

Accessing the E-book edition of **GUIDE TO HEALTH INFORMATICS, THIRD EDITION**

Using the VitalSource® ebook

Access to the VitalBook™ ebook accompanying this book is via VitalSource® Bookshelf – an ebook reader which allows you to make and share notes and highlights on your ebooks and search across all of the ebooks that you hold on your VitalSource Bookshelf. You can access the ebook online or offline on your smartphone, tablet or PC/Mac and your notes and highlights will automatically stay in sync no matter where you make them.

1. **Create a VitalSource Bookshelf account at** <https://online.vitalsource.com/user/new> or log into your existing account if you already have one.
2. **Redeem the code provided in the panel below to get online access to the ebook.** Log in to Bookshelf and click the **Account** menu at the top right of the screen. Select **Redeem** and enter the redemption code shown on the scratch-off panel below in the **Code To Redeem** box. Press **Redeem**. Once the code has been redeemed your ebook will download and appear in your library.



DOWNLOAD AND READ OFFLINE

To use your ebook offline, download BookShelf to your PC, Mac, iOS device, Android device or Kindle Fire, and log in to your Bookshelf account to access your ebook:

On your PC/Mac

Go to <http://bookshelf.vitalsource.com/> and follow the instructions to download the free **VitalSource Bookshelf** app to your PC or Mac and log into your Bookshelf account.

On your iPhone/iPod Touch/iPad

Download the free **VitalSource Bookshelf** App available via the iTunes App Store and log into your Bookshelf account. You can find more information at <https://support.vitalsource.com/hc/en-us/categories/200134217-Bookshelf-for-iOS>

On your Android™ smartphone or tablet

Download the free **VitalSource Bookshelf** App available via Google Play and log into your Bookshelf account. You can find more information at <https://support.vitalsource.com/hc/en-us/categories/200139976-Bookshelf-for-Android-and-Kindle-Fire>

On your Kindle Fire

Download the free **VitalSource Bookshelf** App available from Amazon and log into your Bookshelf account. You can find more information at <https://support.vitalsource.com/hc/en-us/categories/200139976-Bookshelf-for-Android-and-Kindle-Fire>

N.B. The code in the scratch-off panel can only be used once. When you have created a Bookshelf account and redeemed the code you will be able to access the ebook online or offline on your smartphone, tablet or PC/Mac.

SUPPORT

If you have any questions about downloading Bookshelf, creating your account, or accessing and using your ebook edition, please visit <http://support.vitalsource.com/>

GUIDE TO HEALTH INFORMATICS

Third edition



GUIDE TO HEALTH INFORMATICS

Third edition

ENRICO COIERA

Professor and Director
Centre for Health Informatics
Australian Institute of Health Innovation
Macquarie University
Sydney, Australia



CRC Press

Taylor & Francis Group
Boca Raton London New York

CRC Press is an imprint of the
Taylor & Francis Group, an **informa** business

Cover art: A data visualisation of how often clinical trials directly compare different drugs. Arcs above and below the centre line show the differences between trials with or without industry funding. Colours reflect groups in the Anatomical Therapeutic Chemical Classification. Adam Dunn 2014 with permission.

CRC Press
Taylor & Francis Group
6000 Broken Sound Parkway NW, Suite 300
Boca Raton, FL 33487-2742

© 2015 by Enrico Coiera
CRC Press is an imprint of Taylor & Francis Group, an Informa business

No claim to original U.S. Government works
Version Date: 20150128

International Standard Book Number-13: 978-1-4441-7050-4 (eBook - PDF)

This book contains information obtained from authentic and highly regarded sources. While all reasonable efforts have been made to publish reliable data and information, neither the author[s] nor the publisher can accept any legal responsibility or liability for any errors or omissions that may be made. The publishers wish to make clear that any views or opinions expressed in this book by individual editors, authors or contributors are personal to them and do not necessarily reflect the views/opinions of the publishers. The information or guidance contained in this book is intended for use by medical, scientific or health-care professionals and is provided strictly as a supplement to the medical or other professional's own judgement, their knowledge of the patient's medical history, relevant manufacturer's instructions and the appropriate best practice guidelines. Because of the rapid advances in medical science, any information or advice on dosages, procedures or diagnoses should be independently verified. The reader is strongly urged to consult the relevant national drug formulary and the drug companies' printed instructions, and their websites, before administering any of the drugs recommended in this book. This book does not indicate whether a particular treatment is appropriate or suitable for a particular individual. Ultimately it is the sole responsibility of the medical professional to make his or her own professional judgements, so as to advise and treat patients appropriately. The authors and publishers have also attempted to trace the copyright holders of all material reproduced in this publication and apologize to copyright holders if permission to publish in this form has not been obtained. If any copyright material has not been acknowledged please write and let us know so we may rectify in any future reprint.

Except as permitted under U.S. Copyright Law, no part of this book may be reprinted, reproduced, transmitted, or utilized in any form by any electronic, mechanical, or other means, now known or hereafter invented, including photocopying, microfilming, and recording, or in any information storage or retrieval system, without written permission from the publishers.

For permission to photocopy or use material electronically from this work, please access www.copyright.com (<http://www.copyright.com/>) or contact the Copyright Clearance Center, Inc. (CCC), 222 Rosewood Drive, Danvers, MA 01923, 978-750-8400. CCC is a not-for-profit organization that provides licenses and registration for a variety of users. For organizations that have been granted a photocopy license by the CCC, a separate system of payment has been arranged.

Trademark Notice: Product or corporate names may be trademarks or registered trademarks, and are used only for identification and explanation without intent to infringe.

Visit the Taylor & Francis Web site at
<http://www.taylorandfrancis.com>

and the CRC Press Web site at
<http://www.crcpress.com>

In memory of Bob Palese

Contents

Note	xi
Preface	xiii
Acknowledgements	xv
Introduction to health informatics – the systems science of healthcare	xix

Part 1 Basic concepts in informatics

1 Models	3
2 Information	13
3 Information systems	23

Part 2 Informatics skills

4 Communicating	37
5 Structuring	51
6 Questioning	63
7 Searching	77
8 Making decisions	93

Part 3 Information systems in healthcare

9 Information management systems	117
10 The electronic health record	127

VIII CONTENTS

11	Designing and evaluating information and communication systems.....	145
12	Implementation	173
13	Information system safety	195
14	Information economics	221

Part 4 Guideline- and protocol-based systems

15	Guidelines, protocols and evidence-based healthcare	237
16	Computer-based protocol systems	253
17	Designing, disseminating and applying protocols	269

Part 5 Communication systems in healthcare

18	Communication systems basics.....	285
	Interlude – the Internet and the World Wide Web	297
19	Information and communication networks	301
20	Social networks and social media interventions	325
21	Telehealth and mobile health.....	343

Part 6 Language, coding and classification

22	Terms, codes and classification	371
23	Healthcare terminologies and classification systems.....	381
24	Natural language and formal terminology	401

Part 7 Clinical decision support and analytics

25	Clinical decision support systems	417
	Interlude – artificial intelligence in medicine.....	433
26	Computational reasoning methods.....	437
27	Model building for decision support, data analysis and scientific discovery	453

Part 8 Specialized applications for health informatics

28	Patient monitoring and control.....	475
29	Population surveillance and public health informatics	497
30	Bioinformatics.....	515
31	Clinical bioinformatics and personalized medicine.....	537
32	Consumer health informatics	555
	Glossary	571
	References	589

Note

Healthcare is an ever-changing science. As new research and clinical experience broaden our knowledge, changes in treatment and drug therapy are required. The author and the publisher of this work have checked with sources believed to be reliable in their efforts to provide information that is complete and generally in accord with the standards accepted at the time of publication. However, in view of the possibility of human error or changes in medical sciences, neither the author nor the publisher or any other party who has been involved in the preparation or publication of this work warrants that the information contained herein is in every respect accurate or complete, and they are not responsible for any errors or omissions or for the results obtained from the use of such information. Readers are encouraged to confirm the information contained herein with other sources.

Preface

This book is written for healthcare professionals who wish to understand the principles and applications of information and communication methods and technologies in healthcare. The text is presented in a way that should make it accessible to anyone, independent of prior technology knowledge. It is suitable as a textbook for undergraduate and postgraduate training in the clinical aspects of informatics and as an introductory textbook for those pursuing a postgraduate career in health and biomedical informatics.

The text is designed to be used by *all* healthcare professionals, including nurses and allied health professionals, and not just medical practitioners. When I use the term ‘clinician’ in this book I am referring to any healthcare practitioner directly involved in patient care. Those with a background in information and communication technology should find the book a valuable introduction to the diverse applications of technology in health, as well as summarizing the unique challenges in this domain.

With the third edition of the Guide, I have kept the essential backbone of the informatics story the same as in previous editions. Part 1 contains foundational chapters that explain simply the abstract concepts that are core to informatics. Subsequent chapters then build upon those foundations. Part 2 contains a set of chapters that explore all the main themes of the book from the perspective of informatics skills. Practising clinicians must understand how to communicate effectively, structure information, ask questions, search for answers and make robust decisions. Informatics is as much about doing as it is about the tools we use, and these chapters make clear why the study of health informatics is the foundation of all other clinical activities.

We return to each of these information and communication system themes in later chapters, where we take a more technological focus. The book has a strong emphasis on demonstrating what works and what does not work in informatics. I have created a new evaluative framework based on the value of information that runs through the book, to help understand why some classes of intervention appear to work so much better than others. Each chapter ends with questions intended to test the reader’s understanding of the chapter or stimulate discussion of the material. Not all the answers to the questions are easy or obvious, and some are specifically designed to challenge.

Health informatics has undergone many changes since the appearance of the second edition in 2003. New themes have emerged, and new methods and technologies have been adopted. Old ideas have fallen by the wayside. The third edition is thus significantly longer

than earlier editions and contains six major new chapters. The chapters cover implementation, information system safety, social networks and social media interventions, model building for decision support, data analysis and scientific discovery, clinical bioinformatics and personalized medicine and consumer informatics. The new chapters are extensive and focus as much as possible on basic concepts and principles, rather than on simple narrative descriptions of the topics. All the old chapters have been overhauled, most of them significantly restructured, updated and extended. There are very many new sections within the updated chapters, covering diverse topics including health information exchanges, m-health, patient consent models, natural language processing and even augmented reality. Several old chapters have been deleted or merged.

It seemed a foolhardy mission for a single author to write a comprehensive text on health informatics in 1996 or even 2003. In 2014, the task took on Quixotic proportion as I debated which material should appear in an introductory text and what should be excluded. My rule of thumb was to include wherever possible basic principles and organizing structures as a priority and include only information that was likely to have a long half-life. The research base of our discipline grows rapidly, and it is very easy to create chapters that date quickly.

As always, the balance is between creating an introductory work that has some longevity and explores the core concepts needed to understand our discipline with a single and unified voice or writing an encyclopedic multi-author work that tries to do everything, but has too many voices, becomes out of date quickly and overwhelms students. At least for this edition I think we have still managed to keep the book to a 'single voice' overview – although I have had many expert colleagues help me with sourcing, writing and structuring the material and checking what has been written. I hope that the clarity of this text makes up for any limitations in its comprehensiveness.

EC

Sydney, Australia

November, 2014

Acknowledgements

I have been greatly helped, supported and influenced by many people as I wrote this book. To all of you, named and unnamed, I give my thanks.

Health informatics is now old enough as a discipline that we are beginning to lose the first great generation of informaticians. Since the last edition, Branko Cesnik, a pioneer of informatics in Australia, and Mario Stefanelli, a pioneer of artificial intelligence in medicine, have passed away and are both greatly missed by me. For those of you who are starting your careers, the next decade offers you a chance to meet and learn from many of the founding 'greats', and you should not lose that opportunity.

I have spent the last 15 years at the Centre for Health Informatics (CHI). Many of my past research staff and students are now leaders in their own right, and many are distributed across the globe. Others have stayed in Australia and still work closely with me as we investigate the edges of our discipline. You all are a daily inspiration to me and teach me far more than I teach you, I am sure. Denise Tsiros hovers over us all to make sure we are a family and never just an institution. I am also lucky to work with some outstanding senior colleagues who run sister centres to my own within the Australian Institute of Health Innovation – Jeffrey Braithwaite, Johanna Westbrook and Ken Hillman – all wise heads and supporting shoulders.

I have many wonderful international research partners, and their names appear on the many papers we all write together. Thank you for working with me. I spent several productive months on sabbatical in Harvard over 2010 and experienced the gracious hospitality of Ken Mandl at the Children's Hospital Informatics Program. Although this book was not written there, it was the necessary 'pre-contemplation' period I needed to start the project.

This book is based on a rich research tradition, and research cannot flourish without its backers. I am in particular much indebted to Shaun Larkin, without whose support the resources for many of our projects would not have been forthcoming, especially in consumer informatics.

I have a small cheer squad of people who keep me motivated, and you all pop up at different times, but just when I need you. Thank you, Terry Hannan and Louise Schaper. Thank you (again) Bill Caldicott. My cousin Dr Francesco Coiera needs a special thank you because I had to choose between finishing this book and visiting him in beautiful Italy, and the book won.

The ever-patient team at Taylor and Francis needs my thanks as they watched one deadline slip after another, but still believed in this project. I would specifically like to thank Jo Koster for being both delightful to work with and patient and understanding in adversity.

Back in 1999, Bruce Dowton, now the Vice-Chancellor of Macquarie University in Sydney, championed the creation of CHI and created the first Chair in Medical Informatics in Australia at the University of New South Wales. Without his faith and vision, I doubt we would have started or prevailed. Now, many years later, the circle has closed, and I have been gifted to again work with Bruce at Macquarie University and see where we can take this discipline over the next 10 years.

My biggest North American fans have always been my in-laws, Bob and Aline Palese. For this edition, we will deeply miss Bob's efforts to personally increase my sales by door knocking every physician he knew. Rest in peace, Bob.

As it was with the earlier editions but even more so with this one, writing has been a long and sometimes lonely marathon. I have finished only because I have been sustained by the love of my parents and, most of all, my pillars Blair and Lucca.

Contributors

I would like specifically to acknowledge and personally thank the following individuals who have each assisted me in creating this new edition:

Dr Vitali Sintchenko, from the University of Sydney, co-authored Chapters 29, 30 and 31.

Dr Farah Magrabi, from Macquarie University, co-authored Chapter 13.

Dr Annie Lau, from Macquarie University, co-authored Chapter 32.

Dr Zac Kohane, from Harvard University, provided material for the bioinformatics chapter as it appeared in the second edition, and we have retained some of this material in Chapters 30 and 31 of the current edition. The text and illustrations are drawn with permission from Kohane, Kho and Butte's *Microarrays for an Integrative Genomics*, published by MIT press.

Dr Guy Tsafnat and **Dr Blanca Gallego-Luxan** provided me with valuable material that helped in the creation of Chapter 27.

Both Farah and Vitali are dear colleagues, and they worked very closely with me in the early stages of writing this book. They helped me to think through the structure of this edition and specific details of the revisions. Their expertise, assistance and support are very much appreciated, and the book is far better for their many contributions.

Several other individuals with specialized expertise helped me by reading and providing feedback on specific chapters, including Grahame Grieve (Chapter 19), Dr Adam Dunn (Chapter 20) and Dr Guy Tsafnat (Chapter 24).

Margaret Jackson read and checked every chapter as it was finished, in the process becoming an instant health informatics expert. If there are any remaining grammatical disasters, logical inconsistencies or typographic blunders, they rest upon my shoulders.

Illustrations

Several illustrations in this book are either adapted or reproduced from other sources, and we have made our best effort to acknowledge the creators of these images appropriately as the images appear in the text. Many of these illustrations come from open access journals, which generously permit illustration reproduction as long as attribution is clear. Several illustrations appeared in the second edition with permission, and these have again been used. Figure 4.1 is taken from *BMJ* 1999; **318**:1527–1531 and appears with permission. Figure 16.1 is taken from Fox *et al.* (1996) and appears with kind permission of the copyright holder John Fox. Figures 28.6, 28.9, 28.10 and 28.11 appear with permission of W. B. Saunders Company Ltd., London. Figure 15.2 appears with permission of the World Health Organization and Figure 31.5 with permission of *Nature*.

Introduction to health informatics – the systems science of healthcare

Of what value, it may be urged, will be all the theorizing and speculation through which it would profess to guide us, when we come to practise at the bedside? Who has not heard so-called practical men say that medicine is a purely empirical science; that everything depends upon facts and correct experience; or, perhaps, that the power to cure is the main point? All arguments and theories, they say, do not enable the physician to treat his patients more correctly; in an art like medicine they rather do harm, or, at best, no positive good. It is there we are in need of experience - facts, and above all, remedies and their correct employment. All the rest is evil.

F. Oesterlen, Medical Logic, 1855, p 8

If physiology literally means ‘the logic of life’, and pathology is ‘the logic of disease’, then health informatics is the logic of healthcare. It is the study of how clinical knowledge is created, shaped, shared and applied. It is the rational study of the way we think about healthcare and the way that treatments are defined, selected and evolved. Ultimately, it is the study of how we organize ourselves, both patients and professionals, to create and run healthcare organizations. With such a pivotal role, the study of informatics is as fundamental to the practice of medicine and the delivery of healthcare in this century as anatomy or pathology was in the last.

Health informatics is thus as much about computers as cardiology is about stethoscopes (Coiera, 1995). Rather than drugs, x-ray machines or surgical instruments, the tools of informatics are more likely to be clinical guidelines, decision support systems, formal health languages, electronic records or communication systems such as social media. These tools, however, are only a means to an end, which is the delivery of the best possible healthcare.

Although the name ‘health informatics’ came into use only around 1973 (Protti, 1995), it is a study that is as old as healthcare itself. It was born the day that a clinician first wrote down some impressions about a patient’s illness and used these to learn how to treat the next

patient. Informatics has grown considerably as a clinical discipline in recent years fuelled, in part no doubt, by the advances in computer technology. What has fundamentally changed now is our ability to describe and manipulate health knowledge at a highly abstract level and to store vast quantities of raw data. We now also have access to rich communication systems to support the process of healthcare.

We can formally say that health informatics is the study of information and communication processes and systems in healthcare. Health informatics is particularly focussed on

1. Understanding the fundamental nature of these information and communication processes and describing the principles that shape them.
2. Developing interventions that can improve upon existing information and communication processes.
3. Developing methods and principles that allow such interventions to be designed.
4. Evaluating the impact of these interventions on the way individuals or organizations work or on the outcome of the work.

Specific subspecialties of health informatics include clinical informatics, which focusses on the use of information in support of patient care, and bioinformatics, which focusses on the use of genomic and other biological information.

The rise of health informatics

Perhaps the greatest change in clinical thinking over the last 2 centuries has been the ascendancy of the scientific method. Since its acceptance, it has become the lens through which we see the world, and it governs everything from the way we view disease to the way we battle it.

It is now hard to imagine just how controversial the introduction of theory and experimental method into medicine once was. Then, it was strongly opposed by empiricists, who believed that observation, rather than theoretical conjecture, was the only basis for rational practice.

With this perspective, it is almost uncanny to hear again the old empiricists' argument that 'healthcare is an art', and not a place for unnecessary speculation or formalization. This time, the empiricists are fighting against those who wish to develop formal theoretical methods to regulate the communal practice of healthcare. Words such as quality and safety, clinical audit, clinical guidelines, indicators, outcome measures, healthcare rationing and evidence-based practice now define the new intellectual battleground.

While the advance of science pushes clinical knowledge down to a fine-grained molecular and genetic level, it is events at the other end of the scale that are forcing us to change the most. First, the enterprise of healthcare has become so large that it now consumes more national resource than any country is willing to bear. Despite sometimes heroic efforts to control this growth in resource consumption, healthcare budgets continue to expand. There is thus a social and economic imperative to transform healthcare and minimize its drain on social resources.

The structure of clinical practice is also coming under pressure from within. The scientific method, long the backbone of medicine, is now in some ways under threat. The reason for this is not that experimental science is unable to answer our questions about the nature of disease and its treatment. Rather, it is almost too good at its job. As clinical research ploughs ahead in laboratories and clinics across the world, like some great information-generating machine, health practitioners are being swamped by its results. So much research is now published each week that it can literally take decades for the results of clinical trials to translate into changes in clinical practice.

So, healthcare workers find themselves practising with ever-restricting resources and unable, even if they had the time, to keep abreast of the knowledge of best practice hidden in the literature. As a consequence, the scientific basis of clinical practice trails far behind that of clinical research. Consumers struggle even more and have to contend with conflicting messages and information they find online, such as in social media.

Two hundred years ago, enlightened physicians understood that empiricism needed to be replaced by a more formal and testable way of characterizing disease and its treatment. The tool they used then was the scientific method. Today we are in analogous situation. Now the demand is that we replace the organizational processes and structures that force the arbitrary selection among treatments with ones that can be formalized, tested and applied rationally.

Modern healthcare has also moved away from seeing disease in isolation to understanding that illness occurs at a complex system level. Infection is not simply the result of the invasion of a pathogenic organism, but the complex interaction of an individual's immune system, bacterial flora, nutritional status and our social, environmental and genetic endowments. By seeing things at a system level, we come ever closer to understanding what it really means to be diseased, and how that state can be reversed.

We now need to make the same conceptual leap and begin to see the great systems of knowledge that enmesh the delivery of healthcare. These systems produce our knowledge, tools, languages and methods. Thus, a new treatment is never created and tested in intellectual isolation. It gains significance as part of a greater system of knowledge because it occurs in the context of previous treatments and insights, as well as the context of a society's resources and needs. Further, our work does not finish when we scientifically prove that a treatment works. We must try to disseminate this new knowledge and help others to understand, apply and adapt it.

These then are the challenges for healthcare. Can we put together rational structures for the way clinical evidence is pooled, communicated and applied to routine care? Can we develop organizational processes and structures that minimize the resources we use and the harms we create and maximize the benefits delivered? And finally, what tools and methods need to be developed to help achieve these aims in a manner that is practicable, testable and in keeping with the fundamental goal of healthcare – the relief from disease? The role of health informatics is to develop a systems science for healthcare that provides a rational basis to answer these questions, as well as to create the tools to achieve these goals.

The scope of informatics is thus enormous. It finds application in the design of clinical decision support systems for practitioners, consumer decision aids and online health services, in the development of computer tools for research and in the study of the very essence

of healthcare – its corpus of knowledge. Yet the modern discipline of health informatics is still relatively young. Many other groups within healthcare are also addressing the issues raised here and not always in a co-ordinated fashion. Indeed, these groups are not always even aware that their efforts are connected or that their concerns are also concerns of informatics.

The science of what works

I want to let you in on a secret. There are really only three questions that matter in informatics. At the beginning of any new informatics endeavor, you just need to ask:

1. *What is the problem that we are trying to solve?*
2. *How will we know when we have succeeded?*
3. *Is technology the best solution, or are there simpler alternatives?*

If you make sure these questions are asked, then you will be thought of as wise indeed. If you know enough to answer them, you could be held up as an informatics guru.

Reading this at the very beginning of your informatics journey, you may be surprised by the triviality of these questions. Re-reading them at the end of your journey through this book, you may now understand why little else matters and may also understand how rare it is for these questions to be asked in the real world – and what the almost inevitable consequences of not asking them are.

With this framing, we need to understand three things about any informatics intervention – its *possibility*, its *practicability* and its *desirability*. Possibility reflects the science of informatics – what in theory can be achieved? Practicability addresses the potential for successfully engineering a system or introducing a new process – what can actually be done given the constraints of the real world? Desirability looks at the fundamental motivation for using a given process or technology.

These criteria are suggested because we need to evolve a framework to judge the claims made for new technologies and those who seek to profit from them. Just as there is a long-standing, sometimes uneasy, symbiosis between the pharmaceutical industry and medicine, there is a newer and consequently less examined relationship between healthcare and the computing and telecommunication industries. Clinicians should judge the claims of these newcomers in the same cautious way that they examine claims about a new drug and perhaps more so, given that clinicians are far more knowledgeable about pharmacology than they are about informatics and telecommunications.

Overview of the book

The first goal of this book is to present a unifying set of basic informatics principles that influence everything from the delivery of care to an individual patient through to the design of whole healthcare systems. The book is organized into a number of parts that revolve around the two distinct but interwoven strands of information and communication systems. Although

the unique character of each strand is explored individually, there is also an emphasis on understanding the rich way in which they can interact and complement each other.

Part 1 – Basic concepts in informatics

This first part of the book offers an intuitive understanding of the basic theoretical concepts needed to understand informatics – the notions of what constitutes a model, what one means by information, and what defines a system. Each concept is used to develop an understanding of the basic nature of information and communication systems. A recurring theme of the book, first articulated here, is the need to understand the limitations imposed upon us whenever we create a model of the world or use it to design a technology. Understanding these limitations defines the ultimate limits of possibility for informatics, irrespective of whichever technology one may wish to apply in its service.

Part 2 – Informatics skills

Building upon the concepts in Part 1, the second part of the book looks at the practical lessons that can be drawn from informatics to guide everyday clinical activity. Every clinical action, every treatment choice and investigation, is shaped by the available information and how effectively that information is communicated. Five basic clinical informatics skills are explored, each with its own individual chapter:

1. *Communicating* effectively is based upon understanding cognitive models of information processing and is constantly challenged by the limits of human attention and the imperfection of models.
2. *Structuring information*, with a particular focus on the patient record, is shown to be dependent upon the task at hand, the channel used to communicate the message and the agent who will receive the message.
3. *Questioning* others to find information is essential in clinical practice to fill the ever present gaps in every individual's knowledge.
4. *Searching for knowledge* describes the broader strategic process of knowing where to ask questions, evaluating answers and refining questions in the light of previous actions, and it occurs in many different settings, from when patients are interviewed and examined through to when treatment options are canvassed.
5. *Making decisions* requires a clear problem formulation, followed by the assembly of the best scientific evidence, and an unbiased analysis that incorporates the wishes and needs of individuals.

Part 3 – Information systems in healthcare

The chapters in this part provide the technical core upon which all other parts depend. We introduce clinical information systems and their role in supporting the model, measure and

manage cycle. Second, it is shown that it is not always necessary to formalize this cycle completely, especially when flexibility in decision-making is needed. Consequently, many information processes are left in an unstructured or informal state and are more likely to be supported by communication processes.

The electronic health record is introduced next and is the first major technical system discussed in the book. The benefits and limitations of existing paper-based systems are compared with their electronic counterparts. Because the electronic patient record feeds so many different clinical systems, later topics including decision support, protocol-based care, population surveillance and clinical audit are all also introduced here.

The next two chapters cover the foundational informatics topics of how to design, evaluate and then implement working technological systems into complex sociotechnical organizations. It is often a conundrum that well-designed systems do not deliver the benefits expected. The evaluation chapter introduces the concept of the *value of information* and uses it to explain the value chain that starts with information creation and extends to ultimate benefit from its use. Understanding where a system is meant to deliver value along this chain becomes a recurrent motif in later chapters. This theme is expanded in the chapter on implementation, which looks at system implementation as a process of fitting technologies into complex adaptive organizations. The unexpected outcomes of technology sometimes can be explained only by stepping back and taking such a wider system view.

System safety is deeply linked to design and implementation decisions, and the potential downsides of information and communication technologies are explored next, given how closely related the concepts are in these chapters. The final chapter in this section takes another systemic perspective on clinical systems and the value of information, this time coming from economics. Although evaluation methods tell us much about the value of information, economics brings its own equally valid insights.

Having completed this part, one should be able to move on to any of the other parts in any order because each explores a more specialized topic area.

Part 4 – Guideline- and protocol-based systems

In this part, the various forms and uses of clinical guidelines, care plans and protocols are introduced. The different roles that computer-based protocol systems can play in clinical practice are outlined in the second chapter. These cover both traditional ‘passive’ support where protocols are kept as a reference and active systems in which the computer uses the protocol to assist in the delivery of care. For example, protocols incorporated into the electronic record can generate clinical alerts or make treatment recommendations. The growing evidence base for the benefit of such technologies is also summarized, emphasizing that benefits are more likely to be easily demonstrated in process rather than clinical outcome improvements. The third chapter reviews the process of protocol creation, dissemination and application and explores how informatics can create tools to assist at each of these stages.

Part 5 – Communication systems in healthcare

Although interpersonal communication skills are fundamental to patient care, the process of communication has, for a long time, not been well supported technologically. Now, with the widespread availability of communication systems supporting mobility, voice mail, electronic mail and social media, new possibilities arise. The chapters in this section introduce the basic types of communication services and explain the different benefits of each.

The second chapter in this part is probably the most technical of the book, covering information and communication networks and healthcare-specific networks such as health information exchanges. It is the place in this text where interoperability standards are covered in detail, as well as topics that relate to how information is accessed across networks, including privacy and consent.

Social media comprise a different class of communication system, and their importance is underlined with a discrete chapter devoted to them. The chapter introduces basic concepts from social networking theory and the social determination of health and then explores how social media are being harnessed across the spectrum of healthcare services.

The final chapter in this part examines clinical communication from the perspective of telemedicine and m-health technologies. The potential of such systems for different areas of healthcare is described, along with the accumulating evidence base for their success, again using the value chain as one way of understanding sometimes unexpected negative results.

Part 6 – Language, coding and classification

If the data contained in electronic patient record systems are to be analyzed, then they need to be accessible in some regular way. This is usually thwarted by the variations in health terminology used by different individuals, institutions and nations. To remedy the problem, large dictionaries of standardized clinical terms have been created.

The chapters in this part introduce the basic ideas of clinical concepts, terms, codes and classifications and demonstrate their various uses. The inherent advantages and limitations of using different terms and codes are discussed in the second chapter. The last chapter looks at some more advanced issues in coding and describes the theoretical limitations to coding. It introduces natural language processing and text mining methods and explains how the statistical approach to language management is complementary, and sometimes preferred, to the more formal semantic approaches used in clinical terminologies and ontologies.

Part 7 – Clinical decision support and analytics

Clinical decision support systems (CDSSs) are historically among the most powerful classes of informatics intervention we have at our disposal. These computer programs range from systems that simply present data to help a human make a decision to those that generate prompts or alerts when a clinician's decision appears problematic, through to systems with

the capability of making decisions entirely on their own. In the first chapter, the focus is on the different applications for CDSSs, particularly to see where clear successes can be identified. The next chapter takes a more technological focus and looks at the computational reasoning processes that underpin CDSSs. The final chapter in this part looks at how CDSS knowledge is created, through machine learning, data analytic and computational discovery methods.

Part 8 – Specialized applications for health informatics

The final chapters in this book explore some of the specialized ways that decision technologies are applied in clinical practice. These technologies find application in creating intelligent patient monitors or autonomous therapeutic devices such as self-adjusting patient ventilators. Along with communication technologies, CDSSs are essential components of public health and biosurveillance systems. In the field of bioinformatics, human genomic and metabolic knowledge is harnessed using computer techniques and reframes many classes of clinical decision as questions of genetics. When such bioinformatics knowledge is used in clinical practice, it is often described as personalized or precision medicine, and this topic is covered in its own chapter. The book concludes, not on a minor topic, but on one of the most transformational ones both for informatics and for healthcare delivery – the rise of consumer ownership and involvement in the process of care and the role that informatics has to play in making this necessity a reality.

References

- Coiera, E. (1995). Medical informatics. *BMJ* **310**(6991): 1381.
- Protti, D. (1995). The synergism of health/medical informatics revisited. *Methods of Information in Medicine* **34**(5): 441–445.

PART 1

Basic concepts in informatics

Models

A message to mapmakers: highways are not painted red, rivers don't have county lines running down the middle, and you don't see contour lines on a mountain.

Kent, 1978

Man tries to make for himself in the way that suits him best a simplified and intelligible picture of the world and thus to overcome the world of experience, for which he tries to some extent to substitute this cosmos of his. This is what the painter, the poet, the speculative philosopher and the natural scientist do, each in his own fashion ... one might suppose that there were any number of possible systems ... all with an equal amount to be said for them; and this opinion is no doubt correct, theoretically. But evolution has shown that at any given moment out of all conceivable constructions one has always proved itself absolutely superior to all the rest.

Einstein, 1935

The study of human health is based upon a few foundational concepts such as the cell or the notion of disease. Informatics is similarly built upon the basic concepts of data, models, systems and information. Unlike the study of health, in which core ideas are usually grounded in observations of the physical world, these informatics concepts are abstract ideas. For those used to the study of healthcare, informatics concepts often seem detached from the physical realities of the clinical workplace.

This issue is further complicated because we use the same words that describe informatics concepts in everyday language. It is common to ask for more information about a patient, or to question what data support a particular conclusion. In informatics these intuitive ideas need to be more precisely defined.

In this first chapter, we begin our study of informatics by exploring the pivotal concept of a *model*. Whether diagnosing a patient's illness, writing into a patient's record, designing an information system or trying to create a more efficient health service, we use models to shape and direct our actions. Models define the way we learn about the world, interpret what we see and apply our knowledge to effect change, whether that is through our own actions or through the use of technology such as a computer.

A map is not the territory it represents, but, if correct, it has a similar structure to the territory, which accounts for its usefulness.... If we reflect upon our languages, we find that at best they must be considered only as maps. A word is not the object it represents ... the disregard of these complexities is tragically disastrous in daily life and science (Korzybski, 1948).

Box 1.1 Therac-25

Between June 1985 and January 1987, Therac-25 linear accelerators operating in the United States and Canada famously delivered massive radiation overdoses to at least six patients, causing death or serious radiation injury. Patients received doses of up to 20000 rads where a dose of 200 rads was a typical therapeutic dose, and a 500-rad whole-body dose will cause death in 50 per cent of cases. At the time, these overdoses were arguably the worst radiation incidents in the history of radiotherapy.

Medical linear accelerators operate by creating a high-energy electron beam. The beam is focussed onto a patient to destroy tumour tissue and leaves healthy tissue outside the beam focus relatively unaffected. The high-energy beam produced by these devices is focussed by a collimator, usually made of tungsten. This 'flattens' the beam to therapeutic levels and acts like a lens to focus the beam to a tissue depth appropriate for a given patient.

In the Therac-25 accidents, the tungsten shield was not in place when the radiation dose was delivered. As a result, patients received a full dose of the raw 25-MeV electron beam. There were a number of different causes of the various overdoses, but each overdose essentially resulted from modelling errors in the system's software and hardware (Leveson and Turner, 1993).

One critical error resulted from the reuse of some software from a previous machine, the Therac-20. This software worked acceptably in the Therac-20, but when reused in the Therac-25, it permitted an overdose to be given. This was because although the Therac-20 had a physical backup safety system, this had been removed in the design of the Therac-25. The Therac-20 software was thus reused in the Therac-25 on the assumption that the change in machines would not affect the way the software operated. Therefore, software modelled to one machine's environment was used in a second context in which that model was not valid.

Another problem lay in the measurement system that reported the radiation dose given to patients. The system was designed to work with doses in the therapeutic range, but when exposed to full beam strength, it became saturated and gave a low reading. As a result, several patients were overdosed repeatedly because technicians believed the machine was delivering low doses. Thus, the measurement system was built upon an assumption that it would never have to detect high radiation levels.

All these failures occurred because of the poor way models were used by the designers of the Therac-25. They did not understand that many of the assumptions that were left implicit in the specifications of the device would quickly become invalid in slightly changed circumstances and would lead to catastrophic failure.


The motto for this chapter is 'A map is not the territory'. Humans are naturally adept at developing mental models of the world and manage to use them robustly, despite the inherent weaknesses of the models themselves. When these flexible mental models are transferred into a fixed technological system such as a computer, the effects of modelling error can be amplified significantly, with sometimes disastrous consequences (Box 1.1). This is because much of the knowledge used in creating the model has not been transferred along with it. As a consequence, the technological system is unable to define the limits of its knowledge. The 'map' in the computer is not the same as the territory of the workplace in which it is placed. One of the major ideas to be explored in this chapter is that the implicit and explicit assumptions we make at the time a model is created ultimately define the limits of a model's usefulness.

1.1 Models are abstractions of the real world


What then is a model? People are familiar with the idea of building model aeroplanes or looking at a small-scale model of a building to imagine what it will look like when completed. In health, models underlie all our clinical activities. For example, whenever we interact with

patients, we use internalized models of disease to guide the process of diagnosis and treatment.

Models actually serve two quite distinct purposes. The first use of a model is as a copy of the world. The modelling process takes some aspect of the world and creates a description of it. Imagine a camera taking a photograph. The image captured upon the camera's film is a model of the world:



We can generalize from the way a camera records a physical object to the way that all models are created. The process of creating a model of the real world is called *abstraction*:



Abstraction: the process of identifying a few elements of a physical object and then using these to create a model of the object. The model is then used as a proxy representation of the physical object.

The effects of abstraction are directly analogous to the effects of using a camera. In particular, the image captured by camera has four important features characteristic of all models.

- First, an image is simpler than the real thing. There are always more features in the real world than can be captured in an image. One could, for example, always use a more powerful lens to capture finer detail. Similarly, models are always less detailed than the real world from which they are drawn. A city map will thus not contain every feature of the streets it describes. Because models are always less detailed than the things they describe, data are lost in the abstraction process.
- Second, an image is a caricature or distortion of the real world. The three dimensions of the photographed object are typically transformed into a two-dimensional image. Through the use of different filters or lenses, very different pictures of the observed world are obtained. None of them is the 'true' image of the object. Indeed, there is no such thing. The camera just records a particular point of view. Similarly, abstraction imposes a point of view upon the real world, and the resulting model is inevitably distorted in some way. Thus, a map looks very little like the terrain it models. Some land features are emphasized, and others are de-emphasized or ignored. In physiology, one view of the heart is as a mechanical pump. This model emphasizes one particular aspect of the organ, but it is clearly much more than this. It also has a complex set of functions to do with the regulation of blood pressure, blood volume and organ perfusion.
- Third, as a consequence of distortion and data loss, many possible images can be created of the same object. Different images emphasize different aspects of the object or show different levels of detail. Similarly, because we can model a variety of aspects of any physical object, in variable detail, many models can be created. Indeed, the number of possible models is infinite. As we all carry different 'lenses' when we see the world, it is no surprise that different people see the world so differently. Some psychiatrists, for

example, may consider the brain from a Freudian or Jungian perspective. Neurologists may model it as a collection of neurones, each with different functions. Psychologists may model the function of a brain on that of a computer. When a clinician meets a patient, do they see a person, an interruption, a client, a task, a disease, a problem, a friend or a billing opportunity?

Karl Popper would ask his students to 'observe and describe'. They would be puzzled, and eventually ask 'observe what?' That was his point. We always have to observe something in order to describe something. The notion of pure observation, independent of direction, is a myth (Skolimowski, 1977).

- Finally, a camera records a particular moment in time. As the real world objects in an image change with time, their image is frozen at the moment the picture was taken. The difference between you and your photograph increases as you get older. The similarity between any model and the physical objects it represents also degrades with time. A map of a city becomes increasingly inaccurate as time passes because of changes to the city's roads and buildings. Computer programs that 'map' particular work processes slowly become less useful and need to be upgraded as work practices change around them.

As a consequence of these four characteristics of models, a final and central idea now becomes evident. All models are built for a reason. When we create a model, we actively choose among the many possible models that could be created, to try to build the one that best suits our particular purposes. For example, a driver's map emphasizes streets and highways. A hiker's map emphasizes terrain and altitude. Thus, one actively excludes or distorts aspects of the world to satisfy a particular purpose. There is no such thing as a truly 'general purpose' model. This last point is crucial to much of what follows in later chapters because it explains much about the challenges faced when trying to build computer systems that must model real world practices.

So just as a camera cannot capture a 'true' image of an object, one cannot ever build a 'true' model of an object. In philosophy, the argument against models ever being inherently correct is equivalent to arguing against the Platonic ideal. This is the idea that pure forms of physical objects exist outside the realms of the physical world. Although a physical sphere may always have an imperfection, Plato believed there existed an 'ideal' mathematical spherical form. The counterargument is that there is no such thing as ideal or objective truth in the world. There can only ever be our subjective and local point of view, shaped by the input of our senses.

Even in 'pure' geometry, there is no ideal sphere, just an infinite family of possible shapes that vary depending on the rules of the geometric system you choose. We cannot say that only one of these geometries is correct. Rather, they are different explanations of space, based upon different assumptions. We use the one that gives the most satisfactory explanation of the phenomenon we are interested in. For example, Riemann geometry works best for Einstein's relativity theory, rather than classic Euclidean geometry, because it handles the notion of curved space (Figure 1.1).

This philosophical argument continues into the debate about the basis of scientific enquiry and the nature of experimental evidence. A scientific hypothesis is nothing more than a model of some aspect of the world that is to be tested by an experiment. However, if a model can never be correct and objective truth can never be known, then our experiments can never actually *prove* anything to be true. We are never sure that we are right. The best that experiments can do is show us when our models of the world are wrong (Popper, 1976). What remain are theories that have best survived the tests set of them and are in some way more useful than others in interacting with the world.

And though the truth will not be discovered by such means – never can that stage be reached – yet they throw light on some of the profounder ramifications of falsehood (Kafka, 1931).

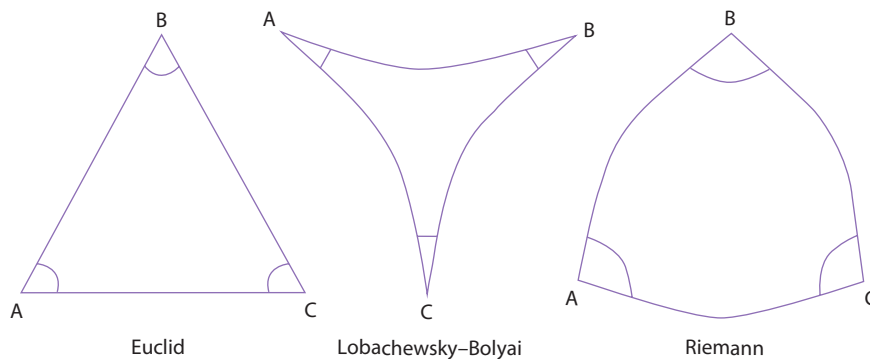



Figure 1.1
There is no 'correct' geometric shape, just an infinite number of possible geometries. The angles in classic Euclidian triangles always add up to 180° , but in Riemann geometry they always add to greater than 180° , and with Lobachewsky-Bolyai always less than 180° .


1.2 Models can be used as templates

So far, we have considered models as copies of the world. There is a second way in which we use models. Some models, rather than being copies of existing things, are used as templates from which new things will be created or that show how things are to be done. An architect creates a set of drawings that will be translated into a building. Economists build mathematical models of a country's economy and then use these to predict the effects of changes in monetary policy. A clinical protocol provides a template for the way a patient is to be managed.

Again, a simple example will make this second use of models clearer. Consider what happens when an image is projected onto a screen:



The stored image is a model of the real world. The projection process uses this model to create a second, slightly altered, display image. We can generalize from this to understand how models can act as templates. The process begins with the creation of a model. This could be a design, perhaps recorded as a set of blueprints or specifications. This is followed by a process of construction or model *instantiation*. A projected image is thus an instance of a stored image. In mathematics and logic, we instantiate the variables in an equation with data values. The equation is a template, and it interacts with supplied data values to arrive at a result. Instantiation uses a model as a template to build an artefact or process that is an instance of the model in the physical world:



Many of the consequences of the instantiation process are similar to those of abstraction:

- Although abstraction loses data to create a model, the process of instantiation adds data to create an instance. The image you see from a projector varies depending upon whether

Instantiation: the process of building an example or instance of a model, using the model as a template to guide the process.

it is projected on a white screen, a wall or the side of a building. The physical surface adds in its own features to shape the final result. The image is the result of the interaction of the projected image and the physical surface it strikes. This explains why implementing the same process in two different organizations yields different results. The local features of each organization uniquely shapes the way the process eventually works. Thus, an artefact is always more complex than the model that it came from because it is situated in the physical world.

- The constructed artefact is a distortion of the original template because the process of instantiation can transform it in many ways. A projected image can be shaped by the use of different filters and lenses to produce a variety of images.
- No two projected images are ever exactly the same because of variations introduced by the physical process of construction. No two physical artefacts are similar even if they are instances of the same template. Even mass-produced objects such as light bulbs, syringes or clay pots have minor imperfections introduced during manufacture that distinguish one instance of an object from another. In contrast, in digital or 'virtual' worlds, we can typically guarantee that the conditions for creating instance copies are identical.
- The effect of the captured image changes with the passage of time as the physical world changes. A movie usually has a greater impact on release than many years afterward as audiences change. A treatment guideline becomes increasingly inappropriate as time passes and new knowledge indicates that other methods are better treatment options.

Thus, the process of creating an instance has a variable outcome, and the impact of any instance of a model in the real world also varies. Two examples will help reinforce these ideas. Despite having identical DNA, two individual biological organisms are never truly identical. If DNA is a model (see Box 2.1), then the process of DNA transcription results in the 'manufacture' of an instance of an individual organism. Even if we take identical DNA as the starting point, local variations in the process of protein manufacture will introduce minor changes that at some level distinguish the clones.

Similarly, although two patients may be treated according to the same guideline, which is a template for treatment, no two actual episodes of treatment are ever exactly the same. Variations in the timing of treatments, availability of resources and the occurrence of other events all can conspire to change the way a treatment is given. Equally, a patient's physical and genetic variations may result in variations to the results of a treatment. The features of the specific situation in which a treatment is given thus result in variations in the way the treatment proceeds and in its final effects upon a patient.

A more general principle follows from these four characteristics of templates. Because the process of creating an instance from a template has a variable result, and the process of doing things in the real world is uncertain because we can never know all the variations that are 'added in' as we follow a template, there is no such thing as a general purpose template. All we can have are templates or designs that are better or worse suited to our particular circumstances and are better or worse at meeting the needs of the task at hand.

As we will see in later chapters, this means that there can be no 'correct' way to treat an illness, no 'right' way to describe a diagnosis and no 'right' way to build an information or

communication system. There can never be an absolutely ‘correct’ design for a treatment protocol or a ‘pure’ set of terms to describe activities in healthcare. This principle explains why clinical protocols will always have varying effectiveness based upon local conditions and why medical languages can never be truly general purpose. What we do have are treatments, protocols, languages, information and communication systems that are better or worse suited to our specific purpose than others at a moment in time.

1.3 The way we model the world influences the way we affect the world

So far we have seen how models act either as copies of things in the world or as templates upon which new things are created. These two processes are deeply interrelated.

In photography, decisions about form and content at the moment an image is created ultimately influence the way it can be used and how useful it will be. Thus, if a picture is intended to be of print quality, it may require a higher-definition image to be taken than one intended for display on a computer screen.

When artefacts are created, it is assumed that they will be used for a particular purpose. If the purpose changes, then a design becomes less effective. Thus, the physical design of the waiting room and treatment areas for a general practice clinic will assume that a certain number of patients are seen during a day and that certain kinds of therapy will be given. If the clinic was bought by radiologists, they would have to remodel the clinic’s design to incorporate imaging equipment and to reflect a different throughput of patients.

If a disease is based upon assumptions about the incidence of a disease in a given population, then it may not work well in a different setting. Treating infant diarrhoea in a developed nation is not the same task as in underdeveloped nations, where poorer resources, malnutrition and different infecting organisms change the context of treatment. Before a model is used, one therefore has to be clear about what has been assumed.

Similarly, a set of rules and procedures may be developed in one hospital and may be spectacularly successful at improving the way the hospital functions. One would have to be very cautious before imposing those procedures on other hospitals, given that they implicitly model many aspects of the original institution. Very small differences in the level of resources, type of patients seen or experience of the staff may make what was successful in one context unhelpful in another.

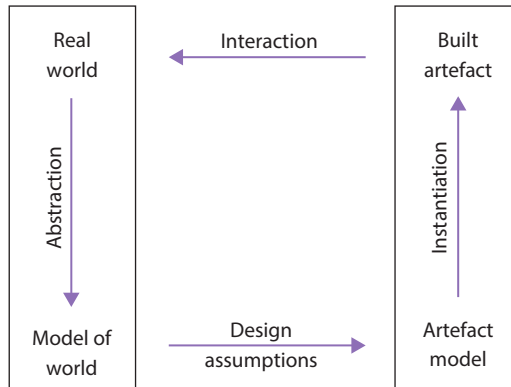
More generally, any designed artefact, whether it is a car, a drug or a computer system, has to be designed with the world within which it will operate in mind. In other words, it has to contain in its design a model of the environment within which it will be used. These specifications constitute its *design assumptions*. Thus, there are direct connections among the process of model creation, the construction of artefacts based upon such models and their eventual effectiveness in satisfying some purpose (see Figure 1.2).

A few examples should make the cycle of model abstraction and instantiation clearer. First, consider a car. The design blueprints of a car reflect both its purpose and the environment within which it will operate. The car’s engine is built based upon the not unreasonable assumption that it will operate in an atmosphere with oxygen. The wheels and suspension

Before the work of the famous physician Galen, it was assumed that the arteries contained air. This was because arteries were observed to be empty after death (Schafer and Thane, 1891). The physicians making these observations had thought they had created a model of arterial function in living humans, but all they had created was a model valid in cadavers. Therefore, the context in which a model is created affects its validity for any other context within which it may be used.

There are no side effects – only effects. Those we thought of in advance, the ones we like, we call the main, or intended, effects, and take credit for them. The ones we did not anticipate, the ones that came around and bit us in the rear – those are the ‘side effects’ (Stermen, 2002).

Figure 1.2
Models of the world are used as a template to define how artefacts such as devices or processes will be constructed.

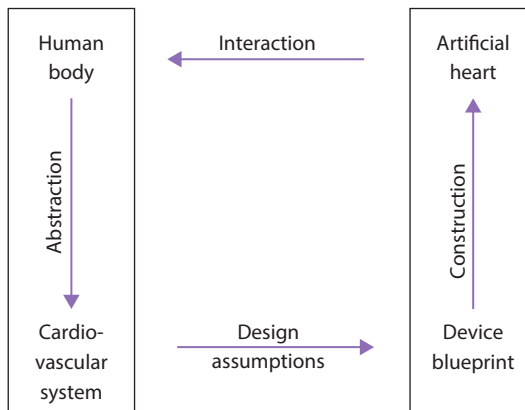


may be designed with the assumption that they will operate on a highway or urban street. If the car was put into another physical environment such as a desert or the lunar surface, it probably would not work very well. Sometimes such design assumptions are left implicit, and they become obvious only when a device is used in a way in which it was not intended, sometimes with catastrophic results (see Box 1.1).

The human body also makes assumptions about its environment. The haemopoietic system adjusts the number of red blood cells needed for normal function based upon the available oxygen in the atmosphere. As a consequence, individuals living at sea level have calibrated their oxygen carrying system differently from those living in high altitudes. An athlete training at sea level will not perform well if moved quickly to a high altitude because these ‘working assumptions’ are no longer met.

Finally, consider an artificial heart (Figure 1.3). Such a device must model the biological heart in some way because it will replace that organ within the cardiovascular system. The artificial heart thus is based upon a model of the heart as a mechanical pump and is designed with the assumption that supporting the pump mechanisms will be beneficial. The better it models all the functions of a true heart, the better a replacement it will be. It is also designed on the assumption that it will need to be implanted, and as a consequence it is crafted to survive the corrosive nature of that environment and to minimize any immune reaction that could be mounted against it.

Figure 1.3
An artificial heart is based upon two kinds of model. First, the cardiovascular system has to be modelled, and second, a mechanical blueprint is used to model the way the artificial heart will be constructed.



Conclusions

In this chapter, the foundational concept of a model is explored in some detail. Models underpin the way we understand the world we live in, and as a consequence they guide the way we interact with the world. We should never forget that the map is not the territory and the blueprint is not the building.

In the next chapter, a second basic concept of information is introduced. These two ideas are then brought together, as we begin to see that knowledge is a special kind of model and is subject to the same principles and limitations that are associated with all other models.

Discussion points

1. 'The map is not the territory'. Why not?
2. Observe and describe. Compare. Why?
3. Biologists have argued whether nature, expressed in an organism's DNA, or 'nurture', via the physical world, is most important in shaping development. If 'nature' is the template and 'nurture' creates instances of an organism, which is more important in shaping an organism based upon the first principles of modelling?
4. In what ways could the limitations of models result in errors in the diagnosis or treatment of patients? Use Figure 1.2 as a template to guide your thinking, if it helps.

Chapter summary

1. Models are the basis of the way we learn about, and interact with, the physical world.
2. Models can act either as copies of the world like maps or as templates that serve as the blueprints for constructing physical objects or processes.
3. Models that copy the world are abstractions of the real world.
 - a. Models are always less detailed than the real world from which they are drawn.
 - b. Models ignore aspects of the world that are not considered essential. Thus abstraction imposes a point of view upon the observed world.
 - c. Many models can be created of any given physical object, depending upon the level of detail and point of view selected.
 - d. The similarity between models and the physical objects they represent degrades over time.
 - e. There is no such thing as a truly general purpose model. There is no such thing as the most 'correct' model. Models are simply better or worse suited to accomplishing a particular task.
4. Models can be used as templates and be instantiated to create objects or processes that are used in the world.
 - a. Templates are less detailed than the artefacts that are created from them.
 - b. An artefact is a distortion of the original template.

- c. No two physical artefacts are similar even if they are instances of the same template.
 - d. The effect of an artefact may change while the original template stays the same.
 - e. The process of creating an instance has a variable outcome, and the impact of the instance of an artefact in the real world also varies. As a consequence, there is no such thing as a general purpose template. All we can have are templates or designs that are better or worse suited to our particular circumstances and tasks.
5. The assumptions used in a model's creation, whether implicit or explicit, define the limits of a model's usefulness.
- a. When models are created, they assume that they are to accomplish a particular purpose.
 - b. When models are created, they assume a context of use. When objects or processes are built from a model, this context forms a set of design assumptions.
6. We should never forget that the map is not the territory and the blueprint is not the building.

Information

The plural of datum is not information.

Anon

To act in the world we need to make decisions, and to make decisions we must have information that distinguishes one course of action over another. In this chapter, a basic framework is presented that defines what is meant by such information. Whether delivered in conversation, captured in a set of hand-written notes or stored in the memory of a computer, the same basic principles govern the way all information is structured or used. The ideas presented here build upon the concept of models developed in Chapter 1. The simple ways that models, data and information interrelate are then unfolded. It then becomes apparent that models and information underpin not just the specialized study of informatics, but also every aspect of the delivery of healthcare.

2.1 Information is inferred from data and knowledge

Informally we know that we have received information when what we know has changed. In some sense this new information must be measurable because intuitively some sources of information are better than others. One newspaper may generally be more informative than others. A patient's medical record may be full of new data, but to the clinician who sees the patient every day it may contain little new information.

Formally, information can indeed be linked to concepts of *orderliness* and *novelty*. The more order in a document, the more 'information' it contains. For example, a patient's record that is broken up into different sections such as past history and allergies is more informative than an unstructured narrative that jumbles the patient's details. Equally, if a patient's record contains nothing new, then it conveys no new information. Statistical measures for the amount of such 'information' communicated by a source are the basis of *information theory* (see Box 4.1). However, these statistical measures of information can only partly help us to understand the computational meaning of information, which underpins informatics.

Terms such as data, information and knowledge are often used interchangeably in common speech. Each of these terms has a quite precise and distinct definition in the information sciences.

Information is constructed by people in a process of perception; it is not selected, noticed, detected, chosen or filtered from a set of given, static, pre-existing things. Each perception is a new generalization, a new construction (Clancey in Steels and Brooks, 1995).

Data consist of facts. Facts are observations or measurements about the world. For example, ‘today is Tuesday’, the ‘blood pressure is 125/70 mm Hg’ or ‘this drug is penicillin’.

Knowledge defines relationships between data. The statement ‘penicillin is an antibiotic’ relates two data elements. The rules ‘tobacco smoking causes lung cancer’ and ‘if a patient’s blood pressure is greater than 135/95 mm Hg on three separate occasions, then the patient has high blood pressure’ are more complex examples of knowledge. Knowledge is created by identifying recurring patterns in data, for example across many different patients. We learn that events occur in a certain sequence or that an action typically has a specific effect. Through the process of model abstraction, these observations are codified into general rules about how the world is and works.

As well as learning generalized ‘truths’ about the world, knowledge can be specific to a particular circumstance. For example, *patient-specific knowledge* comes from observing a patient’s state over time. By abstracting observed patterns, one can arrive at specific knowledge such as ‘following treatment with antihypertensive medication, there has been no decrease in the patient’s blood pressure over the last 2 months’.

Information is the meaning obtained by the application of knowledge to data. Thus, the datum that ‘blood pressure is 125/70 mm Hg’ provides information only if it tells us something new. In the context of managing a patient’s high blood pressure, by using general knowledge of medicine and patient-specific knowledge, the datum may lead to the inference that the patient’s previously high blood pressure is now under control, which is indeed new information.

We can now see how these three concepts are related. By using a piece of knowledge, in a given context, data are interpreted to produce information. The data stored in DNA are interpreted by cellular molecules to create proteins (Box 2.1). Another example may make this even clearer. Imagine someone is speaking to you in a language that you do not

Box 2.1

DNA is just data

Conceptualizing an information system into model, data and interpretation components has a certain universality. In biology, there is a strong information paradigm arising out of our understanding of the role of DNA.

Since its structure and function began to be unfolded, DNA has been seen as some form of master molecule, dictating the development of individual organisms. The doctrine of DNA has perhaps reached its most extreme position in the notion of the selfish gene (Dawkins, 1982). DNA is characterized as clothing itself in cells, which allow DNA to survive and reproduce from generation to generation. DNA, in this view, creates and dictates the development and activity of organisms. The organism is merely the survival machine used by the genetic sequence.

Another view sees DNA as part of a far more complex system. DNA is among the most non-reactive and chemically inert molecules in biology. It thus is perfectly designed for its role, which is to store instructions, much like the memory in a computer. DNA is a kind of database and nothing more.

Thus, although DNA stores the models used to create proteins, it is of itself incapable of making or doing anything. That is the role of the cellular machinery (see Box 30.2). Whereas it is often said that DNA produces proteins, in fact proteins produce DNA (Lewontin, 1993).

The symbolic language of DNA and thus the ability to interpret DNA reside in the surrounding cellular structures. Without these molecules, there would be no way that we could decode the symbolic meaning of DNA – the data stored in the DNA would be uninterpretable. In other words, an organism’s DNA has no meaning outside the context of the cellular structures that contain it.

We can thus regard a complex organism as being the result of a cell’s interpretation of the data stored in the DNA database, by using a language encoded within its proteins and within the context of the data provided by the intra-cellular and extra-cellular environment.

understand. You have received a large amount of data during that conversation, but because you have no knowledge of the language, it is meaningless to you. You cannot say that you have received any information. For the same reason, when two people with different life experiences and knowledge read the same book, they can come up with very different interpretations of its meaning.

2.2 Models are built from symbols


Knowledge is the set of models we use to understand and interact with the world. Sometimes these models are stored as physical analogues of the real thing. A scale model of a township may display planned developments alongside existing structures. The pattern that iron filings make on a piece of paper models the magnetic field created when a magnet is placed underneath. Knowledge is also stored in the heads of people. With the development of language and writing, it became possible for these models to be shared and to evolve from being purely mental constructions to something we can examine and manipulate in the physical world.

A weather map, for example, describes processes that look nothing like its diagrammatic representation. In the realms of science and mathematics, it is common to create models in the form of diagrams or equations. These models are created from a set of symbols which are markings that are used to represent something else. When people talk about knowledge, they usually refer to symbolic models, and that is the sense in which knowledge is used here.

A fundamental characteristic of all symbolic models is that, on their own, they have no intrinsic meaning. The equation $e = mc^2$ is meaningless unless each of the letters in the equation is named and the concepts the letters stand for are explained. Mathematical operations such as equality and multiplication also must be understood. A weather map is equally mysterious without such definitions.

Symbolic models thus gain their meaning when we associate concepts with individual symbols. Specifically, symbolic models are built using a recognized *terminology* and a set of relationships or *grammar* that connects the terms (Figure 2.1). Together the terminology and these grammatical relationships constitute a *language*. In the information sciences, the languages used to create models are usually based upon logic or mathematics.

A terminology contains all the symbols that can be used in building a model, and it maps these symbols to particular concepts, just like a dictionary. For example, in healthcare we have



Terminology: a standard set of symbols or words used to describe the concepts, processes and objects of a given field of study.

Figure 2.1
Symbolic models cannot be understood unless the symbol language and the possible relationships among the symbols are also understood.

Grammar: the set of rules that together specify the allowed ways an alphabet can be put together to form strings of symbols in a language.

specific words or terms that stand for observable events such as diseases or treatments. The term ‘angina’ maps to a data cluster of symptoms and signs that can be observed in a patient.

The set of meaningful relationships between symbols is captured in grammar. In English, for example, the language’s grammar makes sure that when one person strings together a word sequence, others will recognize the intended meaning of the sequence. The same words in two different sequence orders may have very different meanings.

2.3 Inferences are drawn when data are interpreted according to a model

We use symbolic models to reason about the world. A symbolic model is applied to data acquired from the world to come to a conclusion about how things are. Lawyers examine data in a client’s case notes against their knowledge of law to reach conclusions about how likely a client is to succeed in court. Clinicians take data in the form of patients’ observations and measurements and use knowledge about disease and therapy to infer what illness a patient may have and the most prudent course of action.

This process of data interpretation actually requires different kinds of models. Specifically, we need a *database*, a *knowledge base*, an *ontology* and an *inference procedure*. Let us look at each of these in turn.

Interpretation starts with a collection of data. The data may be numbers from a laboratory test. The numbers themselves are just symbols, however, and have no meaning on their own in the same way that points on a graph have no meaning without labels on the axes (Figure 2.2). Hence a terminology is needed, and each datum is associated with a label or term drawn from the terminology. For example, the number ‘7.4’ is associated with the label ‘pH’, giving us the datum ‘pH = 7.4’.

More technically, using what is known as an *entity-relationship* (ER) model for our database design, we can say that the **entity** ‘pH’ is related to an **attribute** ‘test result’ and that the result has a **value** of ‘7.4’. Another attribute of pH may be the date of the test:

pH	
Test result	7.4
Test date	3 January 2015

This set of labels and relationships is called the data model. Together, a collection of data and their associated data model are called a *database*.




Figure 2.2
Data remain uninterpretable in the absence of a language that defines what each datum represents.

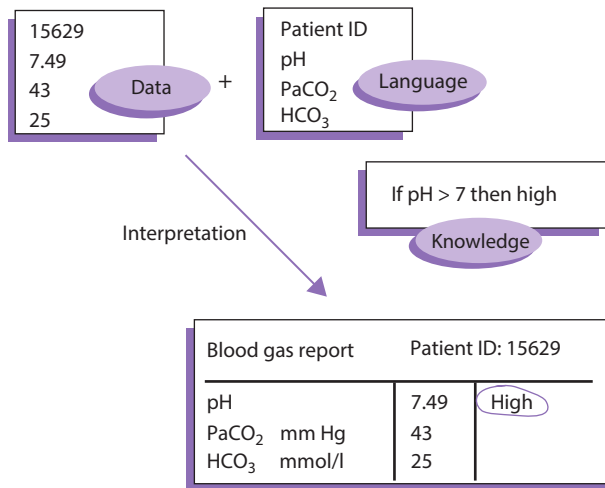


Figure 2.3
Data are interpreted with reference to a data model, a knowledge base and rules of inference.

The process of interpretation next requires knowledge about the different ways the concepts in the database relate to each other. Rules such as ‘if the pH is greater than 7.4, then it is abnormally high’ or ‘if acidosis is present, then treat with intravenous bicarbonate’ may be part of a clinician’s knowledge of acid-base physiology (Figure 2.3). A collection of rules can be thought of as a database containing elements of knowledge, or a *knowledge base*:

Acidosis	
Definition	pH < 7.0
Treatment	Bicarbonate therapy

Just as a database needs a data model to make the data it holds intelligible, a knowledge base itself needs a knowledge model. Called an *ontology*, this knowledge model can be thought of as a dictionary of the allowed concept types in a knowledge base. The ontology may also specify all the ways in which such known concepts can be meaningfully arranged. For example, an ontology may contain concepts such as ‘disease’ and ‘treatment’. It may also specify that any knowledge about these concepts must take the form of a rule with the structure ‘if *disease* then *treatment*’ – this would stop nonsense rules such as ‘if *penicillin* then *acidosis*’. In philosophy, ontology is the study of what ‘is’. In informatics, an ontology pragmatically becomes a model of what is allowed to be said and how it is to be represented:

Ontology: the set of concepts understood in a knowledge base and the rules about how these concepts are allowed to be arranged meaningfully.

Disease	
Caused by	Pathological process
Treated by	Therapeutic agent

A fourth model defines the *rules of inference* that specify how a knowledge base can be applied to a database. For example, a rule of inference could say that a statement of the form ‘if X then Y’ means that when we know X is true, we can then also believe that Y is true. This particular example is the rule of logical *deduction* and is one of many rules of logic that are

used to make inferences (see Chapter 8). There are many different methods of inference beyond classical logic. Health epidemiologists make inferences using the rules of statistics. Lawyers have rules based upon precedent established in prior case law.

Together, a data model and ontology provide the grammar or *syntax* that defines relationships among data. The rules of inference are then used to interpret the meaning or *semantics* of the data.

2.4 Assumptions in a model define the limits to knowledge

In Chapter 1, we saw that assumptions made at the time a model is created affect the way it is used. The way a model is constructed, the context within which it is defined, what is included in it and the purpose for which is intended all affect its ultimate usefulness.

This is also the case for the models that define our knowledge of the world. The implication then is that the inferences we draw from a model are strongly influenced by the assumptions made when the knowledge in the model was first created.

For example, it is common for clinical protocols to define a standard way in which a particular illness is to be treated. Such a protocol is a kind of template model that drives the treatment delivered to a patient. When a protocol is created, its designers make many assumptions, not all of which are obvious at the time. For example, protocol designers may make an implicit assumption that the drugs or equipment they include in the protocol will actually be available and affordable.

What they are actually doing when they make such assumptions is to model the environment within which the designers expect the protocol to operate. This model usually matches their own local environment, and it is only when a designer is forced to check with others facing different circumstances that such implicit environmental assumptions are unearthed.

Thus, a protocol created for a well-equipped modern hospital may not be useful in a primary care clinic, where staff expertise and resources are very different. Equally, a protocol may assume that a patient has no other significant illnesses. This implicit assumption may be exposed when the treatment specified cannot be used because it interacts with a patient's other medications.

With these examples in mind, we can cast the creation and application of knowledge into the same form as the cycle of model creation and application developed in the Chapter 1 (Figure 2.4). First, the process of model abstraction is equivalent to the knowledge acquisition process. Observations made of the world are generalized into a model that describes how different parts of the world interrelate. Recall that such models are always limited, and they emphasize certain observations and omit others. Next, the knowledge model is applied to data. We can view this process as the construction of an inference, based upon a template model that represents our knowledge and a set of data. As we have just seen, design assumptions at the time a model is created affect how it can be used.

One special design assumption associated with a symbolic model is its language. Just as a photographic image cannot be used unless the right display software is available, a symbolic model cannot be used unless the right language and modelling relationships are also available. In other words, the modelling language becomes a design assumption, which needs to be explicitly catered for when the model is used.

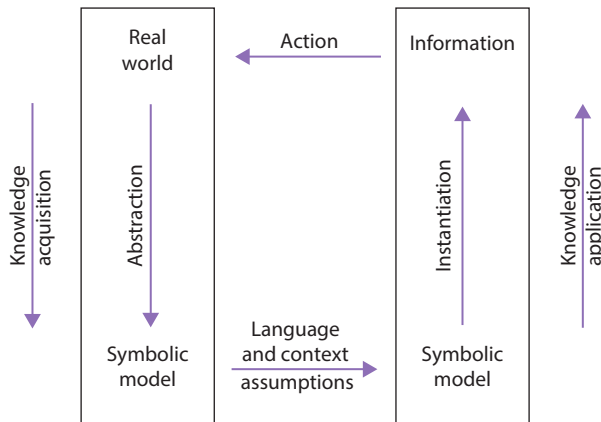


Figure 2.4 Knowledge is acquired through the construction of models, and these models are then applied to data, to draw interpretations of the meaning of the data.

Thus, a treatment protocol designer must assume that the people who read the protocol will be capable of understanding its language and form. If a clinician does not recognize the terms used in a protocol and is not familiar with the concepts and principles it is based upon, then they will be unlikely to understand the intentions of those who wrote the protocol.

2.5 Computational models permit the automation of data interpretation

If the knowledge and data components of a decision problem can be written down, then this problem can in principle be solved using a computer. At other times the task of data interpretation is shared between human and computer. For example, the computer may organize and consolidate data into a graphical presentation, and the human then examines the processed data to make a final interpretation. The proportion in which models are stored either in the computer or as mental models in the head of a human determines where the interpretation takes place. Computer systems thus form a spectrum, ranging from those that have no ability to assist in the interpretation of data to those that are able to carry out a complete interpretation, within the bounds of a given task (Figure 2.5).

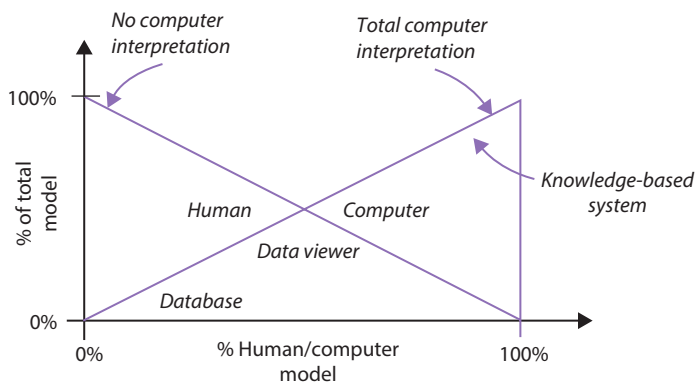


Figure 2.5 Humans and computers can share the burden of data interpretation. The amount of interpretation delegated to the computer depends upon how much of the interpretative model is shared between human and computer.

Computers can act as data stores

If a computer is used solely as a repository for data, it acts as a database. Data are organized according to a data model, so that the origin of each datum is recognizable. Medical data often consist of images or physiological signals taken from monitoring devices. As a consequence, the databases that store complex data from patients can be huge.

Computers can generate data views that assist interpretation

In contrast to passive databases, a computer can carry out some degree of interpretation by generating a view of the data. The computer system shows a user only that portion of data that is of immediate interest and in a way that is best suited to the task at hand.

Consider a database such as Medline, which holds publication data and abstracts of articles from research journals. Given the thousands of papers published weekly, there would be little value for a researcher simply to inspect each record in the database. It would be practically impossible to locate an article within it. What is required is a way of viewing a relevant subset of the data that matches the researcher's question.

For a computer system to provide such 'views' of stored data, a model of the user's needs has to be conveyed to the computer. This communication between the database and the human may be provided by what is known as a *query language*. This is a method commonly used to search library catalogues. Using special words such as 'and', 'or' and 'not', the user constructs a question for the query system about the things that should be displayed (see Box 6.1). Because it recognizes these special words and their meanings, the query system is able to retrieve those records from the database that match the terms provided by the user.

A patient's physiological monitor is also a kind of viewer. If a clinician looked at the raw data stream coming from a measurement device such as an electrocardiogram (ECG), then he or she would be confronted with streams of rapidly changing digits that would be completely unusable in a clinical setting. It is the monitor's role to present sensor data as a set of waveforms for a measurement such as the ECG, or as averaged numeric values for a measurement such as blood pressure. To do this, the computer must have models of the signals and of the kinds of noise and artefact that may corrupt the signals. It also needs a model of the preferred ways of displaying signals to allow humans to carry out the interpretation.

The same process occurs with computer-generated images. Computed tomography and magnetic resonance imaging, for example, both depend on complex models that reconstruct raw data into images that can be interpreted by clinicians. By varying the model parameters applied to the data, the imaging systems can produce different views, or 'slices' through the raw image data.

In general, the degree of division of responsibility for data interpretation between human and computer varies for a number of reasons. It may be that it is inherently difficult to formalize all the knowledge used in the interpretation of data, or it may simply be that the effort involved in modelling is greater than the reward. This is often the case when problems are rare or highly variable.

Computers can be responsible for all data interpretation

As the understanding of how knowledge could be represented in a computer developed, it soon became clear that computers could be used to perform quite powerful forms of reasoning on their own.

Computers are often used as a backup on reasoning tasks that are typically performed by humans. For example, in safety-critical situations such as the operation of a nuclear power plant, the computer watching over the complex system provides a second pair of ‘eyes’ looking over the operator’s shoulder. Computers are also used to interpret data when the task is routine but frequent enough that automation would help. The interpretation of laboratory test results is a common example of this, although typically a human audits the results of such interpretation. Biomedical devices are also often capable of autonomous interpretation. A pacemaker may analyze cardiac activity to look for the development of an arrhythmia.

In all these cases, the interpreting computer does not just possess data and the knowledge that will be used to interpret the data. It also needs a model of the way the computer should ‘think’ about the problem (i.e. the rules of interpretation discussed earlier). For example, a computer’s ‘inference engine’ may use rules of formal logic – most *knowledge-based systems* are built in this way. Sometimes the systems reason with rules of mathematics, probabilities or other more modern techniques such as fuzzy reasoning or neural networks.

Conclusions

In this chapter, we use the idea of a model to define data, information and knowledge and arrive at a rich understanding of everything from the way people draw conclusions to the role that DNA plays in the cell.

Chapters 1 and 2 are a prelude to introducing a third fundamental informatics concept – the notion of a system. In Chapter 3, the discussion introduces the concept of systems and leads to an exploration of what it means to create an information system. In this way, one can begin to understand the ways in which information systems can be forces for good, as well as understand some of their inherent limitations.

Discussion points

1. Take a patient’s laboratory result sheet. Rewrite the information there into a database of facts and show the data model. If the results have been flagged or interpreted in some way, write out what you think was in the knowledge base that was used to make the interpretation. What rules of inference were applied?
2. Compare your answer to the previous question with someone else’s answer. Why may there be differences in your answers? (Think back to Chapter 1).
3. Explain how an individual who reads the same patient’s record on two different occasions can find the first reading full of information and on the second reading find no information at all.
4. There are many ways in which we can model the world, and it is not surprising that there are many different ways that data can be represented in a database. This chapter describes entity-relationship models. What other database models are commonly used, and how might you choose amongst them?

5. Explain the role of hashtags in a Twitter message or of metadata on a Web page.
6. Find the names of some widely used terminologies and ontologies in healthcare or biology.
7. The human genome is the Rosetta stone needed to decipher the origin of human disease. Discuss.
8. A picture is apparently worth a thousand words. Assume each word is six letters. How big a picture do we get, in bits? Hint: How many bits are needed to encode each letter of the alphabet?
9. If everything we understand is subjective, what does it mean to 'know' something?

Chapter summary

1. Information is derived from data and knowledge.
 - a. Data are collections of facts.
 - b. Knowledge defines relationships among data.
 - c. Information is obtained by applying knowledge to data.
2. Knowledge can be thought of as a set of models describing our understanding of the world.
 - a. These models are composed of symbols.
 - b. A symbolic model is created using a language that defines the meaning of different symbols and their possible relationships among each other.
3. Inferences are drawn when data are interpreted according to a model.
 - a. Data on their own have no intrinsic meaning.
 - b. A language identifies concepts within the data.
 - c. Next, the knowledge stored in a model can be used to draw an inference from the labelled data.
4. This process of data interpretation actually requires different kinds of information model. Specifically we need a *database*, a *knowledge base*, an *ontology* and an inference procedure.
5. Assumptions in the knowledge model affect the quality of the inferences drawn from it.
 - a. Assumptions may implicitly define the context within which the model was created.
 - b. These design assumptions include the language used if the model is symbolic.
6. Knowledge acquisition and application are examples of the cycle of model abstraction and template-based construction.
7. Once a model and data have been sufficiently formalized, the interpretation can be automated using a computer.
 - a. Computers can store data according to data models.
 - b. Computers can provide different views onto data according to user models.
 - c. Computers can interpret data when they have a knowledge base and an inference procedure.

Information systems

One doesn't add a computer or buy or design one where there is no system. The success of a project does not stem from the computer but from the existence of a system. The computer makes it possible to integrate the system and thus assure its success.

C. Caceres in Dickson and Brown, 1969, p 207

Computers and automation have captured man's imagination. That is to say, like the psychiatrist's ink blot, they serve the imagination as symbols for all that is mysterious, potential, portentous. For when man is faced with ambiguity, with complex shadows he only partly understands, he rejects ambiguity and reads meaning into the shadows. And when he lacks the knowledge and technical means to find real meanings in the shadows, he reads into them the meanings in his own heart and mind ... Computers are splendid ink blots.

Simon, 1965

A system is commonly understood to be a routine or regular way of working. One can have a system for betting on a horse race or a filing system for storing and retrieving documents. These types of systems are models providing templates for action in the world. Systems can also be models of aspects of the world, such as an ecosystem. In this chapter, we first introduce the general topic of systems and identify a few of their key characteristics. The main focal points of the chapter are what is formally meant by an information system and how information systems are used to control the way decisions are made.

3.1 A system is a set of interacting components

Just as there is ambiguity with the normal meaning of words such as data or information, the notion of a system is equally nuanced. Systems pervade healthcare, and there are countless examples of them. In physiology, one talks of the endocrine system or the respiratory system. In clinical practice, we develop systems for questioning and examining our patients. Indeed,

the whole of healthcare itself is often described as a system. The simplest thing connecting these examples is that *each system consists of a collection of component concepts, processes or objects*.

We saw in Chapter 1 that models are the basis for building artefacts and interacting with the world. We can use these ideas to understand that the collection of entities we call a system can be one of three things:

- A model acting as an abstracted description of a set of objects or processes observed in the real world.
- A model consisting of several interlinked elements acting as a template to action.
- An artefact constructed by the process of instantiating the template in the real world.

So when you read the words ‘the health system’, there are three different possible meanings. The health system could be someone’s description of how he or she sees healthcare, based upon observation of the world. It could be a proposal or plan for how healthcare should work, or it could be the physical collection of people, buildings and infrastructure that collectively come together to deliver healthcare.

As abstract descriptions of the world, like all models, systems help compartmentalize some portion of the world in a way that makes it more understandable. For example, a collection of anatomical structures whose function is closely related may constitute a system. Thus, we speak of the nervous system and collect within its definition organs such as the peripheral nerves, the spinal cord and the brain.

The collecting of such elements into a system is a powerful way of enhancing our understanding of the way things work. When Harvey first proposed the notion of a system of circulation for blood through the body in 1628, he essentially constructed a model that, for the first time, connected the arteries, veins and heart into a functioning whole. So powerful was this model that it was adopted despite its inadequacies at that time. It was, for example, not until 1661 that Malpighi finally demonstrated that the capillaries connected the arterial and venous systems (Schafer and Thane, 1891).

As templates or blueprints, systems allow us to develop clear models of how entities will interact. There are separate blueprints for different systems in a building, including the electrical system, the plumbing system and the air-conditioning system. *Complexity* often refers to the number of interconnections that exist among such system components, and complexity increases with the number of interconnections. A *complex system* is thus one in which there are a relatively large number of connections among components. Complex systems are sometimes informally distinguished from a *complicated system*, which may have a large number of components but relatively few interconnections.

The discipline of decomposing a composite structure into sub-elements that carry out different functions simplifies the design process and will more likely result in something that is well designed for its intended purpose. Most machinery, for example, is designed as a set of complicated modules, to limit interconnection and minimise complexity. Modular systems are common features of modern technology. Cars have separate fuel, suspension and brake systems, and this separation into component systems means that each can be maintained and repaired fairly independently of the others.

3.2 A system has an internal structure that transforms inputs into outputs for a specific purpose

Systems have inputs and outputs

Systems usually have a set of *inputs*, which are transformed by the components of the system into a set of *outputs* (Figure 3.1). The inputs to a coffee maker are ground coffee, water and heat. The output is hot liquid coffee. The inputs to a leaf are water and carbon dioxide, and the photosynthetic system produces as output oxygen and carbohydrates. The inputs to a hospital emergency room are the clinical staff, patients and supplies – the outputs are patients who have been in some way transformed by their emergency room visit.

In the process of transforming inputs into outputs, the *state* of a system may change. For example, a car may have as its state a location, and the input of fuel, oil and a driver transforms the state of the vehicle to a different location. An influx of patients may rapidly change an emergency room from the state ‘able to receive patients’ to ‘full’. Further, the state of a system may determine just how it will change inputs into outputs. If a human’s state is ‘physically fit’ then that person transforms food into energy more efficiently than someone in the state ‘unfit’, who with the same food inputs may tend to deposit more of the food as fat.

Systems have behaviour

For a system to be distinguishable from its environment, it should have a characteristic behaviour. Thus, a weather system may be a set of related pressure bands, whose collective behaviour stands out and demands attention. One of the behaviours of the vascular system is that it can be observed to contain flowing blood. This behaviour distinguishes the vascular system from, say, the nervous system. It is this identification with a particular behaviour that helps conceptually separate a system from other parts of its environment.

Critically, a system’s behaviour cannot usually be predicted by an examination of its individual components, but it emerges from the way the components interact with each other. Thus, the behaviours of social groups, communication networks, roads and neural processes are all emergent properties of the interactions among their components (Box 3.1). The emergence of system behaviour means that the effects of anything we build are not directly predictable from an examination of individual components, but rather from an understanding of how the components interrelate.




Figure 3.1

A system is characterized by a set of inputs and outputs and can be internally decomposed into a set of interacting components or sub-systems.

Physical systems are embedded in an environment

To be distinguishable from everything else, there must be a *boundary* between a system and the rest of the environment. However, that boundary may be difficult to define precisely. Depending upon the integrity of the boundary, three different kinds of systems are possible:

- *Closed systems* have no external inputs and outputs, and they behave like a black box that is unaffected by the external world.
- *Relatively closed systems* have precisely defined inputs and outputs with their environment.
- *Open systems* interact freely with their surrounding environment.

Except for the rare closed system, which does not interact with its environment, a system can never be completely separated from its external world. It is for this reason that emergent

Box 3.1 Braess' paradox

Simple cause and effect analysis predicts that putting more resource toward achieving a goal should improve performance, but this is not always the case. The creation of new roads can lead to greater traffic congestion. The installation of new telephone or computer network elements can lead to degraded system performance. Introducing new workers to a team may actually result in a decrease in the team's performance.

To understand these apparently paradoxical results, one needs to examine events from a system view. Studying the effects of new roads upon traffic, Dietrich Braess discovered that if a new road is built in a congested system, everyone's journey unexpectedly lengthens (Bean, 1996). He explained this result by examining the behaviour of drivers, who made individual decisions about their journey, and the emergent effects of all these individual decisions upon the whole system. Consider a journey from A to D that can follow several routes, such as ABD or ACD.

The delay on any link is a function of f , which is the number of cars on that link. As shown in Figure 3.2 (a), there are 2 equidistant path choices. With 6 cars in the system, they will tend to distribute equally, with 3 cars on each path ABD and ACD. If they did not distribute equally, the congestion on one link would over time cause drivers to choose the less congested path. The expected delay is thus 83 on both routes. As shown in Figure 3.2 (b), a new link BC is added, thus creating a new path ABCD. Assuming previous path costs, drivers think that ABCD is now the quickest route. Users take the new path to try to minimize their journey (Glance and Huberman, 1994), but the choice hurts the whole system. Equilibrium eventually occurs when 2 cars choose each of the paths ABD, ACD and ABCD. This puts 4 cars on the link AB. Now, the new road has increased the expected delay for everyone to 92.

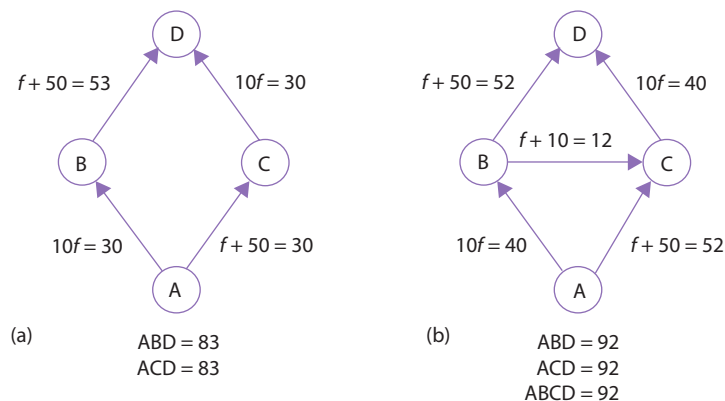


Figure 3.2
(a) and (b) The effect on travel time of building a new road in an already congested system.

behaviour of a physical system is even more difficult to predict because that behaviour is in part determined by the environment in which the system is embedded, and we usually do not know the exact state of the environment. For practical purposes, we usually restrict the description of a system to include a limited set of elements of the environment that are of immediate interest (Box 3.2). Thus, one cannot divorce a physical system from the environment within which it exists because, by doing so, the very context for the system's existence disappears, and its function will alter.

Systems have internal structure

Systems have component parts that together constitute their internal structure. These components are often decomposed into sub-systems for ease of understanding, manufacture or maintenance. A health system is decomposed into community and hospital services. A hospital is decomposed into different departments or services. The amount of internal detail that can be described within a system is probably limitless. The description of the cardiovascular system could descend to the cellular level, the molecular level and beyond. The amount of detail in a system description is usually determined by the purpose of the description.

For a system to function, each of its sub-systems or components must interact or communicate. This means that they may share inputs and outputs, with the output of one component providing the input to another. Inputs to a factory may be raw materials and energy, and its outputs could be manufactured products such as cars or foodstuffs. These products in turn are inputs into the retail sector, which uses them to generate a financial output.

The introduction of antibiotics in the second half of the twentieth century heralded a period of optimism. Infectious diseases were soon to be a thing of the past, and every year saw the discovery of new drugs that attacked an ever-wider spectrum of bacteria. Then, as time went on, the tide started to turn, as individual organisms such as penicillinase-producing bacteria developed resistance to specific drugs. Through a process of natural selection, the drugs that were created to kill bacteria were actually selecting individual organisms that were immune to their effects, thus allowing these organisms to survive and dominate the gene pool. Some organisms now have such widespread resistance that their detection can shut down large sections of a hospital.

Modern medical science is often challenged by critics who cite examples such as the development of antibiotic resistance as proof that its methods ultimately do more harm than good. This 'fight-back' of natural systems following the introduction of technology is not a phenomenon confined to healthcare, but potentially affects every technological intervention made by humans (Tenner, 1996).

From a systems viewpoint, whenever a technology is introduced, the fight-back effect is not so much a fault of the technology as it is an inevitable consequence of the way that we understand systems. When we model the world, we intentionally simplify or exclude whole sections of reality to create a point of view. When a technology is applied, it is aimed at solving a particular problem within a system, with the assumption that its effects are predictable, 'everything else being equal'. This clearly is never completely possible.

The assumption that everything that can be known is known is called the *closed-world assumption* in logic (Genesereth and Nilsson, 1987). Its function is to allow reasoning to proceed even if our knowledge is incomplete. It serves a similar role when a technology is introduced into a system because without this assumption we would never be sure we understood all the possible consequences of the technology's actions.

So rather than being a specific consequence of technology, unexpected outcomes are a result of the imperfect way in which we understand the world. Our only way around it is to assume, at some stage, that we know enough to try things out. The alternative is to do nothing.

Box 3.2 Penicillinase and the closed-world assumption

Systems can regulate their output by using feedback as input

A special case of connecting output and input between systems occurs when some or all of the output of a system is taken back as its own input. This is called *feedback* (Figure 3.3). In this way a system can influence its future behaviour based upon measurement of its past performance.

Feedback systems can become quite complicated in their design and are used to create what are known as *cybernetic* or *control systems*, which are systems that are able to adapt their output to seek a particular goal. Most feedback control systems are more open and use a measurement sub-system to sample their output and use this feedback to determine what the next input should be (Figure 3.4).

The simplest feedback control system consists of three components:

- A *sensor*, which measures the parameter that is to be controlled.
- A *comparator*, which determines whether the measurement deviates from the desired range.
- An *activator*, which creates an output that then alters the environment in some way to change the value of the parameter being measured.

For example, bank account interest is determined by measuring an account value and adding back the interest to the account. A thermostat is part of a feedback system for controlling temperature. The thermostat samples the temperature of a room, which is being warmed by the heating system. As the temperature rises, the system shuts off when a pre-set temperature is reached, or it turns on when the temperature drops below the *set point*. A patient with insulin-dependent diabetes samples his or her blood sugar level with a glucometer and uses the glucose

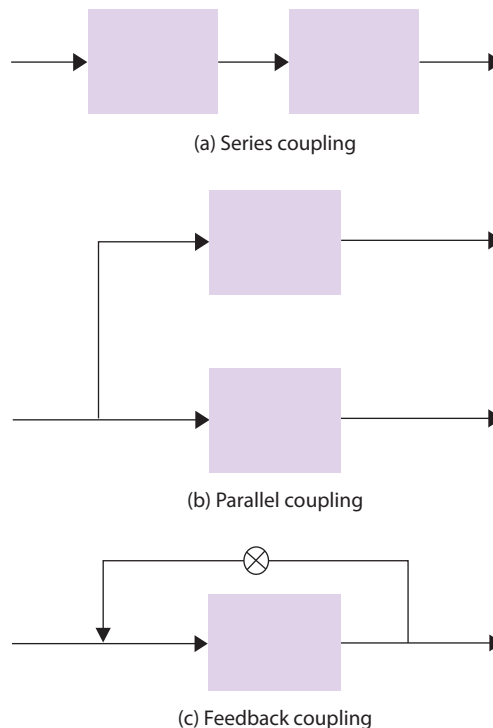


Figure 3.3
(a), (b) and (c)
Sub-systems may be coupled by joining their inputs to their outputs. In the special case of feedback (c), a sub-system takes some of its own output to modify its input.




Figure 3.4
The output of a system may be regulated by a control sub-system that measures the system's outputs and uses that measurement to regulate the subsequent system inputs.

value to adjust the next dose of insulin to be administered, based upon a rule or a formula that aims to keep blood sugar within a set range. The input to the diabetic patient's 'system' is a dose of insulin and food, and the output is an effect on blood sugar that is then used to alter the next input of insulin and food. The thermostat and the glucometer are both measurement devices that form an integral part of a feedback controlled system.

There are two basic types of feedback system (Figure 3.5).

- In a *negative feedback arrangement*, the output of a system is subtracted from the next input. This restricts a system to working within a steady operating range. Physiological systems mostly use negative feedback to provide homeostasis – the maintenance of a desired pre-set physical state, despite variations in the external world. Homeostatic mechanisms maintain everything from body temperature to the concentration of ions within cells.
- In a *positive feedback system*, the output of a system is added to its next input; the result is that the system's output increases with time. For example, a savings account increases in value over time as interest is added to it.



Figure 3.5
System states can be maintained or varied depending upon the type of feedback control being used (After Littlejohn, 1996).

During the French Revolution, a method was devised to create logarithm tables en masse. Individuals called computers each carried out a small calculation, and as a result the team of individuals was able to carry out a large number of intricate calculations. In his seminal treatise On the Economy of Manufacture (Babbage, 1833), Babbage reasoned that because such lengthy calculations could all in principle be broken down into simple steps of addition or subtraction, these calculations could be carried out by machine. This thought inspired Babbage to devise his difference engine, which was the first proposal for a general purpose calculating machine.

Box 3.3 Network systems

- Positive and negative feedback systems can be combined to allow a system to change from one to another of a number of predefined states. The positive feedback component permits the system to move from one state to another, and the negative feedback comes into play when the new state has been reached and keeps the system there.

Systems are arbitrary

We create descriptions of systems to help us understand the observable world. So by their very nature, system descriptions are arbitrary human creations. There can never be something called ‘the correct’ definition of a system, whether it is a description of something in the world or a design to accomplish some function.

There are many possible ways that one could choose to describe a system, and it is often only a matter of convention and practicality that one description is chosen over another. It should come as no surprise that two competing system descriptions may overlap and have common elements. Are the pulmonary arteries more properly part of the cardiovascular system, or are they part of the respiratory system? It depends entirely upon one’s point of view.

Systems are purposive

This brings us to a key point. Descriptions of a system are constructed with a function or purpose in mind (Box 3.3). The intent behind modern descriptions of different physiological systems is to treat illness. The reason that one particular system begins to gain common acceptance over another is that it is inherently more useful for that purpose. Thus, phrenology, the study of bumps on the skull, was replaced by a system of thought we now call neurology because this newer viewpoint proved itself a more useful approach to treating illness. So, over time, whole systems of thought gradually fall into disuse as newer and more useful ones appear.

There are many ways to conceive of systems and to represent their individual elements and the ways in which they interact. Networks are a particularly important and universal systems construct. We can, for example, describe networks of biological and social processes (see Figure 3.6). As templates to action, and as real world artefacts, we invent transportation, communication and energy networks, to name a few.

The basic elements of any network are its *nodes*, which represent the atomic entities that inhabit the network, and the *edges* between those nodes, which record their possible interactions. In a *social network*, nodes are people and edges are their relationships (see Chapter 20). In a *metabolic network*, the nodes are molecules and the edges are their reactions. With a *computer network* such as the Internet, nodes are individual computer machines, and edges are the communication links between them.

Network theory is the study of network behaviour at a general level, and it tells us that despite the wide variety of networks we see in the world, they all share commonalities. Network behaviour is directly related to network structure. The number, distribution and type of links among the nodes in a network strongly influence the network’s properties.

For example, the *average path length* in a network measures the average number of steps it takes to travel from one node to another. Typically, the shorter the average path length, the easier it is for a network to function, whether transporting packets of information across a computer network or minimizing power

loss as energy is transported across a power grid. Many networks share a *small world* property because of commonalities in their structure, meaning that they have a small average path length. The small world effect famously explains why it is usually easy to connect two strangers to each other through shared acquaintances.

Other important network properties that help explain their behaviour include *node centrality* (a measure of how influential a node is in a network, based on how many connections it has to other nodes) and *connectedness* (which describes the degree to which nodes are directly or indirectly connected). The presence of highly connected and central *hubs*, for example, appears to confer stability on networks.

Defining the borders of a network is analogous to finding the boundaries of any system. Some networks are highly isolated. Some are coupled to other networks, upon which they depend, through shared nodes. Dependencies are both a source of system richness and a potential weakness. For example, cascading failures can propagate among networks, as sometimes occurs when blackouts move across connected power grids.

Further reading

Newman M., A-L. Barabási and D. Watts (2006). *The Structure and Dynamics of Networks*. Princeton, NJ, Princeton University Press.

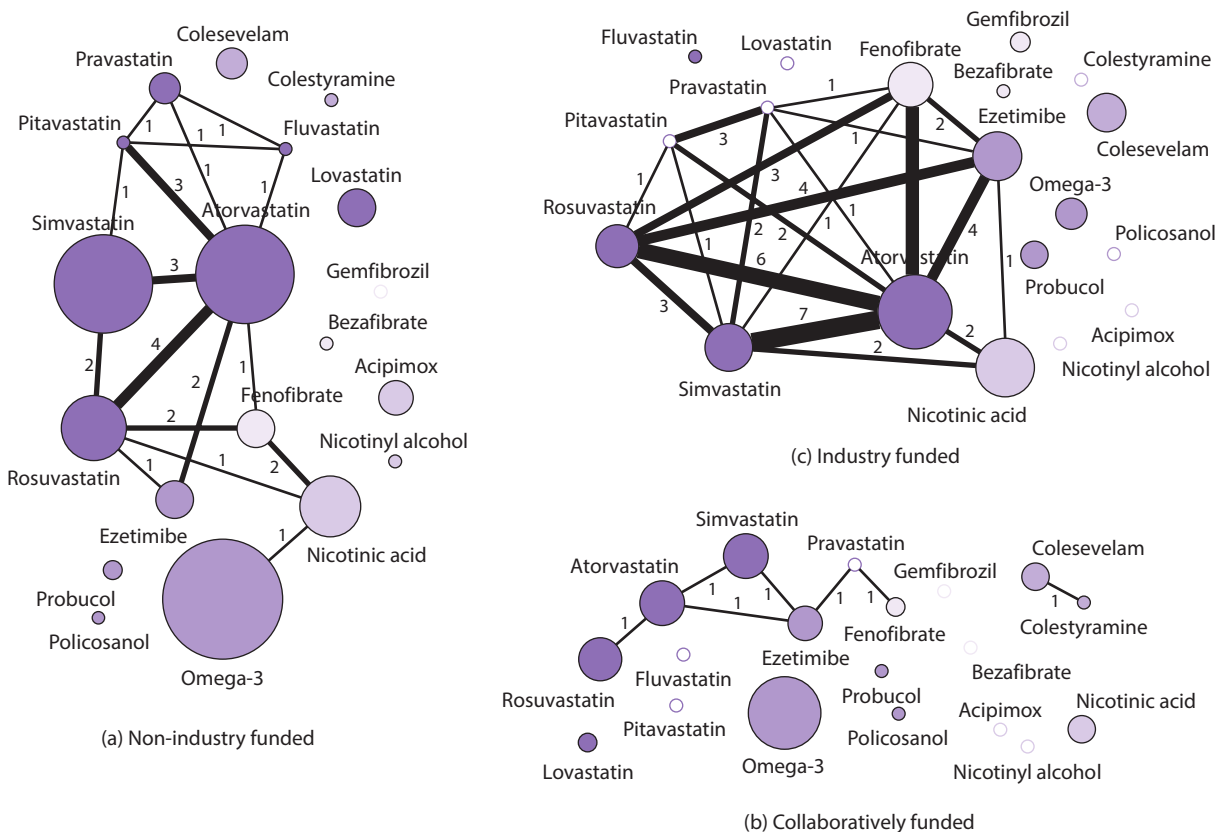


Figure 3.6 (a), (b), (c) Networks showing the number of clinical trials comparing different cholesterol-lowering drugs, depending on whether the comparative effectiveness studies were funded by industry or not. Node size corresponds to the number of trials of a drug. Network structure differences show the different priorities of industry and academic scientists (Dunn et al., 2012).

3.3 Information systems contain data and models

People develop systematic routines because they find themselves doing the same task again and again. We abstract the elements of our actions that recur and give them some objective existence by describing the routine. When the recurring routine is a decision, it should require access to the same kind of data and use the same knowledge. In these circumstances, one can develop a regular process or *information system* to accomplish the decision task. An information system could thus be anything from the routine way in which a clinician records patients' details in a pocket notebook or the way a triage nurse assesses patients on arrival in an emergency department through to a complicated computer system that regulates payments for healthcare services.

An information system is distinguished from other systems by its components, which include data and models. Recall from Chapter 2 that there are different kinds of information models, including databases and knowledge bases. These different information components can be assembled to create an information system. For example, consider a calculator that can store data and equations in its memory. The data store is the calculator's database, and the equation store is its knowledge base. The input to the calculator becomes the equation to be solved, as well as the values of data to plug into the equation. The database communicates with the knowledge base by using a simple *communication channel* within the device, and the output of the system is the value for the solved equation (Figure 3.7).

Many potential internal components could be included within an information system including a database, a knowledge base, an ontology and decision procedures or rules of inference. The different components of an information system are connected by input-output channels, which allow data to be shifted among the components as needed.

A patients' record system is a more complicated example of an information system. Its purpose is to record data about particular patients in some formalized fashion to assist in patients' management. The record system is composed of a database, organized according to data models that are based upon the way clinicians use the data in their decision-making process. A patients' record system also requires an ontology that stipulates the medical language or terminology to be used in the records. The record system may also include access control rules that govern the way that records are to be accessed and possibly the individuals who are permitted to see the records. The record system itself is a component of a larger system, which may be the particular institution within which it exists or the health system as a whole.

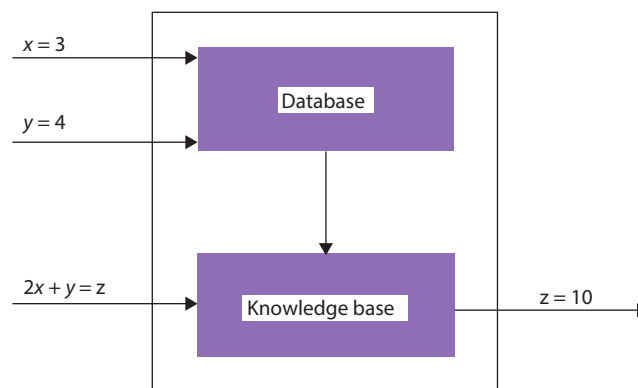


Figure 3.7

A calculator with a stored equation is an information system. The inputs are the equation to be solved and the data values, and the answer output by the system depends on the input values.

We saw at the beginning of this chapter that there are actually three different meanings of the word ‘system’ – a system can be an abstracted description of the real world, a template to action or an artefact constructed in the real world. Consequently, an ‘information system’ may be one of three things:

- A simplified description of an existing set of information processes. For example, one could produce a document or diagram cataloguing the different input and output flows of information that connect a number of different organizations.
- A plan for implementing a new set of information processes.
- An actual physical system. Medline is an information system organized to allow access to the biomedical literature. The Internet is an even larger information system consisting of computer hardware, software, data such as text and graphics stored as Web pages, network connections and the people using these technologies.

Information systems can be created for a number of reasons. In the main, a system is devised because an information process is very common, very complicated or in some way critical. In the first case, the goal of introducing an information system is to reduce the effort of decision making by streamlining the process. In the case of complex or critical decisions, the role of the information system is either to reduce complexity or to minimize the likelihood of error.

Information systems share all the characteristics of systems described in earlier chapters. We collect together decisions, data and models that are of interest for some purpose, and then we look upon them as a system. An information system exists within an environment and is built to interact with an environment. Information systems can never fully capture the richness of the real world, and as with all model-based constructs, they become increasingly less accurate over time as the world changes around them. Information systems built for one set of environmental assumptions may not work as well when transplanted to another environment.

Conclusions

We have now reached a point where we can look at an information system in a fairly rich way, based upon an understanding of the basic principles of models, information structures and systems. Through a variety of examples, we have seen how information systems are often designed to be part of a feedback control system used to manage clinical activities. Whenever a decision is important enough, or is made often enough, an information system is built to manage the process. With this background, it is now possible to move on and look at how information is used to support clinical activities. In Chapters 4 through 8, we will learn how to search for, structure and use information in the support of clinical processes, and in Chapters 9 through 14, we see how healthcare is structured from an informational viewpoint and how information systems reflect and contribute to that structure.

Discussion points

1. Why may a patient object to being considered an input to a ‘health system’?
2. What are the inputs and outputs of a patients’ record system?

3. How many states can you identify for the cardiovascular system, and how does each state change the way the system handles its inputs and outputs?
4. What is the purpose of the health system? What are the components of the health system? What is the purpose of each component?
5. Compare the structure of the health system in the United States and the United Kingdom. If the purpose of both systems is the same, how do you explain the differences in their components?
6. How many people separate you and Kevin Bacon? Why is it such a short list?
7. When is a model not a system? When is a system not a model?*
8. Identify one positive and one negative feedback system in the health system. What are the inputs and outputs? Which components act as sensors, comparators and activators? What is the purpose of the system in terms of control?
9. Identify the different information models used to create the Medline biomedical literature system.
10. Why is it often necessary to upgrade the software in a computer?

Chapter summary

1. A system is a collection of component concepts, processes or objects.
2. Systems transform inputs into outputs and may change their state in doing so.
3. A system has behaviour that cannot usually be predicted by an examination of its individual components but that emerges from the way the components interact with each other.
4. Physical systems are embedded in an environment; closed systems have no external inputs and outputs; open systems interact freely with their surrounding environment.
5. Systems have internal structure.
6. Networks are used to describe the internal structure of many systems. The properties of a network arise out of the way in which the network connects its nodes and edges.
7. Systems can regulate their output by using feedback as input; in a negative feedback arrangement, the output of a system is subtracted from the next input; in a positive feedback system, the output of a system is added to its next input.
8. A feedback control system consists of the following: a sensor, which measures the parameter that is to be controlled; a comparator, which determines whether the measurement deviates from the desired range; and an activator, which creates an output to change the value of the parameter being measured.
9. Systems are arbitrary and purposive.
10. Information systems contain data and models, which include databases and knowledge bases that interact via a communication channel.

* If a model has no decomposable parts, no inputs and outputs, then it is not a system. If a system is an artefact or a process in the real world, then it is not a model. However, if that system has been constructed, then it is probably an instantiation of a model.

PART 2

Informatics skills

Communicating

“The chart is not the patient.”

Gall, 1986

Every clinical action – every treatment choice and investigation made – is shaped by the available information. We can think of this information as the clinical evidence that informs a judgement about the right course of action. Clinicians gather evidence through communication with others, either through what is said now or what has been documented from before.

There are many different sources of clinical evidence used in the routine care of a patient, and these include:

- The patients themselves, who give information about their symptoms and their problems, as well as demonstrate clinical signs through physical examination.
- Clinical colleagues, who exchange messages containing information about the state of patients, their opinions, own workload and needs, or background clinical knowledge.
- The scientific clinical literature, which captures past knowledge about disease and treatment.
- The patient's record, which is a history of the patient's past state, including clinical observations and laboratory and imaging reports, as well any treatments given and their impact on the disease.
- Measurement and imaging devices, from simple instruments such as a blood pressure cuff or glucometer, through to cardiograms, polymerase chain reaction gene tests, ultrasound probes, multisensor patient monitors in intensive care and positron emission tomography scanners.

The information contained in these clinical ‘messages’ is stored in a variety of media and formats and can be delivered in a variety of ways, including face-to-face conversations, letters, e-mail, voicemail and electronic or paper medical records.

When this exchange of information works well, clinical care is solidly based upon the best evidence. When information exchange is poor, the quality of clinical care can suffer enormously. Poor presentation of clinical data can lead to poorly informed clinical practice, inappropriate repeat investigation, unnecessary referrals and waste clinical time and resources (Wyatt and Wright, 1988). For example, the single most common cause of adverse clinical events is medication error, which accounts for about 19 per cent of all adverse events; the

most common prescription errors can be redressed by the provision of better information about medications or the patients receiving them (Bates *et al.*, 2001).

In this chapter we examine this communication process and explore how variations in the structure of clinical messages affect the way in which they are interpreted and can influence the quality of care. If the motto for Chapter 1 was ‘a map is not the territory’, then the motto for this chapter is ‘the chart is not the patient’.

4.1 The structure of a message can influence how it will be understood

What a message is meant to say when it is created and what the receiver of a message understands may not be the same. What humans understand is profoundly shaped by the way data are presented and by the way we individually interpret what is presented. It is thus as

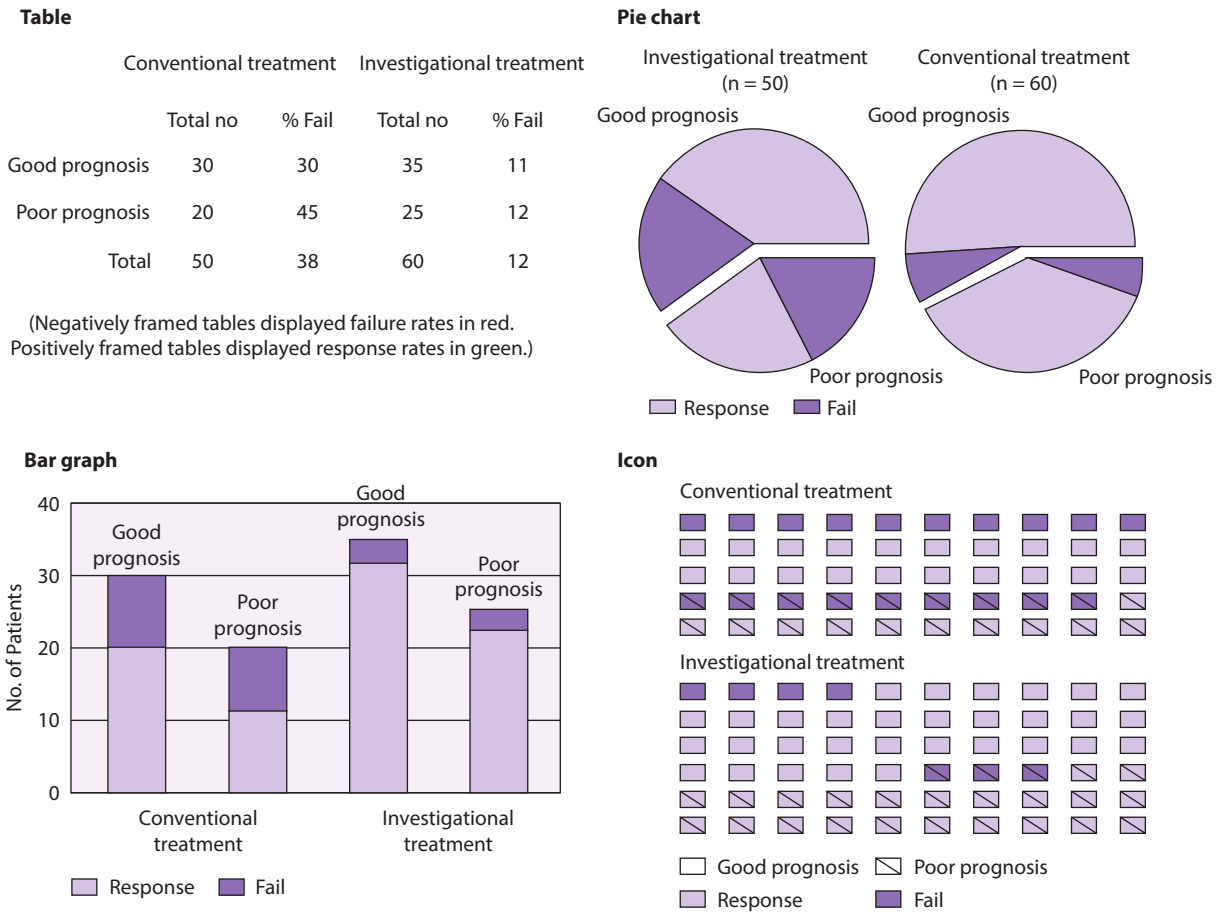


Figure 4.1 Table, pie chart, icon and bar graph displays of the same data from a hypothetical clinical trial each resulted in a different percentage of correct decisions being made. The icon display (bottom right) was most effective for the decision to stop the clinical trial. (From Elting *et al.*, 1999.)


important to structure data in a way that maximizes the chance that it will be understood as it is to ensure that the data are correct.

What a clinician actually understands after seeing the data in a patient's record and what the data actually show can be very different. In Figure 4.1, identical patient data are presented in four different ways (Elting *et al.*, 1999). The data show preliminary results from two hypothetical clinical trials of a generic 'conventional treatment' compared with a generic 'investigational treatment', both for the same condition. In an experiment to see whether clinicians would decide to stop the trial because the data show that one treatment is obviously better than the other, the decision to stop varied depending on how the data were displayed. Correct decisions were significantly more common with icon displays (82 per cent) and tables (68 per cent) than with pie charts or bar graphs (both 56 per cent).

If this example was reflected in actual clinical practice, up to 25 per cent of the patients treated according to data displayed as bar or pie charts would have received inappropriate treatment. This finding underlines just how important the choice of data presentation method is. The way data are structured has a profound effect on the conclusions a clinician will draw from the data. There is an enormous difference between simply communicating a message to a colleague and communicating it effectively.

4.2 The message that is sent may not be the message that is received

Messages can be misunderstood both because of the limitations of the agents interpreting them and because the very process of communication is itself limited. To explore the nature of communication, we will develop a simple general model that describes the process of sending a message between two agents. The agents may be human beings or a human and a computer. A communication act occurs between the two agents A_1 and A_2 when agent A_1 (the sender) constructs a message m_1 for some specific purpose and sends it to agent A_2 (the receiver) across a *communication channel* (Figure 4.2).



The second agent A_2 receives a message m_2 , which can be different from the intended message m_1 . The effectiveness of the communication between the agents (how closely m_1 and m_2 match) depends upon several things – the nature of the communication channel, the state of the individual agents, the knowledge possessed by the agents and the context within which the agents find themselves.

Communication channels distort messages

Many different communication channels are available, from face-to-face conversation, digital channels such as the telephone, e-mail and videoconferencing through to non-interactive channels such as the medical record or letters.

Figure 4.2

When a message is sent between two agents, it is transported over a communication channel. The sent and received messages may not be identical.

The signal to noise ratio measures how much a particular message has been corrupted by noise that has been added to it during transmission across a channel.

A message is sent as a signal across a selected channel (e.g. as sound waves or electronic impulses). Channels vary in their *capacity* to transport such signals. The more limited a channel's capacity, the less of the original message can be transmitted per unit time. Simply put, the thinner the channel 'pipe', the fewer data can flow through at any given moment.

Channels also have different abilities to keep a message exactly as it was sent, and a signal can be distorted during transmission. This distortion is usually called *noise*. Noise can be thought of as any unwanted signal that is added to a transmitted message and that distorts the message for the receiver. It can be anything from the static on a radio to another conversation next to you that makes it difficult to hear your own. Thus, 'one person's signal is another person's noise'. Standard *information theory* describes how the outcome of a communication is determined in part by the capacity and noise characteristics of a channel (Box 4.1).

Therefore, in general, when an agent sends a message, that message may be modified by the chosen communication channel and, through delay or distortion, arrive as a slightly different message for the receiving agent.

Box 4.1 Information theory

Claude Shannon developed the mathematical basis for information theory while working at Bell Laboratories in New Jersey during the 1940s. Motivated by problems in communication engineering, Shannon developed a method to measure the amount of 'information' that could be passed along a communication channel between a source and a destination.

Shannon was concerned with the process of communicating using radio, and for him the transmitter, ionosphere and receiver were all examples of communication *channels*. Such channels had a limited capacity and were noisy. Shannon developed definitions of channel capacity, noise and signal in terms of a precise measure of what he called 'information'.

He began by recognizing that before a message could enter a channel it had to be *encoded* in some way by a transmitter. For example, a piece of music needs to be transformed through a microphone into electronic signals before it can be transmitted. Equally, a signal would then need to be decoded at the destination by a receiver before it could be reconstructed into the original signal. A high-fidelity speaker thus needs to decode an electronic signal before it can be converted back into sound.

Shannon was principally interested in studying the problems of maximizing the reliability of transmission of a signal and minimizing the cost of that transmission. Encoding a signal was the mechanism for reducing the cost of transmission through *signal compression*, as well as combating corruption of the signal through *channel noise*.

The rules governing the operation of an encoder and a decoder constitute a *code*. The code described by Shannon corresponds to a model and its language. A code achieves reliable transmission if the source message is reproduced at the destination within prescribed limits. After Shannon, the problem for a communication engineer was to find an encoding scheme that made the best use of a channel while minimizing transmission noise.

With human verbal communication, the information source is the sender's brain and the transmitter is the vocal cords. Air provides the communication channel, and it may distort any message sent because of extraneous noise or because the message is dampened or *attenuated* the further the distance grows between the communicating parties. The receiver in this model is the listener's ear, and the destination that decodes what has been received is the listener's brain.

Although Shannon saw his theory helping us understand human communication, it remains an essentially statistical analysis over populations of messages and says little about individual acts of communication. Specifically, information theory is silent on the notion of the meaning of a message because it does not explicitly deal with the way a knowledge base is used to interpret the data in a message (Figure 4.3).

Further reading

van der Lubbe, J. C. A. (1997). *Information Theory*. Cambridge: Cambridge University Press.

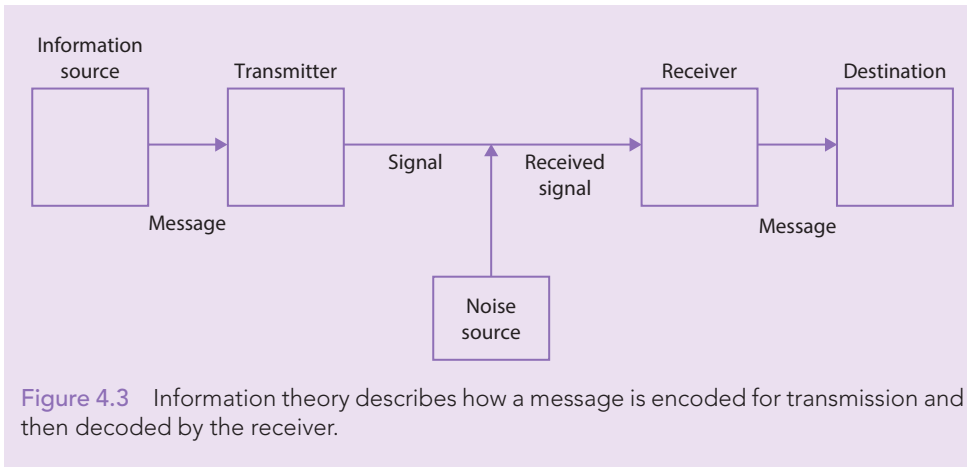


Figure 4.3 Information theory describes how a message is encoded for transmission and then decoded by the receiver.

Individuals do not know the same things

In Chapter 2 we saw that the inferences drawn from data depend on the knowledge base used to make that inference. Because each individual ‘knows’ slightly different things, he or she may draw different inferences from the same data. Variations in diagnosis and treatment decisions, based upon the same data, may simply reflect differences in clinical knowledge among individual clinicians.

When sending a message, we have to make assumptions about what the receiver already knows and shape the message accordingly. There is no point in explaining what the receiver already knows, but it is equally important not to miss important new details. Thus, notionally identical messages sent to a clinical colleague or to a patient end up being very different because we often assume that the colleague requires less explanation than the patient.

The knowledge shared among individuals is sometimes called their *common ground* (Coiera, 2000). This explains why we communicate more easily with others who have similar experiences, beliefs and knowledge. It takes greater effort to explain something in a conversation to those with whom we share less common ground. Conversely, individuals who are particularly close can communicate complicated ideas in terse shorthand. One of the reasons agents create common ground is to optimize their interactions. By developing common ground, less needs to be said in any given message, thus making the interaction less costly and more effective (see Box 21.3).

Returning to our simple communication model, each agent now possesses knowledge about the world in the form of a set of internal models K . Critically, the private world models of the two communicating agents in our model, K_1 and K_2 , are not identical. Thus, agent A_1 creates a message m_1 , based upon its knowledge of the world K_1 (Figure 4.4). A_2 receives a slightly different message m_2 because of channel effects and then generates its own private interpretation of the message’s meaning based upon its knowledge K_2 . Further, agent A_1 makes a guess about the content of K_2 and shapes its message to include data or knowledge it believes agent A_2 will need to make sense of the message being sent. The effectiveness of the message depends upon the quality of the guess agents make about what the receiving agent