the logic of relation-inclusion. The logic of distinction-exclusion exists with parallel branches (sibling relations) in a taxonomy. A <Person> is not an <Organisation> because an <Organisation> is defined as a legally or conventionally constituted group of <Persons>. On the other hand, the logic of relation-inclusion applies to the sub-branches that branch off superordinate branches in a taxonomy (parent-child relations). A <Party> to a creative or contractual relationship can be either a <Person> or an <Organisation>.

'Meaning', says Gee, 'is always (in part) a matter of intended exclusions and inclusions (contrasts and lack of contrasts) within an assumed semantic field' (1996). In natural language, we use rough-and-ready ways of working out whether another person means the same thing as we do by a particular word or phrase. One way is what Gee calls 'the guessing principle' – our judgment or 'call' on what a particular concept means. If we are in the same social, cultural or professional group or community of practice as the communicator of our particular concept, our guess is more likely to be congruent with the communicator's understanding. Another way is 'the context principle', or to add precision to the meaning of a work or phrase by deciphering it in the context of the text and social situation in which it appears (Gee, 1996).

Domain-specific paradigms in the form of tagging schemas are designed to reduce the guesswork and contextual inference required in natural language. The solution is to build a social language which clarifies the exclusions and inclusions. This is achieved in CGML by three overlapping visual and textual techniques: taxonomy, thesaurus and dictionary. Thesaurus and dictionary are the subjects of the next two subsections of this article.

Concentrating for the moment, however, on the general rules of taxonomy or paradigm formation, we need to make distinctions between taxonomic processes of superordination and composition (Martin, 1992). Superordination relations perform the function of sub-classification. They express an 'is a' relationship between one level in the taxonomic hierarchy and another. <Book> is a <Product>, as is an <AudioRecording>. Composition relations, by contrast, connect parts into wholes. They express a 'has a' relation between levels in the taxonomic hierarchy. A <GlossaryItem> and a <GlossaryItemDefinition> are both parts of a <Glossary>. Indeed, a <Glossary> is not functional without both of these parts.

To the superordination and compositional principles identified by Martin, we add the capacity of taxonomies to make a distinction of immanence. This expresses an 'underlies' relationship between contiguous levels in the taxonomic hierarchy. A <Design> underlies a <Work> and a <Work> underlies a <Product>. In CGML, <Design> has just one child, <Work>. However, <Design> and <Work> cannot be conflated even though there are

no multiple children with whom composition (part/whole) or sub-classification functions can be performed. A <Design> may encompass the full scope and essential character of a <Work>. This may be prefigured at the planning or <Proposal> stage. However, a <Design> may never become a <Work>. If it does, however, it does not disappear; rather, it is applied and adapted and remains immanent within the <Work>. Similarly, a <Work> such as the lyrics for a song, remains immanent within its various instantiations as a <Product>, such as a <Book> or an <AudioRecording>, or as a <Performance> at an <Event>. (This logic of immanence in a creative work builds upon, modifies and extends the entity-definition work of the International Federation of Library Associations [IFLA Study Group on the Functional Requirements for Bibliographic Records, 1998].)

Finally, taxonomies need to be cohesive if they are to provide an effective paradigmatic role for a field of practice. Such cohesion is created to a large degree by the proximity of concepts in contiguous levels in the hierarchy. Between one level and another, relations need to be tested to see whether a tag-concept on one level is experientially close enough to be presumed by a tag-concept on another (Martin, 1992). <PrintedBook> and <Design> are not experientially close concepts, and thus would not form a cohesive parent-child relationship. However, the <Design>, <Work>, <Product>, <Book>, <PrintedBook> hierarchy involves contiguous items sufficiently close in an experiential sense to ensure taxonomic cohesion.

Thesaurus

The CGML taxonomy maps synonymous concepts from related tag schemas.

In Figure 5, the CGML open-element tags are represented in green, and CGML fixed-attribute tags are represented in green. For each tag, synonyms are identified in the various tagging schemas against which CGML is currently mapped. The underlined concepts indicate levels of implementation. <Person> data, for instance, can only be collected in the smallest granular units required by any of the mapped tagging schemes. A valid CGML <Person> record (and the IMS, ONIX, XrML, indecs, EAD and MARC synonyms) can only be generated from data recomposed from smaller granular units including, for instance, <GivenNames> and <Surname>.

The CGML Thesaurus takes each tagging schema as its starting point, lists its tags and reproduces the definitions and examples as given by each tagging schema. In this sense, CGML actually works with 17 thesauri, and each new mapping will require an additional thesaurus. Each thesaurus captures the way in which each tagging schema defines itself, and within its own terms. Against each tag, a direct CGML synonym is provided, whose semantics are coextensive with, or narrower than, the tag against which the mapping occurs. Unlike a conventional thesaurus, only one CGML equivalent is given for each mapped tag.

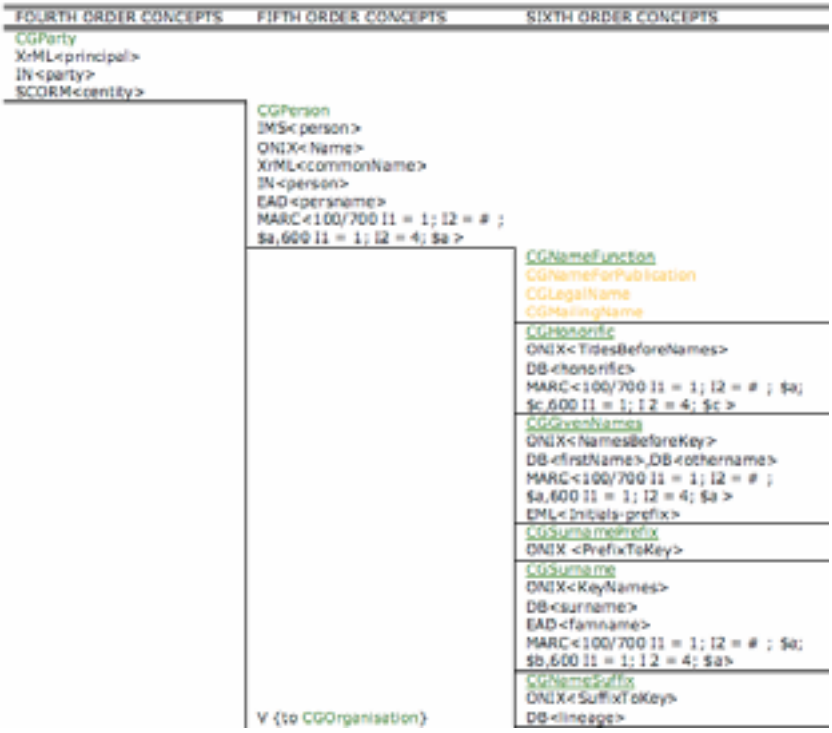

Figure 5. Fragment of the CGML Taxonomy of Authorship and Publishing specifying the concept of <Party> from the fourth to sixth levels.

In combination with its dictionary, CGML uses both what Martin identifies to be the two traditional approaches to the study of lexis in western scholarship: dictionary and thesaurus. Whereas dictionary 'purports to unpack the "meaning" [of lexical items] by means of paraphrase and exemplars', thesaurus is 'organised around meaning'; it 'purports to display the wordings through which meanings can aptly be expressed'. He concludes that '[b]ecause it is organised according to meaning, a thesaurus provides a more appropriate model of textual description for functional linguistics than a dictionary does' (Martin, 1992). In the case of CGML, an additional layer of rigour is added by mapping the 17 thesauri into the paradigm constituting taxonomy.

The effect of these crosscutting processes is the systematic mapping of existing and emerging tagging schemas against each other, and the stabilisation of synonyms between different markup languages through the medium of the CGML interlanguage tag. This has the potential to add functionality to existing schemas, not only by extension of new functionalities to otherwise separate schemas, but also by reinterpreting data created in one framework for

(unanticipated) use in another. CGML thus has the potential to form the foundation for a broker software system within the domain of authorship and publishing.

Figure 6. Fragment of the Dublin Core to CGML Thesaurus.

Practically, this means that CGML provides a simple, transparent, clearly defined natural-language tagging framework which will create data conforming to the schemas against which it is mapped. CGML data can be exported into any XML schema against which CGML has been mapped. The effect is to ensure interoperability between different data collection practices and frameworks – so, for instance, data collected with a CGML-defined framework, can simultaneously become a MARC library catalogue record and an ONIX record for a B-2-B ecommerce transaction. The reverse is only partly the case. Data formatted in any XML namespace against which CGML has been mapped can be imported into a CGML-defined database, and from this it can be exported into XML namespaces other than the one for which the data were originally defined, but only when those data enter CGML at the level of granular delicacy required by the most delicately granular schema against which CGML has been mapped (identified by underlined tags, as illustrated in Figure 5). When a more granular markup is required for interoperability than is available in imported data, this will usually have to be created manually – for example, breaking a <Person>'s name into <GivenNames> and <Surname>, part of which process will involve the complex and highly contextual business of interpreting whether the <Person>'s name appears in English or is structured in the traditional Chinese way.

Dictionary

Fairclough points out that 'it is of limited value to think of a language as having a vocabulary which is documented in "the" dictionary, because there are a great many overlapping and competing vocabularies corresponding to different

domains, institutions, practices, values and perspectives' (Fairclough, 1992). Gee (1996) calls these domain-specific discourses 'social languages'. For instance, we struggled earlier with the concept of 'ontology' central to the notion of the 'semantic web'. We did not think the word was quite right for its application to digital text markup practices, and explained why we preferred the word 'schema'. This is because we were taking seriously the roots of the word in the discourse of philosophy. However, the word may be perfectly fine to describe an approach to digital text markup, but only if we agree that it means something slightly, but importantly, different to what it does in philosophy. It is these slight but important differences that create ambiguity. Such ambiguity needs to be reduced as much as possible in tagging schemas (or, if you like, 'ontologies').

The dictionary solution to the problem of ambiguity is to list the major alternative meanings of a word, although this can only reflect gross semantic variation. No dictionary could ever capture comprehensively the never-ending subtleties and nuances ascribed differentially to a word in divergent social languages.

Dictionary, nevertheless, is one of the fundamental cross-referential tools of CGML, sitting alongside and integrated with the devices of paradigm and thesaurus. However, it is a dictionary of a peculiar kind. It is more like a glossary than a dictionary. In fact the CGML notion of dictionary is best defined by how it is different from a natural language dictionary. For a start, the CGML dictionary does not purport to be about external referents as 'meaning'; rather, it is built via the interlanguage technique from other languages which purport to have external referents. Moreover, insofar as the semantic ground of CGML is meaning itself (and its instantiation in the practices of authorship and publishing), it is a kind of meta-semantics, a language of meaning. It happens to be centred on the realm of semantics in general – the meaning of meaning – and within that realm the social practices and technologies of representation and communication stabilised in the historical tradition of the book.

Furthermore, CGML is not an ordinary dictionary insofar as it develops a specialised 'take' on the world it purports to describe, the world of book meaning. Its meanings are not the commonsense meanings of the life-world of everyday experience, but derivative of specialised social languages which speak in the refined and particularistic way characteristic of the professionals and aficionados of that domain. To apply a pair of concepts of Husserl's, commonsense language in shifting and ambiguous language of the life-world, social languages develop progressively refined (sedimented) and self-consciously reflective (bracketed) discourse more characteristic of science (Husserl, 1970; Cope & Kalantzis, 2000a, b). CGML, in other words, derives from schemas developed in and for professions which have developed high levels of conceptual clarity about what authorship is and what publishing involves.

The CGML *Dictionary of Authorship of Publishing* links a notation (the tag-concept), which may be used in practice as a label for a field in a database or as an XML tag, to a semantically explicit definition, as well as an annotation which explains and exemplifies the tag-concept in terms of subordinate tag-concepts in the taxonomy (the various logics of relation-inclusion discussed earlier), and provides advice where necessary on appropriate and well-formed data entry. The building blocks of the CGML Dictionary are the other tag-concepts of the CGML schema, and these are connected by hyperlinks. The definition builds on parent tag-concepts; the annotation suggests the possible instantiations of a tag-concept by means of illustrative child tag-concepts. The Dictionary is located in maintained in a purpose-built software application, *CommonGroundLEXICOGRAPHER*.

NOTATION Tag-Concept.	DEFINITION Dictionary definition based on parent concepts.	ANNOTATION Examples, related child concepts, data entry advice.
CDCreation	The process of realising Meanings which serve particular communicative or representational Functions.	Involves Creators and Contributors, developing a Design, which is expressed as a Work and is manifest in the form of a Product, Event or Service. A Design has a Title and may vary in Status (as it evolves, for instance, from Proposal, to Draft, to Edition). It may have a Description which summarises or overviews its content. In the case of text, it will be expressed in a particular Language. It may sit in a defined Relation to other Designs, Works or Products, and this may be specified as a Source. Distribution may be the responsibility of a Publisher, on the basis of an assignment or licensing of Rights. It may have an Identifier in the form of a unique string of numbers and/or letters such as an ISBN or a URL. It will have a Format and an inherent and characteristic Structure, and also possibly Extremals which exist specifically in order to relate to the Work, for example by way of Review or Promotional Matter.
CDCreator	A Person who, or an Organisation which, plays a primary role in the conceptualisation of a creative Design and the execution of a Work. An essential Party to the process of Creation.	Includes the roles of Author, Editor, Scriptwriter, Lyricist, Composer, Visual Artist, Photographer and Conference Presenter.

Figure 7. Fragment of the CGML Dictionary of Authorship and Publishing specifying the concepts of <Creation>and <Creator>.

The Dictionary has been constructed using five semantic rules: minimised ambiguity; functional clarity; lowest common denominator semantics; the distinction of silent from active tag-concepts; and comprehensive internal cross-reference.

First rule of dictionary formation: minimised ambiguity. Digital expression languages such those captured by XML (of which CGML is an instance) use natural language tags in the interest of transparency. The appearance of natural language, however, simulating as it does everyday semantics, is deceptive. The further removed from everyday language a digital expression language is, the more effective it is likely to be. For instance, a <Work> may involve some very different kinds of 'editor', obscured by the ambiguity of that word in everyday parlance. CGML defines one kind of <Editor> as a primary <CreatorRole> in relation to a <Work> – a person who pulls together a number of texts by various <Author>s into a coherent work, and maybe writes

an introduction. From a presentational point of view, the <Editor>'s name will appear (via the stylesheet transformation process) on the cover and title page of a <Book>. This <Editor> is distinct from other types of 'editor', such as a <CommissioningEditor> – typically, a person who works for a <Publisher> and who instigates the process which leads to the <Publication> of a <Product>. <Editor> is also distinct from a <CopyEditor>, who identifies textual errors. These latter two kinds of people, frequently simply called 'editor' in everyday parlance, play a <ContributorRole> in the <Creation> process, and need to be clearly and unambiguously distinguished from an <Editor> who clearly and consistently has a <CreatorRole> in the process. In this way, the Dictionary draws explicit boundaries of distinction-inclusion between other tag-concepts, usually positioned as alternatives to each other at the same level in the taxonomy.

NOTATION Tag-Concept	DEFINITION Dictionary definition based on parent concepts	ANNOTATION Examples, related child concepts, data entry advice
CGEditor	A CreatorRole entailing the gathering together into a single, coherent Work, a number of Works by different Creators. Such a work may become a Product in the form of a Book or other, written Non-Book Text.	Not to be confused with a CommissioningEditor, a CopyEditor or a Publisher.

Figure 8. Fragment of the CGML Dictionary specifying the concept of <Editor>.

CGML attempts to achieve a balance between domain-specific concepts which are relatively free of jargon, and the precision characteristic of and necessary to technical and scientific discourses. Except when referring specifically to computers and computer-generated files, publishing terminology is preferred over computer terminology. For instance, <Edition> and <Draft> are preferred over 'version', not only for their familiarity to authors and publishers, but because they reflect an important distinction which is sometimes unclear in version enumeration.

In this process of removing ambiguity, at the furthest reaches of its taxonomic structure CGML may also absorb international standards and controlled vocabularies defining key features of the semantic ground such as ISO 3166 Territory Codes, ISO 4217 Currency Codes, ISO 639 Language Codes, ISO 8601 standard formats for the description of time, Unified Code for Units of Measure (International DOI Foundation, 2002).

Second rule of dictionary formation: functional clarity. The CGML Dictionary is not a description of things-in-themselves. Its purpose is functional – in a primary sense to provide an account of meaning functions, and in a secondary sense to provide a reliable basis for automated rendering through stylesheet transformation languages. Every definition and annotation explains in the first instance what an entity does, rather than what it is. Each tag-concept, moreover, can only do one thing. If a synonymous term in natural language

does more than one thing, as was the case of 'editor' in the previous subsection, a specialised distinction needs to be made explicitly.

Third rule of dictionary formation: lowest common denominator semantics. As discussed earlier, CGML's interlanguage approach means that it takes the common-ground position between broadly synonymous concepts in the tagging schemas against which it maps. Every CGML term or tag translates into an equivalent term in the various other schemas, if and where there is an equivalent. However, these concepts are not always the same. In the nature of social languages characterised by their own particularised 'take' on the world, tag-to-tag equivalents are often not true synonyms. This places a particular semantic burden on the intermediate, interlanguage term and its dictionary definition within CGML. In the case of tag synonyms with roughly equivalent but not identical semantics, CGML either takes the narrower definition in cases when one tag represents a subset of another; or in the case of overlap, creates a new definition restricted to the semantic intersection between the functional referents of the two equivalent tags. This guarantees that data will always be created from within CGML which can be validly exported as content into the database field or XML-tagged content spaces markup by equivalent tag synonyms within the mapped schemas.

The key to CGML's successful functioning as an interlanguage, in other words, is its dictionary definition and data entry rules. If the rule of lowest common denominator semantics is rigorously applied, all data entered within the framework of this definition and data entry rules will produce valid data for each of the standards in which a synonymous term exists. Each interlanguage term represents a semantic common ground – defined in terms which are sufficiently narrow and precise to produce valid data for the tag synonyms in all other standards to which a particular term can be validly mapped at that particular semantic point.

Fourth rule of dictionary formation: distinguishing silent and active tag-concepts. Although certain tag-concepts in CGML map against others successfully using the rule of lowest common denominator semantics, they cannot in practice be implemented at this level because they do not have a sufficient level of semantic delicacy to allow interoperability with schemas that require greater semantic delicacy than is possible at that level. Returning to the example provided in Figure 5, data cannot be entered at the CGML <Person> level even though that would be sufficient for certain schemas against which it is possible to map synonymous <Person> tag-concepts. Data entry must be broken up into the various name elements at the finest level of delicacy required by all of the mapped tag schemas (active tag-concepts, underlined in Figure 5); it can then automatically be recomposed to create valid data to populate the silent tag-concepts. Some of these silent concepts are purely theoretical. There will be very little practical need to 'climb out' to many of the highly abstracted first (root element), second and third-level concepts. Indeed,

some of them are well-nigh useless in a practical sense. Their role is purely to provide an overall system and structure to the schema.

Fifth rule of dictionary formation: comprehensive internal cross-reference. The key to building a resilient and functionally efficient tagging schema is to develop an interlocking system of cross-reference. This is rendered in the CGML dictionary as hyperlinks. Every hyperlinked tag-concept in the dictionary definitions and annotations takes the user to a precise definition and annotation of that tag-concept. Cumulatively, the dictionary definitions and annotations build a systematic account of relations of relation-inclusion and distinction-exclusion, providing descriptive content to the abstract visual representation of paradigm in the taxonomy. The result is that the schema becomes less like a selection of concepts that seem useful to a domain, and more like a theory of that domain.

Crosstalk

Common Ground Markup Language is a Document Type Definition (DTD), but of a peculiar kind. In fact, although it is technically a DTD, it is a DTD of a fundamentally different order to any other. It does not have an independent life as a DTD. Rather, its semantic life is derived solely from other DTDs and whose operational realisation is found within other DTDs. It is the product of crosstalk between DTDs (Common Ground, 2003a, b, c).

This adds another fundamental layer to the bifurcation of DTDs representing structure and semantics and DTDs representing rendering or presentational alternatives (stylesheets). The interlanguage mechanism creates a DTD which does not manage structure and semantics per se; rather, it automatically manages the structure and semantics of structure and semantics. Its mechanism, in other words, is meta-structural and meta-semantic. It is aimed at interoperability of schemas which purport to describe the world rather than immediate reference to the world. We have named its underlying mechanism the 'interlanguage' apparatus. Although developed in the case of one particular instantiation of the problem of interoperability – for the electronic standards that apply to publishing – the core technology is applicable to the more general problem of interoperability characterised by the semantic web and electronic commerce.

By filtering standards that don't talk to each other through the interlanguage mechanism for database and document tagging, a crosstalk mechanism is created that allows conversation and information interchange between unrelated schemas. This produces immediate supply chain efficiencies through the automated transition of digital content from one electronic standard to another. It also provides for the multipurposing of digital content, so that data are fully interoperable across all the full range of functional uses possible in the digital production and transmission of content. Three such applications for this technology are publishing, conference and learning

management software products – three areas in which Common Ground is actively involved. There are many others, within and outside of the domain of textual content.

CGML, in other words, sets out to create functionalities for data framed within the paradigm of one schema which extend well beyond those originally conceived by that schema. It sets out to facilitate interoperability between schemas, allowing data originally designed for use in one schema for a particular set of purposes to be used in another schema for a different set of purposes.

The interlanguage mechanism, in other words, means that metadata are newly created through its apparatus to be interpolated into any number of metadata schemas. It also provides a method by means of which data harvested in one metadata schema can be imported into another. From a functional point of view, some of this process can be fully automated, and some the subject of automated queries requiring a human-user response. Common Ground's current research involves the development and testing of the interlanguage mechanism for interoperability.

The interlanguage mechanism, in sum, is designed to function in two ways:

1. For new data, a filter apparatus provides full automation of interoperability on the basis of the semantic and syntactical rules built into the CGML schema.
2. For data already residing in an XML or RDF schema, data automatically pass through a filter apparatus using the interlanguage mechanism, and passed on into other schemas or ontologies even through the data had not originally been designed for the destination schema. The filter apparatus is driven by a set of semantic and syntactical rules as outlined below, and throws up queries whenever an automated translation of data is not possible in terms of those semantic rules.

The interlanguage apparatus is designed to be able to read tags, and thus interpret the data which have been marked up by these tags, according to two overarching mechanisms, and a number of sub-mechanisms. The two overarching mechanisms are the superordination mechanism and the composition mechanism – drawing in part here upon some distinctions made in systemic-functional linguistics (Martin, 1992).

- The superordination mechanism constructs tag-to-tag 'is a ...' relationships. Within the superordination mechanism, there are the sub-mechanisms of hyponymy ('includes in its class ...'), hyperonymy ('is a class of ...'), co-hyperonoymy ('is the same as ...'), antonymy ('is the converse of ...') and series ('is related by gradable opposition to ...').
- The composition mechanism constructs tag-to-tag 'has a ...' relationships. Within the composition mechanism, there are the sub-mechanisms of meronymy ('is a part of ...'), co-meronymy ('is integrally related to but exclusive of ...'), consistency ('is made of ...'), collectivity ('consists of ...').

These mechanisms are fully automated in the case of new data formation within any schema, in which case, deprecation of some aspects of an interoperable schema may be required as a matter of course at the point of data entry.

In the case of legacy data generated in schemas without anticipation of, or application of, the interlanguage mechanism, data can be imported in a partially automated way. In this case, tag-by-tag or field-by-field queries are automatically generated according to the filter mechanisms of:

- taxonomic distance (testing whether the relationships of composition and superordination are too distant to be necessarily valid);
- levels of delicacy (testing whether an aggregated data element needs to be disaggregated and re-tagged);
- potential semantic incursion (identifying sites of ambiguity); and
- translation of silent into active tags or vice versa (at what level in the hierarchy of composition or superordination data need to be entered to effect superordinate transformations).

The interlanguage mechanism is located in *CommonGroundLEXICOGRAPHER*, an ontology building tool developed by Common Ground. This piece of software defines and determines:

- database structures for storage of metadata and data;
- XML document inputs;
- synonyms across the tagging schemas for each standard against which CGML maps;
- two definitional layers for every tag: underlying semantics and application-specific semantics. In this regard, *CommonGroundLEXICOGRAPHER* creates the space for application-specific paraphrases which can be created for different user environments. The underlying semantics necessarily generates abstract dictionary definitions which are inherently not user-friendly. However, in an application such as *CommonGroundPUBLISHER*, each concept-tag needs to be described and defined in ways that are intelligible in a commonsense way within that domain. Each tag may need to be defined and exemplified somewhat differently in the case, for instance, of conference and learning environments, even though the underlying semantics of creators and works remains the same. It is these application-specific paraphrases that render to the application interface in the first instance;
- export options into an extensible range of electronic standards expressed as XML DTDs.

CommonGroundLEXICOGRAPHER, in fact, manages the superordination and compositional mechanisms described above, as well as providing an interface for domain-specific applications in which interoperability is required (such as publishing or learning management systems).

These transformations preformed by the *CommonGround-LEXICOGRAPHER* software are illustrated in Figure 9. Some of these are

already functional. Others are in the design stages. Various scenarios and exemplars follow.

Following are some examples of how this mechanism may function. In one scenario, new data might be constructed according to a source schema which has already become 'aware' by means of previous applications of the interlanguage mechanism as a consequence of the application of the mechanism. In this case, the mechanism commences with the automatic interpellation of data, as the work of reading and querying the source schema has already been performed. In these circumstances, the source schema in which the new data are constructed becomes a mere facade for the interlanguage, taking the form of a user interface behind which the processes of subordination and composition occur.

In another scenario, a quantum of legacy source data is provided, marked up according to the schematic structure of a particular source schema. The interlanguage mechanism then reads the structure and semantics immanent in the data, interpreting this both from DTD and the way the DTD is realised in that particular instance. It applies four filters: a delicacy filter, a synonomy filter, a contiguity filter and a subset filter. The apparatus is able to read into the DTD and its particular instantiation an inherent taxonomic or schematic structure. Some of this is automated, as the relationships of tags are unambiguous based on the readable structure of the DTD and evidence drawn from its instantiation in a concrete piece of data. The mechanism is also capable of 'knowing' the points at which it is possible there might be ambiguity, and in this case throws up a structured query to the user. Each human response to a structured query becomes part of the memory of the mechanism, with implications drawn from the user response and retained for later moments when interoperability is required by this or another user.

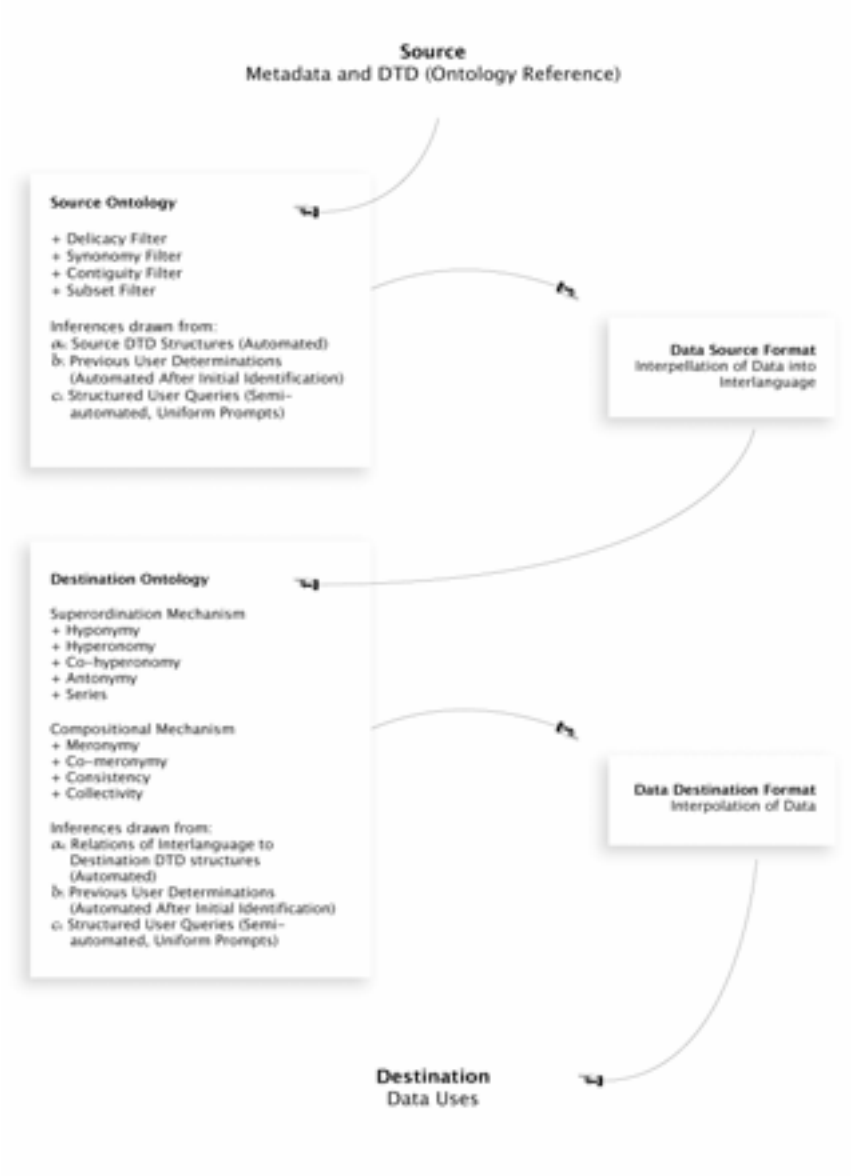

Figure 9. The interlanguage mechanism.

On this basis, the apparatus interpellates the source data into the interlanguage format, whilst at the same time automatically 'growing' the interlanguage itself based on knowledge acquired in the reading of the source data and source schema.

Having migrated into the interlanguage format, the data are then reworked into the format of the destination schema. It is rebuilt and validated according to the mechanisms of superordination (hyponymy, hyperonymy, co-hyperonomy, antonymy and series) and composition (meronymy, co-meronymy, consistency, collectivity). A part of this process is automated, according to the inherent structures readable into the destination schema, or previous human readings that have become part of the accumulated memory of the interlanguage mechanism. Where the automation of the rebuilding process cannot be undertaken by the apparatus with assurance of validity (when a relation is not inherent to the destination DTD, nor can it be inferred from accumulated memory in which this ambiguity was queried previously), a structured query is once again put to the user, whose response in turn becomes a part of the memory of the apparatus, for future use.

On this basis, the data in question are interpolated into their destination format. From this point, the data can be used in their destination context or DTD environment, notwithstanding the fact that the data had not been originally formatted for use in that environment.

Key operational features of this mechanism include:

- The capacity to absorb effectively and easily deploy new schemas which refer to domains of knowledge, information and data that substantially overlap (vertical ontology-over-ontology integration). The mechanism is capable of doing this without the exponential growth in the scale of the task characteristic of the existing crosswalk method.
- The capacity to absorb schemas representing new domains that do not overlap with the existing range of domains and ontologies representing these domains (horizontal ontology-beside-ontology integration).
- The capacity to extend indefinitely into finely differentiated sub-domains within the existing range of domains connected by the interlanguage, but not yet this finely differentiated (vertical ontology-within-ontology integration).

In the most challenging of cases – in which the raw digital material is created in a legacy DTD or ontology, and in which that DTD is not already known to the interlanguage from previous interactions – the mechanism:

- interprets structure and semantics from the source DTD and its instantiation in the case of the particular quantum of source data, using the filter mechanisms described above – for example, in the case of publishing and the Common Ground Markup Language interlanguage, a hypothetical newly introduced digital rights management framework;
- draws inferences in relation to the digital rights DTD and the particular quantum of data, applying these automatically and presenting structured queries in cases where the apparatus and its filter mechanism 'knows' that supplementary human interpretation is required;
- stores any automated or human-supplied interpretations for future use, thus building knowledge and functional useability of this new DTD into the

interlanguage – in this example, into Common Ground Markup Language. These inferences then become visible to subsequent users, and capable of amendment by users, through the *CommonGroundLEXICOGRAPHER* interface;
- interpellates the data into the interlanguage format, in this example Common Ground Markup Language;
- creates a crosswalk from Common Ground Markup Language into a designated destination DTD, for instance a new format for structuring or rendering text, using the superordination and composition mechanisms. These are automated in cases where the structure and semantics of the destination DTD are self-evident to the apparatus, or they are the subject of structured queries where they are not, or they are drawn from the *CommonGroundLEXICOGRAPHER* memory in instances where the same query has been answered by an earlier user;
- interpolates data into the destination format;
- supplies data for destination uses – in this instance, digital rights data applied to a new rendering format.

To give a less challenging example, the source DTD can be already known to the interlanguage, by virtue of automated validations based not only on the inherent structure of the DTD, but also many validations against a range of data instantiations of that DTD, and also numerous user clarifications of queries. In this case, the source DTD might be the e-learning metadata standard such as the UK National Curriculum, and the destination DTD might be a format for structuring learning texts, such as Educational Modelling Language. In this case:

- by entering data in an interface which 'knowingly' relates to an e-learning interlanguage, Learning Design Language, which has been created using the mechanisms of this invention, there is no need for the filter mechanisms nor the interpolation processes that are necessary in the case of legacy data and unknown source DTDs; rather, data are entered directly into the interlanguage format, albeit through the user interface 'facade' of the source DTD – in this case, the UK National Curriculum Standard;
- the apparatus then interpolates the data onto the designated destination format, in this case into Educational Modelling Language;
- the data can be used in the destination format, Educational Modelling Language.

It is possible to use the interlanguage apparatus to construct and apply other meta-markup languages which tie together other semantically overlapping languages of contiguous schemas or ontologies. Our examples here apply to the world of text, but there is no reason why these principles and mechanism, and the *CommonGroundLEXICOGRAPHER* tool, cannot be used to achieve interoperability in other domains.

Narrative

The two key engines of CGML are paradigm (represented visually as taxonomy, and textually by means of dictionary, supported by automated and semi-automated crosstalk mechanisms) and narrative. Paradigm systematically names the domain and interconnects the names; narrative strings these named things into sequences of meaning and action. Paradigm is made of nouns; narrative strings these nouns together with verbs.

Unlike CGML, some digital expression ontologies attempt to add verbs within their expression form, although within a very limited range. Resource Description Framework (RDF) is a highly abstracted propositional language for describing states of being and relationships between entities, typically in the form subject > property > object: person is author; consumer has book. It is made up of simple sentence-like information units, with a propositional verb connecting subject and object (W3C, 2003). Similarly, the *MPEG21 Rights Data Dictionary* uses verbs as well as nouns in its schema (Multimedia Description Schemes Group, 2002).

In CGML, all tags are nouns. CGML, moreover, does not define a noun simply as a kind of word. Rather, CGML tags are defined as lexical items, including pairs or groups of words which in a functional sense combine to form a noun, such as <CopyEditor>. Stringing nouns together into narrative happens in two ways: through the nominalisation of action (when <Author>s *create*, their activity is named the process of <Creation>), and through the necessary or implied actions that bring nouns into meaningful relationship (a <Proposal> means that an <Author> *might write* a <Book>; a <Draft> means that an <Author> *is writing* a <Book>; an <Edition> means that an <Author> *has written* a <Book>).

The first of these two ways of incorporating actions, nominalisation, is typical of specialised social languages. As Martin points out, 'one of the main functions of nominalisation is in fact to build up technical taxonomies of processes in specialised fields. Once technicalised, these nominalisations are interpretable as things' (Martin, 1992).

The second method is the process of creating activity sequences around the name-label tags. This is a process of constructing narrative. We will take a concrete example from CGML in action, and then step back to reflect on the role of narrative in building webs of active meaning from tag schemas.

CGML was originally developed as the basis for the *CommonGroundPUBLISHER* online collaborative writing and publishing environment, in development since 2000, and commercially available since 2002. By way of a brief introduction, *CommonGroundPUBLISHER* is a mixed-medium publishing system by means of which electronic books and digital books are created from a source file archived on a web server. Built on the open source software foundations of *Linux*, *Zope*, a *PostgreSQL* database and XML outputs, the system is a server-based software program which creates an online collaborative working environment through which authors, publishers

and other participants cooperate to create a published work. The system manages file and content flow from the creator to the consumer. The result is a production environment for the production of digitally printed books, one at a time on consumer order, and electronic books which can be downloaded to personal computers and reading devices. Books produced using the system are sold to consumers through the publisher and/or the author's online bookstores, and to bookstores through a B-2-B wholesale portal using the industry standard ONIX XML application profile (Common Ground, 2003a, b, c).

CommonGroundPUBLISHER is driven by a series of online input forms and output portals, through which structured data are created, stored and retrieved. It emulates traditional publishing – which means both its nomenclature and its narrative structure are based on the everyday world of authorship and publishing (and not the world of computing).

Here is one typical activity sequence or narrative that *CommonGroundPUBLISHER* structures and guides:

1. A publisher registers and thereby creates an online bookstore, http://PublisherName.Publisher-Site.com
2. An author registers and thereby creates a personal website, http://AuthorName.Author-Site.com, including an online bookstore for their own works (even if these are published by a number of different publishers).
3. An author creates a proposal and submits it to a publisher.
4. The publisher and author negotiate rights thereby creating a customised copyright agreement (and generating relevant digital rights data, based not only on the commercial variables, but drawing in relevant publisher, author and work/proposal data as captured to this point in the publishing narrative).
5. Publisher, author and referees collaborate to create a publishable work managing a manuscript through the drafting process.
6. Bookdata links and metadata are created, including ISBN, Digital Object Identifier, bibliographical and ecommerce metadata.
7. The work is made available as a published edition in the form of an eBook and pBook, available for direct purchase on both the publisher and the author sites, or for distribution into the conventional book trade and posting to Amazon.com.

Input screens and output portals string the CGML name-label tags into a meaningful activity sequence. In effect, the sequence of actions adds the implied or necessary verbs to the nouns – the paradigmatic tag-concepts which name the players, their work, what their work is about, how their work is going and the form it takes when it is finished.

Scollon describes this process as mediated discourse: a sequence of actions which involve doing things, handing them over, communicating about them (Scollon, 2001). These mediated actions are both discursive and material, creating an object (a book to be held or a text to be read on a screen) and

involving an exchange of symbolic meaning. These form dialogical chains, along the lines of the ATM conversation about which we spoke earlier in this article. The publishing dialogue is a long and complex chain of mediated actions; its aim is not simply something which is to be interpreted in discursive terms; it also involves handing backwards and forwards an artefact which is material as much as it is textual (and an electronic artefact is just as material as a printed one) – Scollon calls this the mediational means. Mediated actions fall within a nexus of practice in which social groups frame or are framed by particular positions – in our case authors, publishers, readers and the like. Underlying the chain is a kind of directedness, a series of intended (or unintended) ends on the part of the parties to the dialogue. In the nature of any dialogic interaction, prospective and retrospective ends may vary; the dialogue may or may not take directions which had been or could have been anticipated. As a chain proceeds it is channelled by a funnel of commitment, a kind of momentum alternately narrowing and reopening the range of possible actions towards their (even if momentary) end. Retrospectively, the chain of actions can be represented as narrative, or an account of how the sequence of actions began, progressed and ended.

This is to speak of actions around the text in the social processes of its construction, filling the active gaps between the text construction stages of <Proposal>, <Draft> and <Edition>, for instance, or the handing backwards or forwards of the developing artefact between <Author> and <CommissioningEditor>, <Referee> and <CopyEditor> in the process creating the sequence of <Draft>s that culminates in a publishable <Edition>.

A funnel of commitment of sorts is also to be found within text itself. Martin calls this genre, or the staged, goal-oriented social process which creates a specific text structure. A scientific report and a novel each has its own characteristic generic purposes, and these purposes provide a kind of direction in the text, reflected in its staging structure. 'Texts typically move through stages to a point of closure and explicitly treated by the speaker/listener as incomplete where closure is not attained' (Martin, 1992). Textual markup points to the underlying structures driving this funnel of textual commitment, and these are the features of the characteristic information architecture's written text, such as that of the book.

Based as it is on an activity sequence, and emulating the collaborative techniques of traditional publishing, *CommonGroundPUBLISHER* rests on a social model of text construction. In this respect, it is quite unlike the World Wide Web, where the text production protocols need not be so systematic or so collaborative. Essentially, the Web rests on a transmission model of text production. A person or organisation publishing to the Web can simply construct and 'post' their material. As a consequence, the quality of the finished text is very variable. By and large, in terms of the level of social interaction required to bring a text to a broad readership, the Web is the equivalent of 'vanity' publishing. *CommonGroundPUBLISHER* applies the discourse and practices of publishing to the digital technologies of text production and

network connectivity. It provides an environment for the collaborative construction of meanings, and for the co-ownership of intellectual property of a conventional publishing variety. Texts are created through a collaborative relationship which involves the multiple perspectives of players who assume various roles, including 'author', 'publisher', 'referee' and 'editor'. The software environment has inbuilt within it a series of social checks in the form of editorial and publishing approval processes. Content is not simply pushed out onto the Web. Rather, *CommonGroundPUBLISHER* establishes a collaborative environment for the construction of text based on the half-millennium-long history of publishing practice. It is these collaborative processes which are the basis for the credibility and authority of published books not so far achieved by the Internet.

CommonGroundPUBLISHER and its underlying expression language, CGML, is an open design scaffold. Herein resides a tension. On the one hand, it makes explicit concepts and processes often implicit in the practices of those who by historical accident have been the powerbrokers of text and culture. And by making these explicit, it hands the means of production of respected meaning to communities of practice, communities with content and a desire to speak, but who have historically been beholden to those who control the means of production and distribution of meaning.

Making things explicit, however, often means that they appear formulaic and rigid. This need not necessarily be the case. In fact, as much as they provide an enabling scaffold, the core concept-tags of CGML can be arranged into any narrative form. Together, a tagging schema such as CGML and a publishing architecture such as *CommonGroundPUBLISHER* open up an infinite number of expressive possibilities. Paradigm can be generative of an infinite number of narratives. CGML supplies the concept-tags of paradigm. Beyond this, the content which is name-tagged is infinitely variable, and the verbs between the names are added by the creators of content by virtue of their actions. CGML, in other words, provides a paradigm for authorship and publishing. Its users create their own narratives – activity sequences in the social negotiation of text and the dynamics inherent in the text structures they create.

Or to put it another way, and to reuse a distinction first made by Sassure, CGML supplies the 'langue' (a set of expressive possibilities in the form of conceptual and practical resources for authorship and publishing), whilst the users of *CommonGroundPUBLISHER* engage in 'parole' (the always unique expressive act, and the always unrepeatably creative conversations involved in the collaborative creation of a published work). *CommonGroundPUBLISHER* not only scaffolds this process; it also records the process (creators' drafts, referees' comments, commissioning editors' suggestions, copy editors' proofs, published editions and the like). Prospectively, the creative possibilities are opened out by *CommonGroundPUBLISHER*. Retrospectively, it maintains a narrative record of the making a specific work.

In this sense, CGML is a resource for meaning, rather than a prescriptive activity sequence for authorship and publishing or a supplied structure of

textual meaning. It is the basis for a process Kress calls transformation and design (Kress 2000, 2001). The design of meaning involves building on resources for meaning available in the world (the designed), appropriating and recombining the elements of those designs in a way that has never been done in quite the same way before (designing) and leaving a residue (the designed) which becomes a new set of resources for meaning, for the design process to begin afresh (Cope & Kalantzis, 2000a, b). This is also the way language itself works. Quoting Halliday, language is a 'resource for meaning making'; as such, it is a system which is open to choice, 'not a conscious decision made in real time but a set of possible alternatives' (Halliday, 1994). This brings us back to the distinction with which we began this section, between formal linguistics, which regards language as a system of rules, and functional linguistics, in which language is understood as a resource for meaning (Martin, 1992). As a scaffold, paradigm is not restricting or constraining. Rather, it is an enabling tool for widening the domain of expressive choice, for creating any number of narrative alternatives.

The Old and the New in Digital Designing

What's Old in the Digital

Much is promised by the enthusiasts of digital information and communication technologies, and much is lamented by its detractors. Hypertext represents a new, non-linear form of writing and reading, they say, in which readers are engaged as they never have been before, as active creators of meaning (Chartier, 2001). And the new technologies create a verisimilitude so striking that, for better or for worse, the elision of reality and its represented and communicated forms warrant the descriptive label 'virtual reality' (Virilio, 1997).

Of all that is claimed to be so new, however, much is not so new. What is claimed to be new to the digital era may have been new to print and the culture of the book, but then it is only new in a millennial frame of reference. Take, for instance, hypertext and virtual reality.

Hypertext, it is argued, is one of the most distinctive features of the digital communications environment, creating the possibility of non-linear readings and reader-chosen navigation paths. Even at first glance, hypermedia technologies are not so novel, tellingly using metaphorical devices drawn from the textual practices of the book such as 'browsing', 'bookmarking', 'pages' and 'index'. Moreover, when we examine the book as an information architecture, its characteristic devices are nothing if not hypertextual. Gutenberg's Bible had no title page, no contents page, no pagination, no index. In this sense, it was a truly linear text. However, within a century of Gutenberg's invention, the modern information architecture of the book had been developed, including regularly numbered pages, punctuation marks, section breaks, running heads,

indexes and cross-referencing. Amongst all these, pagination was the critical functional tool (Eisenstein, 1979).

The idea that books are linear and the Internet is multilateral is based on the assumption that readers of books necessarily read in a linear way. In fact, the devices of contents, indexing and referencing were designed precisely for alternative lateral readings, hypertextual readings, if you like. And the idea that the book is a text with a neat beginning and a neat end – unlike the Internet, which is an endless, seamless web of cross-linkages – is to judge the book by its covers. A book, however, does not begin and end at its covers, despite the deceptive appearances of its physical manifestation. It sits in a precise place in the world of other books, either literally when shelved in a library, and located in multiple ways by sophisticated subject cataloguing systems, or more profoundly in the form of the apparatuses of attribution (referencing) and subject definition (contents and indexes).

As for hypertext links that point beyond a particular text, all they do, albeit much more quickly, is what citation has always done. The footnote developed as a means of linking a text back to its precise sources, and directing a reader forward to a more detailed elaboration (Grafton, 1997). The only difference between the footnote and hypertext is that in the past you had to go to the library to follow through on a reference. Books, in other words, have developed elaborate ways of bursting out of their covers, of always referring to the world outside their covers, including to other books. This relationship to other writing and other books comes to be regulated in the modern world of private property by the laws, conventions and ethics of copyright, plagiarism, quotation, citation, attribution and fair use (Cope, 2001a, b, c).

Certainly, some things are different about the Internet in this regard. Clicking a hypertext link is faster and easier than leafing through cross-referenced pages or dashing to the library to find a reference. But this difference is a matter of degree, not a qualitative difference. In fact, the Internet sorely needs some of the skills of the old book trade. Compared to a library catalogue and a good book index, even the best of search engines seems rudimentary. The Internet is also a place where the quality of texts is at best uneven because copyright questions have been poorly resolved and the practices of editing and publishing have not yet been developed to the extent they have for the printed book. So, for all its dazzle, the Internet is not really that different to a book, and mostly still, not even as smart as a book. And for all the hype in hypertext, it only does what books have always done, which is to point to connections across and outside of a particular text.

Take also the hyperbole of the 'virtual'. There's not much about the virtual in the digital communications era which print and the book did not create as a genuine innovation 500 years earlier. The book as a communication technology brought modern people strangely close to distant and exotic places though the representation of those places in words and images on the printed page. So vivid at times was the representation that the early moderns could be excused for thinking they were virtually there. So too, in their time, the

photograph, the telegraph, the newspaper, the telephone, the radio and the television were all credited for their remarkable verisimilitude – remarkable for the 'real' being so far away, yet here so easily, so quickly, and so seemingly close and true to life. Each of these new virtual presences became a new kind of reality, a new 'telepresence' in our lives. We virtually lived through wars per medium of newspapers; and we virtually made ourself party to the lives of other people in other places and at other times through the medium of the novel or the travelogue. In this respect, digitisation is just another small step in the long and slow journey into the cultural logic of modernity. Digital reproduction and transmission of meaning simply reopens the fundamental questions of aura, authenticity and location raised by Walter Benjamin in the thirties in his discussion of 'the work of art in the age of mechanical reproduction' (Benjamin, 1970).

Once we get over the dazzle of the digital devices, in a semiotic sense, a lot of what is happening is not so new at all. Once we recognise this, we realise that there is value in a kind of reverse engineering of historical practices of authorship and publishing. This is to ensure that the virtues of the new technologies build upon, rather than reduce, the functional purposes and inherent sociability of these practices.

Or, put simply in terms of our own research and development agenda centred around Common Ground Markup Language and *CommonGroundPUBLISHER*, how do we migrate the information architecture of the book, the discourse of authorship and the collaborative practices of publishing into the environment of the Internet? And, in so doing, how do we broaden access to the means of production of valued and valuable meaning?

A book, old dictionary definitions tell us, is a thing. It is a printed volume of pages bound within covers. In one sense, new technologies challenge this definition: the CD-ROM which reproduces, more cheaply and efficiently, texts which have all-but disappeared from the world of print, such as the encyclopaedia; the eBook reading device which allows an endless number of books to be read in a physical space the size of one printed book, more conveniently and with negligible environmental impact compared to a product made from wood pulp; the digital talking book which speaks books to us as we walk or drive; the screen-readable Portable Document Format which deliberately emulates the printed text; and soon, perhaps, flexible substrates that can be read with reflected light. These technologies seem to portend the end of the book in its traditional definition.

However, in another definition, the book is certainly not going away. If we get away from a fixation with its tangible form, we might define the book in terms of its meaning function. And if we do, this would be our definition: *The book is an information architecture with a characteristic textual structure; and beyond its covers, it relates to the universe of books through conventionally legible intertextual devices*. A book is not a physical thing. A book is what a book does (Cope, 2001, a, b, c). And this is what a book does: it does intratextual things, intertextual things and extratextual or semantic and social things.

From an intratextual point of view (taking the unit of text to be the book), a book has a number of devices which mark thematic progress and allow alternative navigation paths. These include nesting devices (parts, chapters, sections and the like, often defined by heads), locating devices (table of contents, list of tables or figures, index), and cross-referencing devices (internal page references, glossaries, keys and devices for highlighting and managing internal redundancy such as summaries and conclusions).

From an intertextual point of view, books are interlinked through referencing conventions (such as citation), bibliographical practices (such as library cataloguing), and the inherent and often implicitly intertextual nature of all text (there is always some influence, conscious or unconscious, something that repeats from other works the authorial conceit of originality).

From an extratextual point of view, books always have semantic and social reference points. Semantically, they direct our attention to an external world, whether that purports to be factual or imagined. This external reference is the basis of validity and truth propositions, such as a pointer to an authentic historical source, or a scientifically grounded semantic reality located in a controlled vocabulary, such as C = carbon. The subject classification systems used by librarians attempt to add structure and consistency to extratextual semantic reference. From an extratextual point of view, texts also have an acknowledged or unacknowledged ontogenesis – the world of supporters, informers, helpers, co-authors, editors and publishers that comprises the socially constructivist domain of authorship and publishing. And capricious readers – texts enter this social world not as authorial edict, but open to alternative reading paths, in which communicative effects are constructed as much by readerly as authorial agendas.

This is the stuff that CGML maps as paradigm, and *CommonGroundPUBLISHER* opens out as narrative or a scaffold for creative alternatives. And although we are thinking books, we are actually talking authorship and publishing more broadly, and these in turn speak a language which could be applied to other creative endeavours.

Defined as an information architecture, the book today is everywhere. Its textual forms and communicative apparatuses are to be found throughout the new electronic formats – the notions of pages, headings, systems for listing contents (buttons and menus), referencing (links), cross-referencing (hypertext) and indexing (searching). In fact, other media are becoming progressively more book-like, such as the printable radio program of which Burrows et al (2001) speak. Or the DVDs which provide alternative viewing sequences never possible in the movie-house, and do this by creating a menu which divides the movie into 'chapters' and provides an 'introduction' in the form of the documentary about the making of the movie.

And so we find ourselves thrust into a new universe of textual media. In one moment, the commentators supply us with utopian readings; in the next apocalyptic. Leaving behind the linear world of the book, they speak breathlessly of hypertext and non-linear readings, of formerly passive book readers whose wilful navigation choices have turned them into active users of texts; and of the representation of virtual worlds in which the distant is brought so close, instantly and palpably. In moments of gloom, they also speak of a new inequality – the information inequality that is the result of the 'digital divide'. And they speak of a world of reduced human interaction, as sedentary persons increasingly find themselves tethered to machines.

Do the new electronic media foretell the death of the book? To answer this question, we need to reflect on the history and form of the book, as well as the electronic texts which, it is alleged, pose a threat. And our conclusion may well be that, rather than being eclipsed by the new media, the book will thrive as a cultural and commercial artefact.

There is no denying that the book itself is undergoing a process of transformation, and this is but one part of a series of transformations in the world of text creation generally. If some of what seems new in the digital era is not so new as the commentators sometimes promise, there can be no denying that some things are indeed new. As we have argued above, on closer examination, what is supposed to be new in the digital media is not so new at all. Hypertext's contribution is mechanical: it automates the information apparatuses that the printed book managed by page numbering, contents pages, indexing, citation and bibliography. And as for the virtual, what more did the written word and the printed image do than refer, often with striking verisimilitude, to things that are not immediately present. Indeed, the information architecture of the book, embodying as it does thousands of years' experience with recorded knowledge, provides a solid grounding for every adventure we might take in the new world of digital media.

Here we will return to the three fundamental changes with which we began this article, and upon which we have subsequently elaborated: the shift in the mechanics of rendering so that the primary focus is not now upon the crafting of meaning form but on marking up meaning function, from which alternative meaning forms follow; the shift towards multimodality, aided by the simple practical and material fact that mechanically reproduced representations of the linguistic, visual and audio are crafted and manufactured from the same raw materials; and the emergence of what we have termed 'polylingualism'. These are, we want to argue, significantly new potentialities opened up by digitisation.

A number of new practical possibilities are opened for the book. One is mixed-medium publication in which books will be available as print, as text to screen, as audio. This extends the range of access – disability access, for instance (Fathers, 2002) – as well as commercially viable channel alternatives.

Another is reader-created books, in which a student or lecturer builds and orders a customised book of course readings, or a poetically inclined lover pulls together a book of other people's poems, or a cook puts together a book of their favourite recipes across a dozen different cookbooks (Burrows et al, 2002). Still another is a burgeoning of titles which would not otherwise have been produced – in print and electronic formats – for which there may well be very small, but nevertheless viable, communities and markets, in areas historically neglected by a publishing industry driven by economies of scale (Cope & Brown, 2002; Cope & Mason 2002a, b; Cope & Ziguras, 2002). Not only can these new publishing opportunities extend the range and diversity of published works, they can also be used more generally as tools to support knowledge management practices (Cope & Freeman, 2002; Cope & Kalantzis, 2002), knowledge economy infrastructure (Vines & Naismith, 2002) and capacity development more generally (Vines, 2002). Indeed, simple high-tech/low-tech solutions may help bridge the digital divide at the same time as revaluing community-based knowledge. With one computer and one small digital book printer, for instance, hundreds or even thousands of copies of a beginning literacy textbook could be produced in any of the dozen or so smaller languages of East Timor, an extremely poor new country of only a million people.

So what is the book's future, as a creature of and conduit for human invention? The digital media, we would argue, represent an opportunity for the book more than a threat. The possibilities are threefold: increased access, greater diversity and enhanced democracy.

Access: as well as the conventional printed book (and there is little doubt that people will always be taking that old printed and bound artefact to the beach or to bed, for the foreseeable future at least), the same text may also be available in a range of alternative media. It could also be available on computer screen or printed to paper on the spot, as there is hardly a computer without a printer. It could be something that is read on an eBook reading device. It could be rendered to audio via speech synthesis. Or it could find itself coming to life through new electronic media currently in development, such as the paper-like plastic substrates that can be read from reflected light. The result will be greater and easier access to books, and new markets: the student who needs to have a chapter of a book tonight for an assignment due in tomorrow; the person who is visually impaired and wants the voice-synthesised version, or another person who wants to listen to the text while driving their car; the traveller who instantly needs just one piece of information from a travel guide and for whom a small piece of text on their mobile phone, about a particular monument or the nearby restaurant, is sufficient; or the teacher who wants to use some textual material as a 'learning object' in an electronic learning environment.

Diversity: the traditional book business ran on economies of scale. There was a magic number, somewhere around the 3000 mark, that made a book viable – worth the trouble to write, print and distribute. Of course, the longer

the print run, the better it was, according to the underlying logic of mass production. Costs reduced the longer the run, and access was at the cost of diversity. Mass production made for mass culture. Supporting this was a cumbersome infrastructure of slow moving inventory, large-scale warehousing, expensive distribution systems and heavily stocked retail outlets – bad business in every respect, and providing little return for anyone who made books their livelihood, least of all authors. It's not only the electronic reading devices that change the economies of manufacturing scale. Variable digital print does the same thing. One thousand different books can be printed in one run, and this entails no more cost than printing one thousand copies of the same book. Small communities with niche markets now play on the same field as large communities with mass markets. Book printing machines the size of a one-hour photo lab will be located in schools, in libraries and in bookstores, all of which will now be able to 'stock' any or even every book in the world.

Democracy: these developments will favour small communities of interest and practice. They will lower the entry point into the world of publishing. Museums, research centres, libraries, professional associations and schools might all become publishers. They'll be more than happy if a title sells a few hundred copies, or is even provided to the world for free – options that were not previously possible. As for quality, publishing decisions will be made by communities who feel deeply for their domain of content, for that is their domain of interest and expertise. It has never been the case that quantity, the traditional mass market measure of success, equates with quality, and that is destined to prove less the case in the future. Thousands of publishers and millions of new titles does not add up to information overload. There's already more than any one person can digest, and we have managed to find ways to locate what suits our particular needs and interests. The result can only be good – a more healthy democracy, a place of genuine diversity. Digital print will also provide a means to cross the digital divide, so that if you can't afford a computer for every person in a readership (a school in a developing country, for instance, or a new literature in a small, historically oral language), the proximity to the computers and digital print will allow cheap printed materials to be produced locally. There will be no need to buy someone else's language and culture to fill a local knowledge gap. This could be a world where small languages and cultures could flourish, and even, as machine translation improves, find that smallness does not mean isolation.

CGML and *CommonGroundPUBLISHER* have been designed to address all these possibilities in a very practical way. Let's hold the digital media to their promise of access, diversity and democracy.

References

ADL/SCORM (2003) *SCORM Overview: shareable content object reference model*. Advanced Distributed Learning. Available at: www.adlnet.org/index.cfm?fuseaction=scormabt&cfid=13586&cftoken=90030509

American National Standards Institute/National Information Standards Organization (2002) *Specifications for the Digital Talking Book*. Available at: www.niso.org/standards/resources/Z39-86-2002.html

Anderson, B. (1983) *Imagined Communities: reflections on the origin and spread of nationalism*. London: Verso.

Australian Government Locator Service (2003) *AGLS Metadata Standard*. National Archives of Australia. Available at: www.naa.gov.au/recordkeeping/gov_online/agls/summary.html

Benjamin, W. (1970) The Work of Art in the Age of Mechanical Reproduction, in H. Arendt (Ed.) *Illuminations*. London: Fontana.

Berners-Lee, T. (1990) *Information Management: a proposal*. Available at: www.w3.org/History/1989/proposal.html

Berners-Lee, T., Hendler, J. & Ora, L. (2001) The Semantic Web, *Scientific American*, May, 184(5), pp. 34-43.

Burrows, P., Loi, D., Wilkins, L. & Coburn, M. (2001) Electronically Mediated Learning Materials, in B. Cope & D. Kalantzis (Eds) *Print and Electronic Text Convergence*. Melbourne: Common Ground.

Burrows, P., Coburn, M. & Loi, D. (2002) The Reader as User in a Re-intermediated Supply Chain, in B. Cope & D. Mason (Eds) *Markets for Electronic Book Products*. Melbourne: Common Ground.

Chartier, R. (2001) *Readers and Readings in the Electronic Age*. Available at: www.text-e.org/conf

CIP4 Organization (2003) *The CIP4 Website*. Available at: www.cip4.org/

Coburn, M., Burrows, P., Loi, D. & Wilkins, L. (2001) eBook Readers: new directions in enabling technology, in B. Cope & D. Kalantzis (Eds) *Print and Electronic Text Convergence*. Melbourne: Common Ground.

Common Ground (2003a) *Common Ground Markup Language*. 6 vols. Melbourne: Common Ground.

Common Ground (2003b) *Creator to Consumer Automated Book Publishing Technology 2003*. Available at: www.C-2-CSystem.com

Common Ground (2003c) *Method and Apparatus for Extending the Range of Useability of Ontology Driven Systems and for Creating Interoperability Between Different Markup Schemas for the Creation, Location and Formatting of Digital Content*. Melbourne: Common Ground.

Cope, B. (1998) The Language of Forgetting: a short history of the word, in M. Fraser (Ed.) *Seams of Light: best antipodean essays*. Sydney: Allen & Unwin.

Cope, B. (2001a) *Globalisation, Multilingualism and the New Text Technologies*. Melbourne: Common Ground.

Cope, B. (2001b) Making and Moving Books in New Ways, from the Creator to the Consumer, in B. Cope & D. Mason (Eds) *Digital Book Production and Supply Chain Management*. Melbourne: Common Ground.

Cope, B. (2001c) New Ways with Words: print and eText convergence, in B. Cope & D. Kalantzis (Eds) *Print and Electronic Text Convergence*. Melbourne: Common Ground.

Cope, B. & Brown, R. (Eds) (2002) *Value Chain Clustering in Regional Publishing Services Markets. Vol. 3.4, Emerging Markets for Books, from Creator to Consumer*. Melbourne: Common Ground.

Cope, B. & Freeman, R. (Eds) (2001) *Digital Rights Management and Content Development. Vol. 2.4, Technology Drivers across the Book Production Supply Chain, from the Creator to the Consumer*. Melbourne: Common Ground.

Cope, B. & Freeman, R. (Eds) (2002) *Developing Knowledge Workers in the Printing and Publishing Industries: education, training and knowledge management in the publishing supply chain, from creator to consumer*. Vol. 4. Melbourne: Common Ground.

Cope, B. & Gollings, G. (Eds) (2001) *Multilingual Book Production. Vol. 2.2, Technology Drivers Across the Book Production Supply Chain, from the Creator to the Consumer*. Melbourne: Common Ground.

Cope, B. & Kalantzis, M. (Eds) (1993) *The Powers of Literacy: genre approaches to teaching writing*. London and Pittsburgh: Falmer Press (UK edition) and University of Pennsylvania Press (US edition).

Cope, B. & Kalantzis, M. (2000a) Designs for Social Futures, in B. Cope & M. Kalantzis (Eds) *Multiliteracies: literacy learning and the design of social futures*. London: Routledge.

Cope, B. & Kalantzis, M. (Eds) (2000b) *Multiliteracies: literacy learning and the design of social futures*. London: Routledge.

Cope, B. & Kalantzis, M. (2002) Manageable Knowledge: communication, learning and organisation change, in B. Cope & R. Freeman (Eds) *Developing Knowledge Workers in the Printing and Publishing Industries: education, training and knowledge management in the publishing supply chain, from creator to consumer*. Melbourne: Common Ground.

Cope, B. & Kalantzis, M. (2003) Significado digital y defensa de una pedagogía de los multialfabetismos [Digital Meaning and the Case for a Pedagogy of Multiliteracies], in M.A. Pereyra (Eds) *Multiliteracies in the Digital Space*. Granada: Consorcio 'Fernando de los Ríos'.

Cope, B. & Mason, D. (Eds) (2002a) *Markets for Electronic Book Products. Vol. 3.2, Emerging Markets for Books, from Creator to Consumer*. Melbourne: Common Ground.

Cope, B. & Mason, D. (Eds) (2002b) *New Markets for Printed Books. Vol. 3.1, Emerging Markets for Books, from Creator to Consumer*. Melbourne: Common Ground.

Cope, B. & Ziguras, C. (Eds) (2002) *The International Publishing Services Market. Vol. 3.3, Emerging Markets for Books, from Creator to Consumer*. Melbourne: Common Ground.

C-2-C Project (2003) *Creator to Consumer in a Digital Age: book production in transition*. Available at: www.C-2-CProject.com

Denny, M. (2002) *Ontology Building: a survey of editing tools*. Available at: www.xml.com/pub/a/2002/11/06/ontologies.html

Dixon, R.M.W. (1980) *The Languages of Australia*. Cambridge: Cambridge University Press.

DocBook Technical Committee (2003) *DocBook Home Page*. Available at: www.oasis-open.org/docbook/

Dublin Core Metadata Initiative (2003) *Dublin Core Metadata Registry*. Available at: http://dublincore.org/dcregistry/index.html

Dunn, R., Hester, R. & Readman, A. (2001) Printing Goes Digital, in B. Cope & D. Kalantzis (Eds) *Print and Electronic Text Convergence*. Melbourne: Common Ground.

Eco, U. (1999) *Kant and the Platypus: essays on language and cognition*. London: Vintage.

EDItEUR (2001) *ONIX Product Information Standards 2.0*. International Group for Electronic Commerce in the Book and Serials Sectors. Available at: www.editeur.org/onix.html

EDItEUR (2003) *ONIX Mappings to MARC*. International Group for Electronic Commerce in the Book and Serials Sectors. Available at: www.editeur.org/onixmarc.html

Eisenstein, E.L. (1979) *The Printing Press as an Agent of Change: communications and cultural transformation in early-modern Europe*. Cambridge: Cambridge University Press.

Encoded Archival Description Working Group (2002) *Encoded Archival Description Tag Library*. Society of American Archivists. Available at: http://lcweb.loc.gov/ead/tglib/

Fairclough, N. (1992) *Discourse and Social Change*. Cambridge: Polity Press.

Fathers, M. (2002) Disability Access to Published Material, in B. Cope & D. Mason (Eds) *Markets for Electronic Book Products*. Melbourne: Common Ground.

Gee, J.P. (1996) *Social Linguistics and Literacies: ideology in discourses*, 2nd edn. London: Taylor & Francis.

Gerber, L. (2001) Translation in a Digital Environment, in B. Cope & G. Gollings (Eds) *Multilingual Book Production*. Melbourne: Common Ground.

Gilster, P. (1997) *Digital Literacy*. New York: John Wiley.

Goldfarb, C. (1990) *A Brief History of the Development of SGML*. SGML Users' Group. Available at: www.oasis-open.org/cover/sgmlhist0.html

Grafton, A. (1997) *The Footnote: a curious history*. London: Faber & Faber.

Halliday, M.A.K. (1978) *Language as Social Semiotic*. London: Edward Arnold.

Halliday, M.A.K. (1994) *An Introduction to Functional Grammar*, 2nd edn. London: Edward Arnold.

Harold, E.R. & Means, W.S. (2002) *XML*. Sebastopol: O'Reilly.

Harvey, D. (1996) *Justice, Nature and the Geography of Difference*. Cambridge, MA: Blackwell.

Hodge, B. & Kress, G. (1988) *Social Semiotics*. London: Polity Press.

Husserl, E. (1970) *The Crisis of European Sciences and Transcendental Phenomenology*. Evanston: Northwestern University Press.

IFLA Study Group on the Functional Requirements for Bibliographic Records (1998) *Functional Requirements for Bibliographic Records*. The Hague: International Federation of Library Associations and Institutions.

IMS Global Learning Consortium (2002) *IMS Learning Design Specification*. Available at: www.imsproject.org/learningdesign/index.cfm

IMS Global Learning Consortium (2003) *Specifications*. Available at: www.imsproject.org/specifications.cfm

<indecs> (2000) *Interoperability Data in e-Commerce Systems*. Available at: www.indecs.org/

International DOI Foundation (2002) *The DOI Handbook*. Oxford: International DOI Foundation.

International Trade Standards Organization for the eBook Industry (2003) *Open eBook Home Page*. Available at: www.openebook.org

Kress, G. (2000) Design and Transformation: new theories of meaning, in B. Cope & M. Kalantzis (Eds) *Multiliteracies: literacy learning and the design of social futures*. London: Routledge.

Kress, G. (2001) Issues for a Working Agenda in Literacy, in M. Kalantzis & B. Cope (Eds) *Transformations in Language and Learning: perspectives on multiliteracies*. Melbourne: theLearner.com/Common Ground Publishing.

Kress, G. & van Leeuwen, T. (1996) *Reading Images: the grammar of visual design*. London: Routledge.

Lash, S. & Urry, J. (1994) *Economies of Signs and Space*. Beverly Hills: Sage.

LaTeX (2003) *LaTeX: a document preparation system*. Available at: www.latex-project.org/

Luo, S. (1998) *An Illustrated History of Printing in Ancient China*. Hong Kong: City University of Hong Kong.

Luria, A. (1981) *Language and Cognition*. New York: John Wiley.

Man, J. (2002) *The Gutenberg Revolution*. London: Review.

MARC Standards Office (2000) *ONIX to MARC 21 Mapping*. Library of Congress. Available at: www.loc.gov/marc/onix2marc.html

MARC Standards Office (2001) *MARC to Dublin Core Crosswalk*. Library of Congress. Available at: http://lcweb.loc.gov/marc/marc2dc.html

MARC Standards Office (2003a) *MARC Standards*. Library of Congress. Available at: http://lcweb.loc.gov/marc

MARC Standards Office (2003b) *Metadata Encoding and Transmission Standard*. Library of Congress. Available at: www.loc.gov/standards/mets/

MARC Standards Office (2003c) *Metadata Object Description Scheme*. Library of Congress. Available at: www.loc.gov/standards/mods/

Martin, J. (1992) *English Text: system and structure*. Philadelphia: John Benjamins.

Mason, D. (2001) Cataloguing for Libraries in a Digital World, in B. Cope & D. Mason (Eds) *Digital Book Production and Supply Chain Management*. Melbourne: Common Ground.

Mason, D. & Tsembas, S. (2001) Metadata for eCommerce in the Book Industry, in B. Cope & D. Mason (Eds) *Digital Book Production and Supply Chain Management*. Melbourne: Common Ground.

Mitchell, W.J. (1995) *City of Bits: space, place and the infobahn*. Cambridge, MA: MIT Press.

Mühlhäusler, P. (1996) *Linguistic Ecology: language change and linguistic imperialism in the Pacific region*. London: Routledge.

Multimedia Description Schemes Group (2002) *ISO/IEC CD 21000 Part 6: Rights Data Dictionary: International Organisation for Standardisation, ISO/IEG JTC 1/SC 29/WG* 11: Coding of Moving Pictures and Audio.

National Curriculum Online (2002) *The National Curriculum Metadata Standard*. Department for Education and Skills. Available at: www.nc.uk.net/metadata/

New London Group (1996) A Pedagogy of Multiliteracies: designing social futures, *Harvard Educational Review*, 66, pp. 60-92.

Nua.com (2003) *How Many Online*. ComputerScope. Available at: www.nua.ie/surveys/how_many_online/

ODRL (2002) *Open Digital Rights Language*. IPR Systems. Available at: http://odrl.net/

Ong, W. (1982) *Orality and Literacy: the technologizing of the word*. London: Methuen.

Online, EdNA (2000) *EdNA Metadata Standard*. Education Network Commonwealth of Australia. Available at: www.edna.edu.au/metadata/

OUL/EML (2003) *Educational Modelling Language*. Open University of the Netherlands. Available at: http://eml.ou.nl/

Paskin, N. (2003). DRM Technologies: identification and metadata, in E. Becker, D. Gunnewig, W. Buhse & N. Rump (Eds) *Digital Rights Management: technical, economical, juridical and political aspects*. Berlin: Springer-Verlag.

Phillipson, R. (1992) *Linguistic Imperialism*. Oxford: Oxford University Press.

PODi (2003) *PODi: the digital printing initiative*. Available at: www.podi.org/

Scollon, R. (1999) Multilingualism and Intellectual Property: visual holophrastic discourse and the commodity/sign. Paper presented at Georgetown University Round Table, 1999, January.

Scollon, R. (2001) *Mediated Discourse: the nexus of practice*. London: Routledge.

Text Encoding Initiative (2003) *Text Encoding Initiative Home Page*. Available at: www.tei-c.org/

United Kingdom Office for Library and Information Networking (2001) *The Biblink Core Metadata Set*. Forum for Metadata Schema Implementers. Available at: www.schemas-forum.org/registry/schemas/biblink/BC-schema.html

Unicode (2003) *Unicode Home Page*.

Vines, R. (2002) Knowledge Management, Capacity Development and Publishing, in B. Cope & R. Freeman (Eds) *Developing Knowledge Workers in the Printing and Publishing Industries: education, training and knowledge management in the publishing supply chain, from creator to consumer*. Melbourne: Common Ground.

Vines, R. & Naismith, L. (2002) Knowledge Management and Publishing Systems Infrastructure, in B. Cope & R. Freeman (Eds) *Developing Knowledge Workers in the Printing and Publishing Industries: education, training and knowledge management in the publishing supply chain, from creator to consumer*. Melbourne: Common Ground.

Virilio, P. (1997) *Open Sky*. London: Verso.

Vygotsky, L. (1962) *Thought and Language*. Cambridge, MA: MIT Press.

Wacquet, F. (2001) *Latin, or the Empire of the Sign*. London: Verso.

Williams, S. (2002) *Free as in Freedom: Richard Stallman's crusade for free software*. Sebastopol: O'Reilly.

World Wide Web Consortium (1999) *HTML 4.01 Specification*. Available at: www.w3.org/TR/1999/REC-html401-19991224/

World Wide Web Consortium (2002) *XHTML 2.0 Working Draft*. Available at: www.w3.org/TR/2002/WD-xhtml2-20020805/

W3C (2002) *Semantic Web*. Available at: www.w3c.org/2001/sw/

W3C (2003) *RDF Semantics*. Available at: www.w3.org/TR/2002/WD-rdf-mt-20021112/

XrML (2003) *Extensible Rights Markup Language*. Available at: www.xrml.org

---

BILL COPE has been a Director of Common Ground Publishing since 1984, and is currently developing experimental mixed-medium print and Internet publishing. He is a former First Assistant Secretary in the Department of the Prime Minister and Cabinet and Director of the Office of Multicultural Affairs. His academic research and writing crosses a number of disciplines, including history, education and sociology; and examines themes as varied as Australian immigration, workplace change and literacy learning. Since 2000, he has worked on a joint Common Ground/RMIT research initiative, 'Creator to Consumer in a Digital Age: book production in transition', funded as part of the Infrastructure and Industry Growth Fund of the Book Production Enhanced Printing Industry Competitiveness Scheme of the Australian Government's Department of Industry, Tourism and Resources. The outcome of this work is a series of research reports in book form examining changes in book production technologies and markets, as well as a programme of national and international seminars and industry training and tertiary education materials (www.c-2-cproject.com).. *Correspondence*: The Centre for the Workplace Communication and Culture, PO Box 463, Altona, Victoria 3018, Australia (bill.cope@commongroundpublishing.com).

MARY KALANTZIS is Dean of the Faculty of Education, Language and Community Services at RMIT University and President of the Australian Council of Deans of Education. She has been a Commissioner of the Human Rights and Equal Opportunity Commission, Chair of the Queensland Ethnic Affairs Ministerial Advisory Committee and a member of the Australia Council's Community Cultural Development Board. With Bill Cope, she is co-author of a number of books, including: *Productive Diversity,* Pluto Press, Sydney, 1997; *A Place in the Sun: re-creating the Australian way of life,* Harper Collins, Sydney, 2000; and *Multiliteracies: literacy learning and the design of social futures,* Routledge, London, 2000. *Correspondence*: Faculty of Education, Bundorra West Campus, PO Box 71, Bundoora, Victoria 3083, Australia (mary.kalantzis@rmit.edu.au).