

# ***The Knowledge Web***

***Learning and Collaborating  
on the Net***

**MARC EISENSTADT & TOM VINCENT**  
**Knowledge Media Institute**


[kmi.open.ac.uk/knowledgeweb](http://kmi.open.ac.uk/knowledgeweb)

# ***The Knowledge Web***

***Learning and Collaborating  
on the Net***

MARC EISENSTADT & TOM VINCENT  
Knowledge Media Institute

[kmi.open.ac.uk/knowledgeweb](http://kmi.open.ac.uk/knowledgeweb)



**KOGAN  
PAGE**

***To Jacqueline, Amelia, Nathan and Leo (ME)***

***To Ollie, Bernadette and Christine (TV)***

---

First published in 1998

Reprinted with a new Introduction in 2000

Reprinted 2000

Apart from any fair dealing for the purposes of research or private study, or criticism or review, as permitted under the Copyright, Designs and Patents Act, 1988, this publication may only be reproduced, stored or transmitted, in any form or by any means, with the prior permission in writing of the publishers, or in the case of reprographic reproduction in accordance with the terms of licences issued by the CLA. Enquiries concerning reproduction outside these terms should be sent to the publishers at the undermentioned addresses:

Kogan Page Limited  
120 Pentonville Road  
London N1 9JN, UK

Stylus Publishing Inc.  
22883 Quicksilver Drive  
Sterling VA 20166-2012, USA

© Knowledge Media Institute, 1998, 2000

---

**British Library Cataloguing in Publication Data**

A CIP record for this book is available from the British Library.

ISBN 0 7494 3178 4

---

Printed and bound in Great Britain by Biddles Ltd, *www.biddles.co.uk*

# Acknowledgements

## Preparation of the book

We are grateful to the following people who have made significant contributions to producing the electronic version of this book: Simon Keats for compilation, composition, layout and editing; Jon Linney for cover design, and advice on book design, printing and graphics; Jon Linney and Tony Seminara for specific figures throughout the book.

## Reprint rights

Much of Chapter 6 was first published in the *Proceedings of the 1997 IEEE Frontiers in Education* conference under the title: 'Innovations in large-scale supported distance teaching: transformation for the Internet, not just translation', and is reused here by kind permission of the IEEE. We are grateful to the following organizations and individuals for permission to reprint figures and/or Web pages: Electronic Text Center, University of Virginia (Figure 2.2), Dan Kaveney, Prentice-Hall (Figure 2.3), Iain Gilmour (Figure 2.5), Marianna Bultjens, Scottish Sensory Centre (Figure 2.12).

## Sponsorship

We are grateful to the following organizations that have provided support, sponsorship and funding for research and development projects, the outcomes of which appear within the book. They include:

Andersen Consulting; Apple Computer; BT; Department of Trade and Industry; Department for Education and Employment; European Union (COPERNICUS Programme and ESPRIT IBROW3, PATMAN, and HC-REMA projects); Higher Education Funding Council (England); Gatsby Charitable Foundation; GEC Plessey Telecommunications, Ltd.; Guide Dogs for the Blind Association; IBM; IT World; London Mathematical Society; Macromedia Europe; Nuffield Foundation; RealNetworks, Inc.; Relay Systems Ltd.; Royal National Institute for the Blind; Sun Microsystems Inc.; UK Science and Engineering Research Council; JISC (JTAP programme); UK Science and Engineering Research Council jointly-sponsored CASE Award (with BT); Voxware Inc.; The Open University's Development Fund, INSTILL Programme, Research Committee, Office of Technology Development, IET Research Committee, and Social Sciences Faculty.

## People

Many people have been involved in the projects described throughout the book. We would particularly like to acknowledge the following for their contributions:

Rachel Bellamy, Mike Brayshaw, Jenny Bull, Marianna Buultjens, Paula Cole, Chris Denham, Shamus Foster, Adam Freeman, Prue Fuller, Iain Gilmour, Caroline Gray, Jerzy Grzeda, Ann Hanson, Tony Hasemer, Ben Hawkridge, Stephanie Houde, Ralph Keats, Simon Keats, Andrew Law, Ches Lincoln, Mike Lewis, Jon Linney, Fred Lockwood, Ross Mackenzie, Heather Mason, Chris McKillop, Mark Miller, Diane Mole, Philip Odor, Roger Penfound, Barbara Poniatowska, Kevin Quick, Mike Richards, Andy Rix, Neil Robinson, Craig Rodine, Ortenz Rose, Tony Seminara, Malcolm Story, Mary Taylor, Debra Thompson, Chris Valentine, Hannah Walton, Ben Whalley, Matthew Whalley, Dave Williams, Mike Wright.

## Affiliations

Most of the authors are members of The Open University's Knowledge Media Institute, located in the Berrill Building at The Open University, Milton Keynes MK7 6AA, UK. The exceptions are as follows (all affiliations are still at The Open University unless otherwise stated):

Sir John Daniel (Vice Chancellor); Josie Taylor and Matthew Stratfold (Institute of Educational Technology); Marian Petre, Linda Carswell, Pete Thomas (Faculty of Mathematics and Computing); Marion Phillips (Open University Oxford Regional Office, Foxcombe Hall, Boars Hill, Oxford OX1 5HR, UK); Alberto Riva (Department of Computer and Systems Science, University of Pavia, Italy).

# List of colour plates

Colour plates appear between pages 152 and 153. Figure numbers refer to the black-and-white figures within the text, where relevant.

- I The Knowledge Media Institute in The Open University's Berrill Building
- II Figure 1.1. Sir John Daniel at IDLCON-97
- III Figure 2.6. World map interface to Webcams
- IV Figure 2.7. The Virtual Microscope
- V Figure 2.9. First Flight – Wright Brothers plane simulation
- VI Figure 3.9. The weigh mode of the Virtual Spring
- VII Figure 4.1. The interactive course map
- VIII Figure 5.8. A video segment in the Advisor prototype
- IX Figure 6.2. Example of a marked assignment
- X Figure 6.3. Audiographic tutorial student screen dump with inset photos of student and tutor
- XI Figure 8.3. Output of D3E toolkit
- XII Figure 8.7. JIME slide interface
- XIII Figure 9.6. A question screen in the Pub Quiz
- XIV Figure 10.2. A screen snapshot of the KMi Planet frontpage
- XV Figure 11.3. The final state of a movie made by a tutor for a student
- XVI Figure 14.2. Tadzebao from World Wide Design Lab

# Introduction to the 2000 printing

The *fin de siècle* 'Internet Gold-Rush' during which we write this new introduction offers us a wonderful opportunity to reflect on the change and the *pace* of change surrounding Knowledge Media. After all, the very phrase 'Knowledge Media', which we have chosen as the theme for our research work and the name of our Institute, is itself at the very crossroads of the most intensive activity the world has ever known in areas of learning, education, media development, communication, collaboration and management rethinking. Change happens in 'Internet time', where developments that were unforeseen during one print run of a book may have arisen from scratch, been influential on an earth-shattering and paradigm-shifting scale, and be gone and forgotten by the time of the next print run (cognescenti might like to reflect, for instance, on the fate of 'push technologies' during the short life cycle of this very book!). Yet we're still delighted that 'good old print' is the chosen medium for this discursive description and overview of our research work for a large audience: it is still the appropriate medium for the content we want to convey and the style in which we want to convey it. And yes, the accompanying Web site at [kmi.open.ac.uk/knowledgeweb](http://kmi.open.ac.uk/knowledgeweb) still works to provide hands-on experience and fast-turnaround updates, as one would expect from a research group like ours.

As we argued originally (and infuriatingly to some), not only is it true that the three important things to understand are *still* the Web, the Web and the Web, but it is also true that the Web remains a poor and frustrating mediator of learning. Naturally, it is getting a lot better, and with increasing speed, and the work herein is an important harbinger of the way forward (in our humble opinion). Re-examining our work with the hindsight of 'Internet time', ie a few years after some of the original work was undertaken, we feared that it would appear positively 'Stone Age', and require a total rewrite of the book. Happily, this has not been necessary: the work stands to this day on its own merits, and in many cases (as we describe below) has either had significant impact on a large scale, has been the recipient of an award or prize recognizing its innovation, or has led to further innovation that drives the exciting pace of research in today's successful and expanding Knowledge Media Institute (KMi).

KMi today has grown to more than 50 researchers, funded by research grants from the European Union, UK research councils, private sector sponsorships, the UK Millennium Commission and The Open University Development Strategy Group. We have

created educationally relevant research prototypes that have impacted on hundreds of thousands of adult learners, thousands of local schoolchildren, and tens of thousands of others who have downloaded and explored the content from our Web sites.

During our own period of growth as a research lab, the 'penetration' or percentage of households and educational establishments with Internet connections, good quality Web access and self-maintained Web sites continues to skyrocket. High-speed or broadband connectivity, fuelled by increasing demand, lower prices and vastly improving multimedia content, continues to gain in prominence. Interestingly, while some users and developers rush to pump increasing video content down the line as we enter the broadband-enabled new millennium, we continue to argue that we can do *better* than 'mere video', precisely because the constraints of narrow-band (ie typical dial-up modem speeds of the old millennium) made us concentrate, throughout the work described in this book, on two phenomena that have paid huge dividends: a) media-rich end-user experiences that involve elements other than, or in addition to, video streams, including simulation models, graphically sophisticated animations supported by high quality audio, shared group activities and photo-realistic virtual reality; and b) 'pumping content down the line' is a completely misguided metaphor for anyone involved in education in the first place – it belies an embarrassingly old and largely discredited approach of 'pouring knowledge from a vessel', patently absurd to anyone motivated enough to pick up this volume in the first place.

Our approach to media-rich learning experiences derives from constructivist models of education, where a key underlying theme concerns ownership of content. Whether we are working with schoolchildren, adult learners, or corporate employees and their managers, we find that empowering individuals to *create their own content* provides a powerful motivating influence that completely transforms the learning experience. Thus, while the NASA Web site is a fantastic achievement and learning resource, we are confident that the children who built their own Web-enabled Mars buggy (described in Chapter 2) will remember that experience for the rest of their lives, and they certainly continue to take pride in the 'hits' their school Web site received from NASA! Asynchronous discussion forums, of the kind described in Chapters 7, 8 and 16, are significant because ownership of the content and direction of the discussion typically rests with the students themselves. Indeed, authorship of original journal content in the D3E environment (Chapter 8) is made simple enough that content creation and distribution becomes an easy



and motivating source for the spread of new ideas throughout a school, university or workplace. Our environments, such as Tadzebao (Chapter 14), ensure that discussions are shared in an easy and highly visible manner.

To our mantra of 'Web, Web, Web' we add a further thought from our colleague Stuart Watt – as establishments embark on whole-sale Web site creation, they should really ask themselves three questions: 1) Who is going to maintain the site?; 2) Who is going to maintain the site?; and 3) Who is going to maintain the site? The significance of the 'maintenance burden' cannot be overstated. Not only can everyone access your Web site now, but they can see how out of date it is with respect to recent developments. We face this dilemma ourselves, and must stress to educators and managers embarking on this route that they need to consider the problem from day one in order to address the longevity and scalability of their activities.

We highlight the nature of current developments by means of a quick 'KMi Tour' based on the main section headings of the book: Learning media, Collaboration and presence, and Knowledge Systems.

## **Part I: Learning media**

The rapid changes in information technologies that have been experienced in the past two decades have continued since this book was first written. In parallel, this challenges educationalists as to how best to exploit the new media associated with these technologies. Others have been motivated to introduce these technologies and media into their institutions because there are attractive economic benefits. These approaches are often diametrically opposed. For example, it is technically very easy to transfer Open University printed course materials to the Web. The savings in print and distribution costs are obvious and significant. But would it benefit students? The most likely outcome would be irritated students who spend many hours printing out the material! This is not to say that there can't be simultaneous benefits for learning and cost. However, the domination of the latter is only likely to have short-term financial gains.

In his book *Mega-Universities and Knowledge Media* (Kogan Page, revised 1999), John Daniel asserts that the biggest danger facing those using knowledge media today is that poor quality online education will return distance learning to the low status of its paper-based equivalent of earlier days. In pursuing their own renewal, universities should recognize that technology presents dangers as well as opportunities.

Universities have thrown themselves with varying degrees of enthusiasm and preparation into online teaching. Early assessments of these initiatives suggest that they are enriching the learning environment for existing campus students rather than drawing in substantial numbers of newcomers. Furthermore, some students are voting with their feet against purely Web-based courses, arguing that the Web is better seen as a component of a wider teaching strategy. In the Knowledge Media Institute, we remain optimistic, yet critical, about the Web as a teaching and learning resource. Evidence continues to confirm that the Web, as with other technologies and media, can be successfully exploited provided that the educational need to which it is applied is identified first. This may seem obvious, but how many times do we see computer-based applications looking for a need!

Consideration of the pros and cons of knowledge media was a feature of the developments that illustrate Part 1 of the book. This continues to be the case and it is exemplified by developments that are focused on the needs of disabled learners. The paradox remains that many of the applications of the new information technologies that provide new opportunities for learning can, at the same time, introduce barriers to learning for some. However, it is also the case that taking into account the needs of disabled learners at an early stage of development can benefit all learners. The recent Europrix Knowledge and Discovery award for the OU's Geology course acknowledged the quality and role of the Virtual Microscope and Virtual Field Studies (Chapter 2). This development originated to meet the needs of disabled students unable to undertake these activities. The Virtual Microscope also serves to illustrate the rapid changes in information technology. It was not that long ago when it was a significant challenge to create a CD ROM, or even hard disk, version with the quality and quantity of images that are required to meet the standard set by earth scientists. During the past two years, innovations in streaming media have enabled a Web version of the Virtual Microscope to be produced.

Research has continued at the Knowledge Media Institute into how the Web can be used to support the teaching of young children. Developing the curriculum around historical and contemporary experiments was under way when this book was first written. Since then, this area of development has flourished. In particular, the history of flight has proved to be an exciting and valuable topic for several aspects of the curriculum, including history, control technology and communication and writing skills, as well as developing new media and information technology skills. This

topic includes numerous examples of early experiments associated with flight that in many cases never went beyond drawings and, if they did, often resulted in spectacular failures. Understanding why and how these experiments succeeded or failed can be very informative. The work of the aviation pioneer Sir George Cayley (1773–1857) has been used with 10- to 12-year-old children for the past two years. After undertaking a wide range of associated physics experiments, children built model versions of the Cayley Glider, and several were taken to the actual hill in the north of England that Cayley used. All of the basic experiments were filmed, with children explaining what they were doing, and the making and flying of the glider. The final stage is the building of a Web site that brings everything together. The children have become authors in the knowledge society rather than merely recipients of information. One clear outcome from this project in a local school is the confidence and skills gained by the children in using the new media and associated technologies, and their use becomes transparent, with the focus being the curriculum content – a highly desirable goal in any aspect of using knowledge media in education.

Another outcome of projects in schools has been the identification of the increasing gap between the knowledge and skills of children and parents in the educational applications of information technologies. This has led to the establishment of a major project at the Knowledge Media Institute into how this gap can be reduced and, in turn, enable parents to contribute more to their children's learning when the use of computers is involved. The project involves 300 parents, working in teams of five (CLUTCH CLUBS – Computer Literacy Understanding Through Community History) at different schools in the Milton Keynes area, researching a local history topic of their choice. As the information gathered is being published on the Web, it is necessary to provide training and support in a range of media and IT skills. An interesting aspect of the project is how such a large group of people with limited skills can collectively produce a high quality and integrated Web site. This has led to the development of numerous server-side tools and templates that support the parents at a distance in acquiring appropriate skills for researching their topics, collecting information, manipulating digital media and publishing their results on the Web. The outcomes of the project have wider implications where the Internet and the Web are used for teaching and training, which is considered in Part 1 of the book.

## Part II: Collaboration and presence

The Open University and the Internet were made for each other: ‘distance teaching’ + ‘streaming Webcast lectures’ = a marriage made in heaven, right? Wrong, wrong, wrong!

First, ‘distance teaching’ is an old mis-statement of what the Open University actually does. We prefer to talk about the OU’s trademarked ‘supported open learning’ to emphasize its vast tutorial and support infrastructure and 30 years’ of experience in the provision of individualized learning. Second, ‘streaming Webcast lectures’, as a technology, are a sure-fire path to ‘virtual classrooms’, which replicate the *worst* aspects of conventional education! As we point out in detail in Chapters 1, 4, 6 and 9, the technology is, and always has been, an *adjunct* to the methods deployed throughout the Open University. Third, the research under way at the Knowledge Media Institute is *not* about ‘marriage-brokering’ between glitzy new technologies and distance teaching challenges – on the contrary, it is about evolving new ideas and new media by working with educators, designers, cognitive scientists, knowledge engineers and students, and looking at how learning needs can best be met in a creative way.

The projects described in Part II of this book have tried consistently to push back the frontiers of what could be achieved by looking at collaboration technologies *in a new light*. Four projects in particular have flourished in ways that we correctly foresaw, and that are worth describing briefly here: D3E, Lyceum, WebSymposia and Planet.

The Digital Document Discourse Environment (D3E) described in Chapter 8 continues to be the supporting infrastructure for the highly successful *Journal of Interactive Media in Education*, as well as the London Mathematical Society’s *Journal of Computation and Mathematics*. It is being adapted for mainstream use by several large courses within the Open University, and is the foundation for several ‘sharing best practice’ activities in UK industry. In this latter capacity, one of the key strengths of D3E has always been its role as a knowledge management tool: where corporate memories reside largely (and inaccessibly) in the tacit knowledge of employees and a mountain of (raw text) documents, D3E lends rational support by facilitating a principled structuring of the raw documents. More importantly, it delivers *ownership* of the content back to the originators, and provides a low-stress entry point into the world of knowledge representation, which is vital for constructing robust and reusable repositories of knowledge.

The KMi Stadium projects described in Chapter 9 have branched into two significant strands. The very large-scale one-to-

many Webcasts have evolved into a strand that we now refer to as WebSymposia, while the smaller many-to-many voice-mediated group tutorials have continued to evolve under the Lyceum banner that we introduced in that chapter. WebSymposia emphasizes very high production values, large scale and rich media including audio, video, animations, simulations, virtual reality and real-time audience polling. It provides a bespoke 'wrapper' environment, built upon a modern multiplayer games architecture that is technologically agnostic in the sense that it works with any of the three leading streaming media products (RealPlayer, Windows Media and Apple QuickTime at the time of this writing). The intuition behind WebSymposia is that we can provide presenters and audiences with a *better* overall experience, even using today's bandwidth constraints, than they could have even with a dedicated satellite TV station of their own. The key is to concentrate on a Web-centric rich-media mix rather than focusing entirely on video, and this is why our 'Knowledge Web' philosophy holds such promise.

Lyceum, in contrast, takes to heart the comments of Open University Associate Lecturers (our 8,000+ adjunct faculty members providing mentoring and tutoring worldwide), who work with groups of no more than 20 geographically dispersed students per tutor. Lyceum has been moved from the KMi research environment to the Open University's Centre for Educational Software, and turned into a robust group audio-conferencing and graphics-sharing tool customized for the requirements of OU course teams. The Open University's Centre for Modern Languages has been an enthusiastic supporter and early user of Lyceum, exploiting its voice and graphics-sharing capabilities to deliver French and German practical tutorials, exercises and examinations via the Internet. Another such team is Open University Business School course B823, 'Knowledge Management', which has deployed Lyceum widely for use on a group concept-mapping course assignment. Over 1,000 students are registered for the Knowledge Management course at the time of this writing, which makes Lyceum one of the key large-scale voice-on-the-Net activities in a teaching context.

*KMi Planet*, our newspaper-on-the-Web mediated by intelligent agent software, has been adapted for use by several organizations, including the British Educational Communications and Technology Agency (BECTa) and the 'Marchmont distance learning observatory' of the UK's LearnDirect project (formerly the University for Industry). The Planet philosophy (easy story production) has also been influential in the way Web-centric news

is developed and delivered throughout several of the schools with whom we are working, and in the CLUTCH project described earlier. The key ingredient from the perspective of collaboration is that *ownership* of the original content resides in the hands of individual grass-roots members of the organizations concerned, once again providing an easy and highly motivating entry point into the world of knowledge sharing.

### **Part III: Knowledge systems on the Web**

An important premise underlying the creation of the Knowledge Media Institute, and indeed the very work captured by our title *Knowledge Web*, is that researchers in the areas of knowledge systems and artificial intelligence can benefit greatly from the integration of rich media and excellent graphic design into their work. Similarly, we've always felt that researchers in the areas of new media can benefit from the integration of intelligent software into their work. This double-ended integration continues to bear fruit through a variety of projects that have evolved from the work described in Part III.

One large European Union-funded project that began after the first printing of the book, and that incorporates many of the book's themes, is entitled 'ENRICH: Enriching Representations of Work for Organizational Learning'. The ENRICH project aims to harness situated workplace learning, where the knowledge resides tacitly in the heads of employees, and provide a sound basis for reusing that knowledge in a principled and accessible way across an entire organization. It uses the Planet and D3E architectures described above as the low-stress entry point into getting 'stories' and 'articles' from motivated employees onto the Web, and then encourages the same employees to build an informal taxonomy of their stories and articles using our WebOnto tool, introduced in Chapter 14. This taxonomy is our 'Trojan horse', for it yields a major leap into the world of knowledge representation. From this informal taxonomy, the more heavyweight knowledge representation tools come into play, and create a rich representation that can then be subject to modern automated reasoning tools. Thus, rather than simple 'raw text' that must be (frustratingly) accessed by conventional search engines, a representation full of paraphrases and inferences can be reused by different employees in different contexts, and corporate memory can live on beyond the shallow collection of 'bits and bytes' in which it is otherwise hopelessly stuck.

Another EU-funded project is RichODL, which seeks to deploy the ENRICH approach in an engineering context. RichODL is

designed to assist in the training of students and employees as dynamic system modellers and simulators, and integrates rich simulation environments within a D3E-style discussion environment and a WebOnto-style indexing and searching environment.

Another important Knowledge Systems activity in KMi is The Bayesian Knowledge Discoverer (BKD), a data-mining tool that can extract reusable knowledge from databases, using sound and accountable statistical methods, even when data is missing or incomplete. BKD allows estimation of conditional probability from data, extraction of the graphical structure from data, goal-oriented propagation of evidence and missing data handling. It has now been distributed to over 3,000 sites worldwide.

The ‘psychological agents’ described in Chapters 15 and 16 continue to play a prominent role in our research. The meeting-management agent known as Luigi has evolved through constant user feedback into a more powerful environment that we refer to as Meet-O-Matic. Meet-O-Matic continues the theme of keeping control entirely in the hands of the meeting proposer, but emphasizes the kind of coarse-grained perspectives (‘which days are good for you?’) relevant to organizers of large meetings. The Virtual Participant of Chapter 16 has attracted additional funding, and promises a magnificent means of managing the mountains of e-mail that threaten our collective sanity at the moment!

## **The future**

We said in the original introduction (page 17), ‘Knowledge Media looks to us as if it will be the key not only to education in the next Millennium, but also to many facets of society’. Alas, the ‘next Millennium’ is now here as the presses roll on this edition. Other than replacing ‘next’ by ‘this’, we’d like to stand firmly by our original claims, and as before we invite readers to ‘join us on this exciting journey’!

# Contents

Acknowledgements and affiliations	v
List of colour plates	vii
Introduction to the 2000 printing	viii
<b>Knowledge Media takes off on the Web: introduction and overview of the book</b>	<b>1</b>
<i>Mark Eisenstadt and Tom Vincent</i>	

---

<b>Part I    Learning media</b>	<b>19</b>
1. Can you get my hard nose in focus? Universities, mass education and appropriate technology	21
<i>Sir John Daniel</i>	
2. The Web: enabler or disabler	31
<i>Tom Vincent and Peter Whalley</i>	
3. Collaborative learning in networked simulation environments	47
<i>Peter Whalley</i>	
4. Media integration through meta-learning environments	63
<i>Tamara Sumner and Josie Taylor</i>	
5. Developing Web-based student support systems: telling student stories on the Internet	79
<i>Peter Scott and Marion Phillips</i>	
6. Innovations in large-scale supported distance teaching: transformation for the Internet, not just translation	97
<i>Marian Petre, Linda Carswell, Blaine Price and Pete Thomas</i>	



<b>Part II Collaboration and presence</b>	<b>117</b>
7. Promoting learner dialogues on the Web	119
<i>Matthew Stratfold</i>	
8. New scenarios in scholarly publishing and debate	135
<i>Simon Buckingham Shum and Tamara Sumner</i>	
9. Exploring telepresence on the Internet: the KMi Stadium Webcast experience	153
<i>Peter Scott and Marc Eisenstadt</i>	
10. KMi Planet: putting the knowledge back into media	173
<i>John Domingue and Peter Scott</i>	
11. Sharing programming knowledge over the Web: the Internet Software Visualization Laboratory	185
<i>John Domingue and Paul Mulholland</i>	
<b>Part III Knowledge Systems on the Web</b>	<b>199</b>
12. Accessing artificial intelligence applications over the World Wide Web	201
<i>Alberto Riva and Marco Ramoni</i>	
13. Knowledge modelling: an organic technology for the knowledge age	211
<i>Arthur Stutt and Enrico Motta</i>	
14. The World Wide Design Lab: an environment for distributed collaborative design	225
<i>Zdenek Zdrahal and John Domingue</i>	
15. Psychological agents and the new Web media	239
<i>Stuart Watt</i>	
16. The Virtual Participant: a tutor's assistant for electronic conferencing	249
<i>Simon Masterton</i>	
References	267
Index	287