



UNIVERSITY OF LEEDS

This is a repository copy of *The Language Machine*.

White Rose Research Online URL for this paper:
<http://eprints.whiterose.ac.uk/81779/>

Version: Published Version

Book:

Atwell, ES (1999) *The Language Machine*. The British Council . ISBN 0-86355-438-5

Reuse


Unless indicated otherwise, fulltext items are protected by copyright with all rights reserved. The copyright exception in section 29 of the Copyright, Designs and Patents Act 1988 allows the making of a single copy solely for the purpose of non-commercial research or private study within the limits of fair dealing. The publisher or other rights-holder may allow further reproduction and re-use of this version - refer to the White Rose Research Online record for this item. Where records identify the publisher as the copyright holder, users can verify any specific terms of use on the publisher's website.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



eprints@whiterose.ac.uk
<https://eprints.whiterose.ac.uk/>



Eric Atwell

**By opening
this document,
you are confirming
that you accept
the terms of the
licence agreement**

[VIEW LICENCE](#)

[QUIT DOCUMENT](#)

For centuries, scholars and scientists have dreamed of a mechanical device with the power of language. That dream is now a reality. The language machine that will listen, understand, translate and speak is rapidly becoming a part of everyday life. This book, commissioned by The British Council from Eric Atwell at the University of Leeds, explores some of the technological, social and educational implications of language machines in the years to come. Will we need to learn languages in the 21st century?

Copyright © The British Council 1999

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, or otherwise, without the prior permission of the copyright holder.

Published by The British Council
<http://www.britishcouncil.org>

10 Spring Gardens
London SW1A 2BN
United Kingdom

Digital book production by Catchline
<http://www.english.co.uk>

The Catchline team:

Editorial director: Margaret Keeton
Publishing manager: Sjanoché Wakely
Design consultant: Carlton Larode

Original cover design by Carlton Larode


Recommended citation reference: Eric Atwell (1999) *The Language Machine*. London: British Council

ISBN 0 86355 439 3

(This digital book is also available in conventional print format ISBN 0 86355 438 5)

Contents

The language machine	6
Research labs to desk-tops	14
Turning theory into practice	18
The problems with machines	28
UK advantage?	38
Real-life situations	44
Future scenarios	50
Glossary	58
References	60
Questions for debate	64



What are the benefits to the UK of making the language machine work? See page 43



This season's must-have: the wristwatch translator
See page 53



Language machines helping with disabilities
See page 54

I propose to consider the question 'Can machines think?' This should begin with definitions of the meaning of the terms 'machine' and 'think'.

– *Alan Turing,*
Computing Machinery and Intelligence (1950)

Imagine picking up the telephone to speak to someone in another country. You have no common language but you are still able to communicate, each of you speaking and hearing your own language.

This is the power of language engineering.

– *Linglink Project Report to the EU (1997)*

Preface

How will computers be used for language learning in 2020? Will we still be learning languages at all? This book goes some way towards answering those questions.

Potentially, computers let you choose your own speed of learning, your own method, and your own vocabulary to reflect your own interests. The concept of 'the classroom', where learning takes place in groups for economic reasons, is itself under threat. The success of individual computer-based learning in subjects such as maths and physics has already proved that.

But isn't language a social skill? Language, yes. Language learning, not necessarily. Teachers may soon be able to concentrate on productive tasks, while the donkey work of practice takes place between individuals and their computers.

So language *use* will never be replaced by machines? We can't be sure of that either. This book shows how close we are to Douglas Adams' babelfish becoming reality. There is no point in suggesting that diplomats will never negotiate via a translation machine, or that politicians will never rouse a crowd of people with translation machines stuck to their ears. If the technology is available, people will use it.

The technology of the language machine may have begun crudely, perhaps even comically. But if one day, I find myself speaking Welsh while my wife answers in Slovak, who will need a world language?

Watch this space: this is how the motor car began.

Rhodri Jones
Manager, English 2000
The British Council

Wealth creating?
The business
opportunities
See page 19

Progress and Prospects

'The success of language engineering research and technological development is bound to have an impact on our economic future because it can be applied across such a wide range of information systems and services with such significant benefits.

... The information society will permeate virtually every area of life involving interactions between people and organisations, in both the public and private spheres. [Language Technology] will enable the information society through intuitive, human-centred modes of interaction with products and services. These will include spoken interaction, removing the need for keyboards and keypads, the use of many different languages to process information and interact with devices – as well as the ability to communicate across language barriers.'

Language Engineering:
Progress and Prospects '98,
The EU Commission (1998)

The language

machine

Speaking to machines? Talking to books? Conversing with a computer?

Is the language machine – listening, thinking, talking back, translating – a fantasy? Or reality?

Surveying the territory

Language machines – machines that in some sense simulate human language and human language behaviour – are used daily around us. Indeed, such machines already have a long history, and have acquired a place in our culture and our expectations.

This book provides a survey of the current state of speech and language technology, focusing particularly on machine translation and speech recognition, highlighting the histories and academic disciplines contributing to their development; it examines the components and technologies; possible pitfalls; main developers; current and potential uses; predicted developments; and paints some likely scenarios for the future impact of the language machine. Finally, this book focuses on current and prospective developments that could have an impact on demand in the marketplace for language teachers and language-based professions in the foreseeable future.

It may help to define what constitutes 'the foreseeable future'. *The Future of English?* (Graddol, 1997) contains projections for as far ahead as 2050; the BT Technology Calendar (page 12) includes predictions to 2045. The developments alluded to in this book are based on an arbitrary forward look of approximately 20 years.

Precise predictions are of course, dangerous: nobody knows for sure what the rate of progress will be in language machine development. Yet what we can be sure about is that personal computing technology is developing – at an almost alarming rate. We need only to look around to see that many applications which once seemed impossible or fantastical, are possible today.

The impact?

The question, 'What will it all mean?' is a fundamental one, and is addressed at many points throughout this book. Here, it may be easier or preferable to explain by example. The following passage from the *Linglink Project Report to the EU* (1997) illustrates some prospective uses and markets for the language machine, which may themselves

form a starting point for a consideration of possible implications:*

* **Think** how much more easily you could use your computer, by simply telling it what to do. Suppose also that many of the features of your car were voice operated, such as using the telephone, selecting a CD, opening the windows, and learning about traffic conditions ahead.

Consider how much time you would save if a computer could find exactly the information you are looking for, from multilingual sources, and do it much more efficiently than you.

Imagine picking up the telephone to speak to someone in another country. You have no common language but you are still able to communicate, each of you speaking and hearing your own language.

Access to systems and services through natural language will make them available to everyone. The barriers will be gone. Machines will understand what we tell them and we shall understand each other better.

Through language engineering we can find ways of living comfortably with technology. Our knowledge of language can be used to develop systems that recognise speech and writing, understand text well enough to select information, translate between different languages, and generate speech as well as the printed word.

By applying such technologies we have the ability to extend the current limits of our use of language. Language enabled products will become an essential and integral part of everyday life.

This is the power of language engineering; using the power of language.

Linglink Project Report to the EU (1997)

What's in a name?

This book is a survey of language machines – what they are; where they've come from; the present-day research; implications for the language industries; and the impressive range of UK and EU collaboration on speech and language technology.

This book also explores another theme – how language machines may evolve from their present use in high-cost, high-specification systems, to practical, everyday contexts in education, industry, business and the home.

Short forms

AI Artificial Intelligence
 CALL Computer Assisted Language Learning
 HLT Human Languages Technology
 IT Information Technology
 MAT Machine Assisted Translation
 MT Machine Translation
 PC Personal Computer
 SLE Speech and Language Engineering
 SR Speech Recognition
 SALT / SLT Speech And Language Technology
 WWW World Wide Web

Glossary?
 See pages
 58 & 59

describing words

the technology

Imaginative ideas
from the Linglink
Report to the EU
See page 7

What makes a language machine?

The term language machine is used throughout this book, but it is really an umbrella term for several developments in computer technology. Such developments will provide us with opportunities to carry out familiar tasks in new ways. For example, we will be able to do the following:

- * talk to machines and have them transcribe accurately what we have said (speech recognition);
- * talk to machines and have them understand the basic content of what we have said, so they can respond appropriately to a question about information that they contain, or to which they can provide access (speech understanding);
- * use machines that translate words automatically from one language to another (whether we input words by speech, handwriting, by scanning text or using a keyboard), or use machines to translate our words to several languages (machine translation);
- * produce a draft of our words in another language for subsequent refinement (machine-assisted translation);
- * read out loud screen-based text such as email, Web pages or word-processed documents (speech synthesis);
- * search through databases of multilingual documents without needing the ability to speak, type or read all of the languages involved;
- * issue commands to 'intelligent devices' – such as a car radio or a mobile phone – using only our voices;
- * speak to a computer and get feedback on the 'correctness' of pronunciation or the appropriateness of lexical choices and grammar style;
- * use computers interactively in the teaching of languages.

At the heart of each of these lies a model of language that understands not simply the *vocabulary* of one or more languages, but also *grammar*, *syntax* and *semantics*.

A language machine is therefore a computer that is capable, in one way or another, of extracting and acting on the meaning of words.

Is it a language machine?

To be called a language machine, a language model must lie at the heart of a computer system, enabling natural language to be understood and processed in an intelligent way. The essential components of a language machine are outlined on page 20.

Perhaps predictably, jargon and acronyms litter the field, making it difficult for newcomers to find their way. To help, a glossary is included on pages 58–59, and for more information, Web and book references are given on pages 60–63.

... So is it SALT?

SALT – Speech And Language Technology – is a term used to describe the growing interdisciplinary field which brings together theory and practical applications of linguistics, computing, engineering, and a host of allied areas. SALT is only one of several rival names and acronyms for the field – HLT (Human Language Technology), and SLE (Speech and Language Engineering) are others – but SALT has been the term favoured by the UK Department of Trade and Industry in promoting industrial applications of the technology.

One of the more difficult areas to define – partly because of the interdisciplinary, exploratory nature of the subject – is the limits of Speech and Language Technology. The question ‘does this technology count as SALT?’ is, at the very least, a starting point.

From childrens’ toys which respond to spoken instructions, to a digitised voice giving a telephone caller push-button choices, it is clear that many current computing technologies have acquired the ability to handle natural language in some way. Indeed, we can see technology and language meeting in an ever-growing number of domains – the Internet facilitates transmission of email messages and publication of text in World Wide Web pages; word processors are ubiquitous in the production of print; computer-assisted language learning (CALL) systems can augment traditional modes of language teaching and learning through exercises in vocabulary, grammar, reading, writing, listening and speaking skills. As these technologies develop and converge, we may even increasingly expect voice to be part of our interaction with machines, perhaps anticipating the usefulness of the language machine in practical contexts such as language teaching and learning.

The same technologies apply not only to English but to other languages. Email messages need not be in English: technology can be used to transmit and receive any sequence of characters – ;-) or 8-[*] – and the World Wide Web allows a wide variety of different data types and formats to be transparently interwoven, so that a Web page can include images, sound, and text.

Yet these technologies are not true SALT technologies: they do not process or model language in an ‘intelligent’ way. Intelligence is required of the person using the system, rather than the system itself. Recent word processors may provide spell-checkers and guidance on grammar and points of style, but incorporate only rudimentary ‘Artificial Intelligence’ models of English. Likewise, most Computer Assisted Language Learning (CALL) systems (particularly those restricted to drills of vocabulary or points of grammar) cannot truthfully be described as having a ‘deep’ knowledge or model of English beyond a list of items to be learnt: the software would work just as effectively if it were equipped with drills of chemical formulae or historic dates and events. (For information about the CALL research which does use language machine technology, see page 49.)

Thus to answer ‘yes’ to the question ‘Is it SALT?’, a computer system should incorporate and make use of some sort of language model.

telling histories

1791 – 2020

Getting to grips
with basics
See pages 8 & 9

1791 Kempelen's speaking machine
1863 Bell's 'talking head'
1872 Edison invents electric typewriter
1876 Bell patents telephone
1887 Sound recording machines invented
1939 Bell Telephones demonstrates Voder
speech synthesis machine
1950 'Turing test' proposed
1970 Successful synthesis of female voice
1976 EC begins development of Machine
Translation technology
1977 Texas Instruments produce Speak &
Spell toy with voice synthesis
1984 Apricot computers in UK build
speech recognition into portable PC
1993 Apple Newton Personal Digital
Assistant uses handwritten input
1994 British National Corpus compiled
1997 Altavista/Digital provide free web
translation service
2020 Maria Callas sings the Sex Pistols?

Dreaming of speaking machines

For centuries, scholars, scientists, magicians and even tricksters have dreamed of mechanical devices with the power of language. The language machines now available – which recognise natural language, process it in some intelligent way, and communicate the result in natural language – are the outcome of a long process of research and development in several related fields.

Machines that talk – what we now know as speech synthesisers – are the oldest of the technologies that make up the language machine. One of the earliest documented accounts relates to the young Isaac Newton, who in the 17th century, observed that filling a glass with beer gave rise to a series of vowel-like sounds from 'oo' to 'ee'.

In the 18th century, several attempts were made to produce mechanical 'talking heads'. In 1791 for example, Wolfgang von Kempelen – an inventor from Bratislava – created a device with bellows which could create human-like sounds. Although the device might now be called a rudimentary speech synthesiser, the machine was not taken as seriously as it might have been. Von Kempelen had earlier attempted to create an intelligent machine that could outplay a human at chess. The machine turned out to be a

From Turing's child-machine to a 25 thousand dollar race

'In 1950, Computing Machinery and Intelligence, a famous paper by British pioneer Alan Turing, postulated a program that could imitate human conversation so well that "an average human interrogator will not have more than 70% chance of making the right identification [as between human and computer] after five minutes of questioning". He predicted this 70% level for around 2000. ...

Each year \$2000 and a bronze medal goes to the most human-like approximation to Turing's definition, thanks to the Loebner Prize competition ... \$25,000 is up for grabs if the program's chance of being rumbled gets below Turing's 70%. This year's Loebner winner scored 89%. ...

Turing's prescription was: stage one, build an educable system; stage two, educate it. (Educable means that professional teachers can take over without further help from implementers.) Turing called instances of stage one 'child machines'. They could be achieved today by an integration of new conversational agent technologies with state-of-the-art machine learning.' (The Guardian, June 24 1999.)

fraud, powered by a chess-playing dwarf.

Some years later the British scientist Charles Wheatstone created an improved model of the Kempelen machine, which was later to capture the imagination of Alexander Graham Bell – the inventor of the telephone. In 1863 Bell produced his own ‘talking head’. The device could articulate the sound ‘mama’ sufficiently clearly to cause a neighbour to enquire after the baby’s welfare.

Perhaps it is unclear whether these experiments and developments belong to the realm of science, technology, music, magic, or public entertainment (Bell’s brother Melville once took to the Edinburgh stage with a false beard and foreign accent to demonstrate one machine), but all these early attempts at speech synthesis arose from attempts to reproduce the articulatory apparatus of a human vocal tract, using bellows and resonating chambers.

It was not until after World War II, when the science of acoustic analysis was more developed, that an alternative approach of synthesising speech sounds directly became possible. And it was not until the 1970s that synthesis of a female voice was successfully achieved.

The state of the art **today**,* which includes ‘speech morphing’, allows the speaking and singing voices of particular individuals to be mimicked by computers. Like the discreet digital changes to photographs that can remove and add objects or body parts to the finished picture, the modern ‘speaking machines’ raise important ethical questions: how much trust should we have in what we hear? Do we have any rights to our own vocal **identity**?*

Future scenarios
See pages 50 to 53

**‘The evolution of human intelligence keeps pace with the evolution not only of language, but of technologies supporting and processing language.’
(O’Donnell, 1998)*

Remember your reactions on learning to use a word processor, and seeing your words appear – and disappear – on a monitor screen? ‘The computer not only made it easier for me to write; it had also changed the very substance of what I was writing, and in that sense, I suspect, it had an enormous effect on my thinking.’ (Johnson, 1997)

**Voices form a part of our identity. How then will language machines change our speech – or thoughts?*

When destiny calls, have the confidence to answer

‘In 1876, Bell and Watson demonstrated their new device to Western Union, but the company’s executives failed to see its potential. “Mr Bell,” they wrote to him, “after careful consideration of your invention, while it is a very interesting novelty, we have come to the conclusion that it has no commercial possibilities,” adding that they saw no future for “an electrical toy”. ... In the next twenty years ... Bell’s telephone company, renamed American Telephone and Telegraph, became the largest corporation in America, with stock worth \$1,000 a share. The Bell patent became the single most valuable patent in history.’ (Bryson, 1994)

seeing futures

emerging

**Artificial Intelligence* researcher Minsky has warned: 'Do not be bullied by authoritative pronouncements about what machines will never do. Such statements are based on pride, not fact.'
(Cochrane, 1997)

**From first imaginings ... technological breakthrough; lowering of manufacturing costs; easy integration into equipment; experiments with consumer electronics; fashion items; adoption by designers and digerati; product copying; rock-bottom cost; high-street retail; global distribution; worldwide acceptance*

...

**'Speech-synthesis* chips which turn text into speech will soon be possible and could, for example, be incorporated into fax machines to read messages aloud to blind and partially-sighted recipients.'
(Philips *Vision of the Future* online magazine, 1997)

The pursuit of the speaking machine
See pages 10 & 11

To 2045: the BT Technology Calendar

British Telecom Laboratories keep a close eye on current and **potential*** future developments in IT*. They have even produced a *Technology Calendar* for 1997–2045 (Pearson, 1997) with predicted technology developments 'derived from literature searches, discussion with experts and logical deduction'. The predictions of greatest relevance to this book are listed opposite. Note that these predictions cover technical developments – computers will get smaller, faster, smarter – alongside social developments – computers will be increasingly integrated into, and essential for, everyday life and work.

The later predictions may seem far-fetched, or at least difficult to have confidence in, although the same might have been true if a typical home or office of today could be described to a reader in the 1950s. Whether or not predictions about thought-input and **direct brain-links*** are plausible, it is noticeable that key developments of language machines come relatively early in the calendar.

A widespread expectation

Another source of predicted developments in IT is Philips' *Vision of the Future* online Web magazine. The Philips Corporate Design team has developed more detailed descriptions of a number of future devices and technologies; again, language machine technologies such as **speech synthesis*** and machine translation play a key part in their vision.

A more academic look at the future of IT is presented in *Computing Tomorrow* (Milner and Wand, 1996), a collection of papers by computer science professors from UK universities, in which a chapter is devoted to natural language processing. In the most recent UK university Research Assessment Exercise (RAE), the journal *Computational Linguistics* was one of the ten journals most cited by computer science researchers.

Together, these are clear indicators of the centrality of the language machine in the UK computer science academic community.

- 1997** *wristwatch telephone
machine recognition of body language and gestures
automatic text summarisation and abstracts*
- 1998** *wrist computer and/or PC Card wallet computer
natural language home information retrieval and
interaction*
- 2000** *electronic newspaper
hand videophone
visual computer personalities on screens*
- 2002** *one-chip, multi-speaker voice recognition*
- 2003** *information technology literacy essential for any
employment*
- 2004** *real time language translation for print and voice*
- 2005** *voice synthesis quality to human standard
full voice-interaction with machine
processing vague information by 'common-sense
inference'*
- 2007** *domestic robots: small, specialised attractive
portable translation device for simple conversation
systems to understand text and drawings (e.g.
patent information)*
- 2012** *robots for almost any job in home or hospital*
- 2018** *Artificial Intelligence technology, imitating the
thinking processes of the brain*
- 2025** *thought recognition as everyday input means
learning superseded by transparent interface to smart
computer*
- 2030** *more robots than people in developed countries
full direct brain link intelligence enhancement by
external means*

* the wildcards?

The BT Technology Calendar organised its dated predictions under a number of headings, but in one category – Wild Cards – it offered a list of ‘things that could happen almost anytime...’

The items listed are eclectic, but illustrate the difficulties involved in predicting what will happen next:

- ? Stock market crash
- ? End of the nation state
- ? Climatic instability
- ? New age attitudes blossom
- ? Human mutation
- ? Humans access net directly
- ? Life expectancy approaches 100
- ? Nanotechnology takes off
- ? No-carbon economy worldwide
- ? Rise of an American dictator
- ? Self-aware machine intelligence
- ? Time travel invented
- ? Major technology accident
- ? Asteroid hits earth
- ? Social breakdown in US or Europe
- ? Whole generation unable to effectively read, write, think or work

Web addresses
for Philips and BT?
See page 63

Research labs

to desk-tops

What stage of development has the language machine reached?

How long before a language machine is bundled with your desk-top PC?

An academic heritage

The history of research within various academic disciplines into speech and language technology is a long one. Within university linguistics departments, for example, computational linguistics is a growing sub-field. Some linguists use computers – in either research or teaching – as powerful tools for analysing or modelling aspects of language. The grammar of English, for example, can be modelled by a set of formal rules; used by a computer program to generate new sentences or parse a corpus of test sentences. This application is comparable to geneticists' use of computer models of gene-sequences or geographers' use of computer maps and databases.

In a computer science department, natural language processing is seen – similarly to Artificial Intelligence, human-computer interaction, and multimedia – as a sub-field of computer science. Speech and language technology is a potentially powerful way to give us 'natural' communication with computers, and to facilitate various types of 'computer-mediated' communication between people.

Language is also important in related fields, for example of telecommunications, or 'telematics' – a term conceived by European Union (EU) research funding agencies – within which language engineering is concerned with building speech and language processing resources, and components for integration into telematics systems.

Where languages meet computer science

Speech and language technology has a number of different names and perspectives. No single academic viewpoint is more 'valid': there are at least as many language researchers in computing as in linguistics departments. Speech and language technology researchers additionally span a wide range of related fields; at the University of Leeds, for example, the Centre for Computer Analysis of Language And Speech (CCALAS) has had members working in fields as diverse as Artificial Intelligence, communication studies, computer-based learning, computer studies, education, English, electrical engineering,

linguistics and phonetics, modern languages, operational research and information systems, philosophy, psychology, and transport studies. Each contributing discipline underpins the language machine in different ways. Briefly highlighted below are the contributions of just two: linguistics and **computer science**.*

Linguistics The language machine needs a theoretical framework within which the components of a language model and their interrelation can be defined, much of which can be derived from the field of **linguistics**.* One example is the division of language processing into a hierarchy of levels – phonetics, lexis, syntax, semantics, pragmatics and discourse modelling (see box, page 16).

Computer science As a language machine is a software system running on a computer, it can draw on many areas:

- * theoretical computer science provides theory on the use of formal simulations and mathematical models of computation;
- * natural-language processing is now a specialist sub-field, providing algorithms such as statistical language modelling techniques for computer processing of language (including analysis of English syntax and semantics) and dialogue management systems;
- * software engineering contributes concepts of modular software design, re-usable language engineering components and resources, and an emphasis on robustness to cope with realistic, unrestricted natural language;
- * Artificial Intelligence also makes an important contribution. A robust language machine must have a large-scale linguistic knowledge base, knowledge of language processing encoded in data, and rules for manipulating and making inferences from that data. It must also be able to adapt or 'learn', improve its own processing on the basis of 'training' or past behaviour. Knowledge-based systems and machine learning are key Artificial Intelligence research topics (Atwell, 1993).

**Computer science* undergraduate courses often include software engineering, theoretical computer science, natural language processing, and Artificial Intelligence: students now encounter language machines as practical problems as well as theoretical ideas. See for example, courses at:

<http://www.leeds.ac.uk/students/ugmodules/comp.htm>

**Language?* Not as we expect it: 'Aaron Williamson's installation and performance *Hearing Things (The Oracle)* connected a microphone in the gallery space with a computer running speech recognition software. The microphone picked up all sounds, speech or not, and the indiscriminating computer was left to do its job. The resulting "creative mishearings" were a testament to the way human attributes resist simulation and systematisation. Williamson, who is profoundly deaf ... added another layer of translation by interpreting the text output with body-heavy movement'. (Haskel, 1999)

Practical uses of language machines?
See pages 44 to 47

creating popular markets

**Developments in speech and language technology are much more likely to be influenced by the spread of personal computing, the consumer market, and take-up of the Internet than by any academic developments in theory, models or algorithms in linguistics or computer science.*

**Widespread commodity product? The language machine may have an academic heritage, but it is popular culture, SF, fashion, and visionary product development that will give the technology wide and trendy appeal for a youth market: 'Instead of carrying your laptop, wear it ... In the future, computer displays may be sold by the gallon and painted on, CD-ROMs may be edible, and parallel processors may be applied like suntan lotion.'*
(Negroponte, 1995)

The contribution
of computing
See page 15

A use for everyone?

New technology typically enters the mass market from origins in research and development – often military or scientific. The language machine is likely to find a similar route: out of the science labs, via the niche markets which have supported and funded its development, combined into third-part products, and made available through resellers and distributors in the wider consumer [markeplace](#).*

Personal computing and the Internet Any widespread take-up of language machines is tied to the broader take-up of related new technologies, and in particular to the rapid developments in personal computing and the Internet – which together are currently driving the rapid 'socialisation' of IT. As these combined forces bring many forms of technology to a wider public, so we can expect language machines to become [widespread](#).*

The language of language processing

Phonetics: the study of speech processes, including the production, perception, and analysis of speech sounds from both an acoustic and a physiological point of view.

Lexis: the study of words or vocabulary items in a language, including all forms having individual meaning or grammatical function.

Syntax: the study of the grammatical arrangement of words and morphemes in the sentences of a language or of languages in general.

Semantics: the study of meaning in language, and the principles that govern the relationship between words and sentences and their meanings.

Pragmatics: the analysis of language in practice, taking account of the context of language use.

Discourse modelling: the analysis of linguistic phenomena that range over more than one utterance or sentence.

Education The practical applications for speech and language technology are many, and several are explored in this book. But what of one of the first uses many people think of for a language machine? Teaching and **learning**.*

Although it is natural to assume that the needs of teaching and learning should contribute to the development of language machines, this has not been the case, and indeed researchers and teachers still struggle to marry their different disciplines. One reason is that the practical tasks involved in teaching languages to humans are fundamentally different from those involved in teaching computers.

Teaching a machine – machine learning – is a sub-field of Artificial Intelligence. A machine learns behaviour through exposure to a ‘training set’ of examples. **Natural-language learning*** systems thus learn how to analyse and process natural-language by being given examples of expected input together with appropriate analyses (Charniak, 1993; Atwell, 1996).

Office and home In some systems, voice has already arrived at the desk-top: Apple Computer’s recent operating systems include English speech-to-text, which can read back screen-based text or ‘alert’ messages in 25 different voices. While some users may struggle to find an application for a talking computer, others – such as the visually impaired – will find this technology beneficial.

What of applications such as translation? To develop a machine-translation system, a researcher could use a corpus of English sentences and their French equivalents to derive an **algorithm*** for mapping from English to French. Machine translation systems are now widely used for gist or ‘first-draft’ translation.

Whatever the application area, language practitioners and linguistic professionals should be involved in helping to plan and integrate the future use of language machines, whether as a substitute for conventional human-to-human interaction, or for the delivery of language teaching.

**English language teachers and learners, according to a report on a survey of user requirements for the ISLE project (ISLE, 1999) desired, above all, a ‘conversation practice machine’ as an aid to learning: with such a system a learner could ‘chat’ in ordinary spoken English to practise everyday conversational language skills.*

**Natural-language learning systems often use sophisticated statistical pattern-matching models (for example Hidden Markov Models or neural networks), oblivious to any possible lessons from education or language teaching theories of human language learning.*

**Algorithms ... Natural-language learning ... Hidden Markov Models ... neural networks?*

See Glossary, pages 58–59

Too expensive
to be useful?
See page 28

Turning theory

into practice

Why do researchers dedicate lifetimes? Why do universities work to provide resources? Why do funding agencies support speech and language projects?

This section looks at what the language machine aims to do, and how researchers are trying to get there.

Why create language machines?

Speech and language technology researchers, like practitioners in many scientific disciplines, can find it difficult to agree on a simple, comprehensive definition of their discipline and its boundaries. It is, however, possible to outline some of the objectives and interests that language machine researchers have in common:

Creating computer models of language One common core aim is to build computer models of human language and speech, such as models of grammar, vocabulary, speech production and recognition, and language learning. Computer models do not have to use graphics to 'visualise' aspects of language: they can be conceptual as well as visual. A set of formalised grammar rules defining the phrase-structure of part of the English language constitutes a 'model'. Such models have applications in several disciplines and research areas.

Building computerised language resources Researchers generally make use of computerised language resources in developing language machines, either directly or indirectly. These resources may include: corpora (collections of texts); lexical databases (computerised dictionaries); and software 'tools' to access such resources and link them to other software systems (page 27 lists available resources). In themselves these resources are a huge investment supporting potential uses in many different disciplines.

A corpus can help not only in direct research: it may also help indirectly. One example would be in guiding the development of a new English dictionary. Once completed the dictionary itself could be subsequently incorporated into a language processing system.

Assisting natural communication between people and computers Language machine developers have a common interest in using language to make communication between people and computers more 'natural'.

With a natural-language spoken dialogue system, for example, people can access information on a computer

using the language of ordinary, everyday **conversation**.*

This has several practical applications. Routine telephone enquiry services or call centres can be augmented or replaced by spoken dialogue systems that access computer databases; the Dutch national railway telephone enquiry service (tel: 00-31-900-9292) already offers this option to enquirers held in a queue (Bod, 1998). With implications for staffing costs and service 'round the clock', it is easy to see why many telecommunications companies are investing in language engineering.

Assisting communication between people A related objective for language machine developers is to assist or enrich direct human-to-human communication by using language machines as intermediaries. **Communication*** between people may be either spoken or based on written text – and language machines can have a part to play in both types of interaction.

Machine translation systems already help people to communicate across language barriers. Yet the potential for this use exists not only across continents: the language machine also has a role to play within multilingual countries. The UK internal market is one such opportunity highlighted by The Engineering and Physical Science Research Council (below).

Wealth creation Speech and language technology has significant commercial potential. In the UK, the Department of Trade and Industry actively promotes its take-up and commercial exploitation by UK **industry**.* It is a view shared by funding bodies. The EPSRC highlights the economics of multilingualism:

Supporting the languages of [the UK's] ethnic minorities in this way could have rich economic rewards ... Thus Urdu, Gujarati, Punjabi and Cantonese become important targets for the domestic translation industry to be able to handle. They are also the languages of some of our major trading partners and the language of trade is the language of the customer. (EPSRC, 1998)

**Today's travelling professionals use a combination of lap-top computer, modem and phone to access large computer systems located elsewhere. Tomorrow, they may reach their computer, to issue commands and requests for information, simply with a mobile phone – and their voice.*

**Communication is a vital part of any workplace – so expect to see language machines built into off-the-shelf word processors and software. Some professions are already using the technology: publishers have employed speech and language techniques to help build language teaching dictionaries and textbooks.*

**Industrial uses? Manufacturers and suppliers of language machines can find a route to market through companies looking to embed speech and language technology in third-party products or services. The UK DTI is a point of information and support.*

Global trading
benefits to the UK
See page 43

what's involved

input to output

Can the language machine really help people to communicate?
See page 19

Defining shared approaches

In 1995, the US National Science Foundation and the European Commission commissioned a comprehensive survey of the state of the art in *Human Language Technology* (Cole et al, 1995), which reviewed the field in terms of the components or modules that are needed for an overall, holistic speech and language processing system.


These components – what makes up a language machine – are broken down into the various steps of the technology: input of words (spoken or written); analysis; understanding; spoken output or text generation; dialogue; document processing; multilingual ability; multimodality; mathematical models; language resources; storage; transmission; and evaluation.

Figure 1 shows how these components interact: mathematical models, language resources, transmission, storage and evaluation modules can be called by any of the other elements.

Figure 1: The essential components of a system defined as SALT – Speech And Language Technology

The components in this diagram are definitive of a SALT application: any true language machine must include one or more of these modules and any software system built on one or more of these components constitutes a language machine.

Language technologists work within each of these areas, and universities and funding agencies seek to support research activities here – these areas may also indicate the type of potential uses for the language machine that we can expect to emerge in the office workplace of the future.



Spoken language input The challenge of spoken language input is to take acoustic signals (the soundwaves of speech) and convert them into characters, words, and sentences that a computer can analyse. There are two approaches to handling spoken input:

- * Speech recognition: this converts input spoken language into the standard computer text representation (Ascii characters) by 'transcribing' as faithfully as possible;
- * Speech understanding: extracts the 'gist' or 'meaning' underlying the input speech, passing these key words and concepts to another application.

Speech input 'add-ons' to word processors, such as IBM ViaVoice or Dragon NaturallySpeaking, use speech recognition to attempt faithful transcriptions of dictated **speech**.*

A speech understanding interface to a computer-based information system, on the other hand, may not need to transcribe every word as long as it can spot the key words and phrases of the query. In the example of a speech input/output Loebner prize contestant (pages 10, 24), for example, spotting the word 'mother' might trigger the response: 'Tell me more about your family.'

Early speech recognition research systems were tuned to the voice of a specific **user**.* Nowadays, speech recognition systems are generally trained to expect standard American or British English, until or unless taught otherwise. Many can adapt to the pronunciation and accent of an individual speaker if that person has the patience to read out training sentences.

While speech recognisers can adapt their original language model, they can have particular difficulties where a user's speech varies significantly from the standard model: pronunciation, heavy accents or varying styles of English as spoken by non-native learners can each affect recognition. Yet speech recognition systems have found educational uses: the ISLE system (see page 62) has a spoken language model based on British English, and aims to detect and diagnose non-native learners' pronunciation

**Speaking to books? The author's original manuscript for this book was prepared using IBM ViaVoice speech recognition software, rather than a traditional keyboard and mouse. The software had difficulty with the phrase 'ELT practitioners', offering instead:*

*eel teeth practitioners
reality practitioners
the empty practitioners
guilty practitioners
ancillary Capriati*

Speech recognition allows authors to dictate documents directly; grammar and style checkers can then be used to polish documents and bring to the author's attention text that might benefit from more of their time.

**IBM's UK speech recognition research team was at one time led by a Yorkshireman, and some of their demonstrations featured this regional accent.*

Problems with the language machine?
See pages 28 to 35

listen, analyse

understand

Medieval style or machine stylus?

'As a calligrapher, I was fascinated by the palm-top computer "assistants" that replaced a keyboard with a rectangular writing surface and stylus.

Despite 20 years experience as a calligrapher, and daily approving comments on the regularity and attractiveness of my handwriting, the successful use of these "assistants" alludes me. My handwriting, according to exasperated sales staff and exhibitors, is 'too good' for the basic recognition techniques used by the machine. My classical italic hand joins too many of the letters for the "assistant" to readily distinguish them. Does a style of handwriting developed to remain legible at greater speed – so that court clerks could keep records of court proceedings in medieval Italy – have to be abandoned to help a machine?' (Friedrich Sporle)

Despite the demise of the Apple Newton which used handwriting recognition, IT suppliers have not abandoned the challenge. BT are developing the SuperQuill: a pen-sized handheld computer that can 'read handwritten words' – even when the holder is writing not against a pad but in the air.

difficulties by highlighting specific phonemes where the learner deviates from the model.

Written or printed language input Another way to produce written text – or, rather, printed text – is to type it into a word processor, email or desktop publishing package. For many people, however, a lack of keyboard skills can make using a computer a daunting, slow or potentially embarrassing task. The text may also exist in typed or printed form, but the 'electronic' copy may not be available or in a usable file format.

To address these scenarios, language machines need techniques for recognising handwritten script and incorporating optical character recognition (OCR). In the case of OCR, text is scanned and the image is then interpreted and converted into an Ascii text file. Most OCR packages can be trained to recognise particular fonts or typefaces to reduce errors in interpretation: applying a spell-checker to the initial output will often further reduce transcription errors.

Handwritten text can present more daunting problems: few people's handwriting is as consistent – or as clearly legible – as a typeface, and the recognition techniques must also take into account joining strokes and other variables.

Language analysis and understanding Whether language is input as speech, by scanning, in writing or typed via a keyboard, the next stage involves linguistic analysis of the character-sequences.

Words must be first identified in a computer dictionary using a 'lemmatiser', which computes the root (or lemma) for any input word. The dictionary may also be used by a part-of-speech tagger (see page 48), which computes the grammatical word-class of each word, and by a semantic tagger, which annotates each input word with a representation of its 'meaning'. A parser then adds grammatical phrase-structure to each input sentence. Generally, the system has to build a representation of the 'semantics' or meaning of each sentence. This representation is passed on to an application; in a database interface, for example, it

could be used as a query to retrieve information.

Linguists and language engineers have not yet been able to agree a standard set of part-of-speech tags, semantic tags, phrase-structure trees, or semantic representations to use in language analysis and annotation. As a result, the EU has funded research and development projects to try to agree standards (EAGLES and MATE: see Web references, page 62): examples of annotation at each level are available in the reports from these projects.

Language generation If the application calls for a reply or response in English, this response needs to be generated. Unless this subsystem can ‘know about’ and use grammar and meaning, it will be limited to simply producing ‘canned text’ or predefined sentences. The peacekeeper’s translator, for example (page 45), uses a speech recogniser for spoken English input, but the output translation is selected from a set of digitally **pre-recorded phrases**.*

Spoken output technologies A speech synthesiser converts Ascii text into spoken language output. This requires a variety of subsystems. A ‘text chunker’, for example, breaks running text into meaningful ‘chunks’ or ‘tone units’, telling the synthesiser where to insert pauses, raise or lower pitch, and so on in order to make the synthesised speech intelligible and easier to listen to.

Current speech synthesisers, such as BT Laureate, IBM ViaVoice Outloud and Home Page Reader, can output simple sentences and mimic a range of different speakers. Despite these advances, however, they can produce ‘unnatural’ intonation and occasional mispronunciations when attempting to deal with unrestricted natural **dialogue**.*

Discourse and dialogue Some language machines attempt to combine language input and output to allow a ‘natural conversation’ or dialogue with the computer. This calls for a dialogue management system that includes all the above components, along with models of overall discourse structure, principles of pragmatics and so on, to ensure

*‘*Computerised voice synthesisers often have a tendency to sound rather, well, robotic. For years, researchers have worked to improve such systems’ intonation so that it matches the “prosody” rules that people use to add grammatical meanings to strings of words – for example, the way that the pitch of a voice falls at the end of a sentence, or rises before a question-mark. But even the most well-spoken computers tend to sound bored, and their endless droning can irritate human listeners.*’ (The Economist, February 27 1999 p. 96)

**Many researchers believe that we’re more likely to be using natural speech input than listening to synthesised natural speech output: ‘One-chip multi-speaker voice recognition’ will be achieved before ‘voice synthesis quality is up to human standard’.* The BT Technology Calendar. (Pearson, 1997)

Language machines in action
See pages 45 – 47

talk back

and translate

**Thinking computers? The annual Loebner Prize, evolved in part from Turing's ideas, offers a prize for the 'most human computer'. Contestants are currently limited to keyboard input and on-screen replies, but the longer-term aim of the contest is to progress to spoken dialogue input and replies – in effect, an English language conversation.*

Taking the
Turing test?
See page 10

**Redundant? MT 'only works well with simplified source texts, or to produce a "gist" version ... of documents that, for economic reasons, would otherwise never be translated at all. There is certainly no prospect in the foreseeable future of human translators being made redundant by such systems!'*

With Machine-Assisted Translation however: 'the computer provides sophisticated aids to enhance the efficiency and consistency of the translation, but the human remains in charge and makes all the decisions.' (Rothwell, 1998, describing the University of Leeds MA in Applied Translation Studies.)

smooth interaction. The ideal spoken dialogue system could meet the Turing Test (Turing 1950), designed to determine whether an AI system was truly 'thinking'. Turing's suggestion was that, if a human engaged a computer in conversation and the responses from the computer were indistinguishable from those of a human, the computer could be said to be **thinking**.*

Document processing Dialogue modelling is less relevant to written language, but systems for processing written texts do need other types of 'overarching' models: intelligent tools for text editing, retrieval, information extraction, summarisation, and so on. As word processors become increasingly sophisticated and incorporate these types of functionality, the boundaries between language machine and word processor are starting to blur. English language learners can clearly benefit from access to 'language-aware' functions in word processors, although they do need to be aware of their possible fallibility.

Multilinguality A machine translation (MT) system translates text from one language into another, and may even deal with spoken input or output. They may also involve language identification: determining which language is being used given a printed or spoken signal. Most research has concentrated on English and, to a lesser extent, other European languages. For languages and language-pairs that have no comprehensive machine-translation system, tools for machine-assisted translation – such as bilingual dictionaries and term-banks – may be useful.

An obvious potential concern for language teachers and language professionals, such as interpreters and translators, is whether MT (and the language machine in a broader sense) could make language-learning – and hence language teachers – **redundant**.*

Current MT systems remain imperfect, and there are still many social, business and political situations in which being able to communicate directly in a shared language is greatly preferable to computer-mediated communication.

Multimodality With rapid advances in multimedia computing, techniques for combining language with image, gesture, and facial movement are increasingly being explored. Many Web-based chat facilities now allow the use of cartoon figures, often referred to as 'avatars'. These raise issues relating to role-play and the creation of online identities as much as language machine technology. Some speech synthesisers, for example IBM ViaVoice Outloud, provide a cartoon face to accompany the synthesised speech, so you can 'see' someone speaking to **you**.* Manufacturers continue to search for likely new multimedia products to create and develop emergent markets. One example is BT's domestic videophone, withdrawn from sale in 1999 due to lack of demand: the new Orange videophone is already planned for release.

**Virtual immortality?*

A multimodality system 'would capture images of your face and examples of your speech before you died, so that after death your bereaved family could see and talk to you via a "virtual video-phone".' (Howard, 1998)

The European dilemma ...

In the 1980s, the EU funded the Eurotra project to develop a machine translation system for translating between languages of member states. Much of the funding went to individual research centres in member states, along with a free hand to develop their theories. The result was a range of different system modules from different sites, based on different models or theories of language. While potentially interesting as a vehicle for comparisons between rival linguistic theories, it was arguably not a good engineering approach to issues such as algorithm and component reusability. Eventually, the EU bought in technology – Systran – developed outside the Eurotra project. For some researchers, Systran bore the stigma of being composed of a bundle of techniques that may have worked in practice but were not theoretically well founded. The EU, however, decided it was better to have an inelegant system that worked than an elegant one that didn't.

... with no clear solution?

The translation demands for the European Commission are huge – the translation department, based in Luxembourg, now employs nearly 2000 staff. Not surprisingly, the centre continues to seek translation methods such as voice activation techniques, to take it into the future, with varying degrees of success:

'The Commission has been developing machine translation techniques since 1976, and it now operates a machine translation procedure called EC/Systran, based on the product developed by Systran. "Any official in any of the European institutions can send his or her text for translation to this system by e-mail," explains Mr. Theologitis. "Turnaround time for any document, virtually regardless of length, is around ten minutes."' Voice activation trials, on the other hand, remain 'inconclusive.' (International Herald Tribune, January 25 1999)

knowledge

communication

Systran in the EU
See box page 25

**Will the language machine affect the way we write?*

'Enabling machine translation from English into other languages – requires planning, preparation, precise agreement on specific word use and allowing for languages that take up more space. Using icons, drawings and diagrams can also help, while jargon, colloquialism, ambiguity and grammatical sophistication can all confound the process.'
Fiona Mitchell, Technical Author

Transmission and storage In the real world, applications have to deal with noisy environments, transmission over networks, and other signal processing difficulties. A specialist topic of research covers issues of speech coding and enhancement to reduce noise and transmission costs.

Mathematical methods Some theoretical techniques apply at several 'levels' in language machines: statistical language models or neural networks can 'learn' by being exposed to a corpus of exemplar language data. Machine learning based on mathematical pattern-matching models is, however, very different from human learning.

Evaluation Another overarching issue, touching all other sub-fields, is evaluation: comparison of rival systems in terms of accuracy, speed, and more intangible qualities such as 'usefulness' and 'user acceptability'.

While quantitative evaluation (of aspects such as accuracy and speed) can be done using a 'test corpus', qualitative assessment calls for people to use the machine. Any valuable evaluation of language machine applications in language teaching, for example, calls for teachers and students. Although speed and accuracy of language-processing modules are important, the key concern must be whether people really feel that they **benefit**.*

Language resources: hundreds, thousands, millions – and growing

Many areas of research and development require large-scale resources, such as language corpora, lexical resources, and software tools and components. A major research issue relates to the collation of reusable, large-scale resources that can serve the needs of many different language machines. One of the first corpora, collected in the 1970s, was LOB – Lancaster-Oslo/Bergen – which contains one million words. Some 20 years on, the BNC – British National Corpus – was compiled, with one hundred million words. These much larger collections are called VLCs – Very Large Corpora.

For the speech and language technology researcher, such a text corpus can be used to 'train' a language machine: a machine-learning algorithm or program can be applied to the corpus for example, to extract a language model automatically. A text corpus may also be used to test a language machine and evaluate the coverage of its language model – for example, by counting how many sentences a grammatical analysis program can parse correctly.

Resources for developing the language machine

Spoken corpus resources

Recordings of words, sentences, scripted speech, spontaneous monologues, dialogues.

Collations of written texts.

Part-of-speech word-tagged text (each word is annotated with a Part of Speech tag).

Treebanks: sentences annotated with syntactic phrase structure.

Spoken corpora: digitised acoustic signal aligned to transcription.

Parallel corpora: English sentences aligned to translations in another language.

Learner corpora: English learners' text or speech with errors (pronunciation, spelling, grammar, etc.) marked, and annotated with corrections.

Lattice corpora: the uncertain output of a speech recogniser, including a set of Ascii 'candidates' for each spoken word.

Parallel-annotated corpora: a single text sample annotated with several of the above levels of linguistic analysis.

Lexical resources

Wordlists for a wide range of languages.

Machine-readable dictionaries (MRDs): headwords, parts-of-speech, pronunciation, inflected forms, and meaning definitions.

Lexical resources tailored for specific systems, e.g. specialist terminology banks.

Bilingual dictionaries for machine translation and computer-assisted translation.

Lists of collocations, idioms, proverbs, compound terms, and so on.

Thesauri showing synonyms and other lexical relations.

Wordnet: a network of semantic relationships between words.

Software tools and demonstrators

Lemmatiser: gives the root or lemma for any input word.

Part-of-speech tagger: gives the grammatical wordclass of each word in running text.

Parser: adds grammatical phrase-structure to each input sentence.

Semantic tagger: annotates each input word with 'meaning' (e.g. from a dictionary).

Concordancer: examples of selected keywords, as they appear in context in a corpus.

Machine translator: translates text from one language into another.

Speech recogniser: converts input spoken language into Ascii text.

Speech understander: extracts the 'meaning' underlying the input speech.

Speech synthesiser: converts Ascii text into spoken language output.

Text chunker: breaks running text into meaningful 'chunks', for example, to make synthesised speech more intelligible.

Pocket electronic dictionaries and translators.

The problems

with machines

If we have the common aims and approaches in place, why isn't the language machine in common use?

What problems are slowing development? What holds it all back?

What slows development?

There are several problems facing the researcher – and indeed the potential customer – for the language machine, some of which may continue to hold back the development of the technology well into the future. This section examines seven key problems.

Problem 1

The language machine is expensive

Speech and language technology has, in the past, had a reputation for requiring computers with powerful processors and large amounts of memory and disk storage space, allied to in-house research and development teams who develop and install language machines and integrate them into current systems.

This situation is now changing, thanks to rapid increases in computing technology. Much like advanced graphics and Internet access, technologies previously restricted to university research laboratories are increasingly available on home and office computers.

One concrete example of this is speech recognition technology. IBM, like many others, have used the basic algorithm used in speech recognition – known as Hidden Markov Modelling – since the 1970s. In the 1980s, they developed impressive speech recognition demonstrators in their research laboratories, although showing them at conferences and exhibitions required advanced computing hardware and speech research experts.

The same demonstration (or perhaps a more convincing one) now can be taken into small business premises, colleges, schools or homes – without the need for a researcher or a lorry full of hardware. You can buy a speech recognition package off the shelf, or may even get one free as a promotional item with a PC magazine.

This move, from research laboratories to **everyday use,*** leaves some academic researchers feeling redundant. Speech recognition is no longer a technology research issue, as the theoretical problems are (supposedly) largely

solved. Further evidence of this trend is provided by the British EPSRC* responsible for providing funds for fundamental, theoretical research in UK universities. A recent policy discussion document on the future of research on Human Factors in Information Technology (EPSRC, 1998) proposes that university researchers can no longer make a useful contribution to speech technology theory, and that the EPSRC therefore should stop funding research in this field.

But even though some speech and language applications have reached the level of availability and price of everyday PC software, this is still beyond the reach of many language professionals.

Take, for example, English language teachers and learners in developing countries,* where class sizes of 50 to 100 students are common, and a blackboard with chalk is often a luxury. Even in Western schools, where computers are more commonly available, speech and language technology remains difficult to use in whole classroom teaching. Language machines are usually designed for an individual: a speech recogniser, for example, is limited to a single headset-microphone and using it in classroom situations is difficult beyond occasional demonstrations.

Manufacturers of CALL software that incorporates speech recognition – for example, Talk To Me (Auralog 1996), ISLE Interactive Spoken Language Education (ISLE, 1999) – recognise these difficulties and deliberately target individual language learners with home PCs rather than classroom groups. This development, alongside moves to greater learner autonomy and the increase in open learning approaches, may in time offer new ways of constructing the learning experience, while fundamentally changing the balance between classroom and individual learning. Inequality of home PC access may however become a pressing issue, and perhaps makes more important the role of study centres and libraries* offering access for all. In providing points of access to new technology, such centres may perhaps in the process reinvent themselves as technology junctions.

*EPSRC – Engineering and Physical Science Research Council – is one of seven agencies funding research in UK universities. Its role is to promote and support strategic and applied research, advance knowledge and technology that contributes to economic competitiveness and life quality, and promote understanding in the fields of engineering and the physical sciences.

*Compare two countries: in 1998, an IBM PC cost the same in London and Bucharest. Average monthly salaries in GBP:
British – £1600
Romanian – £45

*'Information technology is important in tackling social exclusion because anyone can walk in to a library, sit down and start tapping away – at absolutely no cost.' – Chris Smith, the Secretary of State for Culture, Media and Sport. (The Independent, June 19 1999)

Language machines for disabilities
See page 54

inappropriate

technology

How many million words are collected in the BNC?
See page 26

Problem 2

The language machine has been developed without consultation

On the whole, developers have not asked users what they want from speech and language technology. Research has been driven by the interests of researchers rather than potential customers, and has tended to focus on tackling theoretically interesting problems rather than examining usefulness. In the past, for example, much computational linguistics research was led by linguists whose primary aim was to implement linguistic theories.

The EU learnt from its experience in the 1980s with the Eurotra project (box, page 25), and is now funding a programme of language engineering research where users must be consulted at all stages. In the case of the ISLE project, designed to respond to German and Italian learners' pronunciation of English, teachers and learners have been involved at every stage, from initial drafting of research plan through to specification of requirements and verification of the final system. (ISLE project, see page 62)

Will languages teachers be redundant?
See page 24

Problem 3

The language machine doesn't match the expectations of customers

Some potential customers expect the language machine to be able to understand everything and to match human language-processing, even though people themselves often display a considerable talent for misunderstanding!

Although the language machine clearly falls well below human **standards*** for language processing at present, there are a growing number of tasks for which it is arguably adequate. In practice, the best current machine translation systems are at least as good as a fairly competent second language learner. This may be quite acceptable for some tasks: imperfect translation may be adequate for the reader to gather the 'gist' of a text, or for a first draft to be polished and corrected by a professional, human translator.

The problem with calligraphers
See page 22

Problem 4

The language machine is not appropriate for some tasks

Some computational linguists might suggest that speech and language are obviously always the best means of communication – with or via computers. But even if the language machine worked with a high degree of accuracy, it is still not clear that everything should be done via language and speech.

The computer industry has put much effort into developing graphical user interfaces (GUIs, pronounced ‘goo-eez’) that use icons, 2-D and 3-D graphics, menus, and pointers; these may not readily convert into speech-controlled applications, and even if they could be converted it may not be the best approach to the task. Most of us, for example, manage to use lifts quite happily by pressing one of a small number of buttons to get to our selected level in a building: it is doubtful whether speech control could add much value to the experience. The automobile and computing industries are putting large amounts of money into developing in-car command and control via speech for non-driving functions that are not safety-critical, such as opening the windows and turning on the radio (see for example, Schillo et al, 1996; Tyler, 1998). Again, the added value of such language machine gadgetry remains dubious.

**The standards and skills of computer software are already better than many people realize. ‘It is frequently my experience that when demonstrating recent advances in, say, speech or character recognition, observers are surprised at the state of the art. For example, a typical computer user’s last experience with speech-recognition technology may have been a low-end freely bundled piece of software from several years ago that recognized a limited vocabulary, required pauses between words, and did an incorrect job at that. These users are then surprised to see contemporary systems that can recognize fully continuous speech on a 60,000-word vocabulary, with accuracy levels comparable to a human typist. (Kurzweil, 1998)*

I will compare the thee with the day of an été?

If you read ‘This book examines the current state of technology of languaggio and the speeach and observes in particular two technologies: automatic translation and acknowledgement of the word’, what would be your reaction? A word-processing glitch? A liquid lunch for the editors? Actually, this is how a phrase from this book looks after a return trip to Italian through AltaVista’s Babelfish translation service. You may have got the basic meaning – apart perhaps from ‘acknowledgement of the word’. Compare this with the line of Shakespeare in the title of this box to judge how conventional prose stands up to machine translation in comparison with poetry. Yet we can be sure that translation systems will improve: expect upgraded and extended corpora to allow more continuously successful translations with fewer untranslated or mangled words. A panel of experts give their verdict on the current usefulness of machine translation on page 37.

human-machine

mismatched

*‘*Speech recognisers were wholly inappropriate where misinterpretations could imperil life and limb. The safety critical problem is that two quite different instructions can sometimes sound very similar: “descend to a hundred feet” and “descend 200 feet”. ... improving the linguistic model and using contextual constraints may lead to more acceptable results. ... While acoustic pattern-matching has made great advances to the stage where sophisticated continuous speech recognition packages are available ‘off the shelf’, there is still a need for further research into higher-level linguistic models of grammar and dialogue structure for practical enterprises such as ours.*’
(Churcher et al, 1997)

Safety critical computing is a key area where language machines may be not just unwise but dangerous, at least until error-rates are much lower than at present. Consider, for example, the use of **speech recognition*** as an aid to Air Traffic Control around Leeds-Bradford Airport.

Problem 5

To use the language machine successfully, we need to rethink how we approach tasks

If we simply try to add a language machine to an existing system, the fit may not be perfect. Using a speech recogniser is a quite different process from typing into a word processor: people can speak much faster than they can type; words will not be recognised and transcribed correctly, leaving the speaker to review and correct mistakes. The time taken to dictate a book could be longer than the time taken to type it from scratch.

With the advent of word processors in the 1980s, typing pools diminished as more white-collar workers did their own word-processing. Tyler (1998) suggests that to make best use of speech recognition, we should rethink our working practices and perhaps turn back the clock.

This need to review approaches to existing tasks and processes may have more specific **implications*** for some

Know less, understand more

Michael Dertouzos, Director of the MIT Computer Sciences Lab, here describes the Pegasus airline reservation system, which uses speech recognition, an artificial (but physical) ‘ear’, phoneme checking against a dictionary, analysis of pitch, emphasis, grammatical and linguistic construction, and both screen-based and spoken output. Pegasus works because it specialises; its knowledge is restricted to what is needed to deal only with airline reservations, rather than attempting to recognise any human utterance:

‘Without a narrow domain, a system like Pegasus would get about 95 percent of each spoken word correctly. Sounds good. However, compounding these errors over five-word sentences would cause Pegasus to correctly understand only about two-thirds of the sentences. Not so good. With the corrections supplied by a narrow domain and language constraints, [it] correctly recognises almost 90 percent of the spoken sentences. Though it is not the 99.5 percent we humans achieve, it is good enough.’ (Dertouzos, 1997)

professions. For example, language teachers who wish to use CALL software incorporating speech recognition technology may need to rethink approaches, and decide which aspects of learning could be achieved more successfully by the individual using the technology alone, and which learning should remain within the classroom as a shared, group activity. Similarly, whole-school policies – of funding new equipment, providing new learning spaces, and managing the politics of departmental resourcing – may need to come under similar scrutiny if language machine technology is to be used more widely and successfully in education.

Problem 6

Newcomers need training and time to learn to use the language machine

One supposed strength of the language machine is that communicating with a computer in spoken English is 'natural' and straightforward. However, current language machines are far from reaching the distant goal of providing a general-purpose system with linguistic capabilities sufficiently advanced to pass the Turing Test (pages 10, 24). To yield useable results, people need training and time to learn the constraints and limitations of a particular system. Furthermore, a language machine is likely to be one component in a much larger IT system: people need time to develop an understanding of how the language machine fits into the overall system.

In learning to use speech recognition, for example, a speaker must plan ahead and enunciate clearly, and learn (or re-learn) how to use software using their voice rather than a mouse and keyboard.

Many of today's computer-literate professionals are now comfortable with a mouse and keyboard; to these users, speech is not as self-evidently more 'natural' as it might have appeared to a previous generation less familiar with using computers.

**'The drawback of speech input for normal office dictation, from a management services point of view, is that time spent correcting the errors which the speech/text system committed cancels out the time saving won by speaking the words instead of keying them in ... Simply creating a speech recognition software package that transforms general correspondence speech input into text will not serve the professional end user ... We (Philips) allow the author to dictate the document and send a speech file for transcription to staff who can listen to the text while correcting. Having the author do the correcting undermines the essence of speech recognition as the author wastes time on administrative tasks and the typing pool sits idle.'* (Tyler, 1998)

What are the drawbacks of a conversation machine for the ELT business? See pages 34 & 35

expensive

and useless?

**Practical applications for the language machine have remained largely hidden from popular view, used instead in large-scale, third-party systems. However, the technology is being adopted in niche markets for specialist tasks.*

We can also expect to see it in common use for: creating commercial documents in the appropriate language; creating multilingual documents; providing translations; localising technical and instructional manuals.

Conversing with
your computer?
See pages 23 & 24

Problem 7

Applications face all of the above problems

Attempts to use the language machine in practice have been dogged by a combination of several or all of the above problems, as its application in English language teaching illustrates. Relatively few language teachers, let alone individual learners, can afford a computer powerful enough to use the latest language machine software. IBM ViaVoice Gold, for example, requires a PC with a Pentium 166MHz processor, 32Mb RAM and 125Mb free disk space. While this is no longer the sort of machine confined to research laboratories, it is still not easily within the budget of every teacher or student – even in Western countries.

Neither have very many language teachers been consulted by language machine researchers to discover what they would **use*** the technology for: the Project ISLE user consultation exercise described on page 30 is the exception rather than the rule. One of the main requests of language teachers was for a ‘conversation practice machine’: a computer system with which a learner could ‘chat’ in ordinary spoken English to practise everyday conversational skills. Unfortunately – as speech recognisers are not accurate or robust enough to cope with general unconstrained dialogue and conversational skills, particularly where the unconstrained dialogue comes from

Tower of Babel? Or Cubicle of Babble?

One potential scenario for the language machine is that it will become part of our everyday interactions with many kinds of equipment – from PCs to wristwatches. What will this mean for the office? The open plan office is a design that many people find distractingly noisy – sound barriers are already becoming more common – and if more of us spend more time talking than typing, barriers may become standard fittings.

At present, sound levels remain a problem to routine use in the office – noisy environments or telecommunication networks can make speech recognition less accurate. Noise-reduction technologies, specialist microphones, signal enhancement, and keyword spotting can all help, but accurate recognition depends on excluding other voices and sounds to prevent the language machine’s attempts to decipher them.

inexpert speakers – the current project has been forced to focus on pronunciation tutoring. **English language teachers*** also wanted the ability to correct errors in learners' conversations and give guidance on their causes. Our current language machines are not up to this task; in fact, we need to debate whether it is appropriate – or perhaps whether it is dangerous – to use today's technology in this way. With the current levels of performance from language machines, it is quite likely that the learner could say something quite reasonable to the system and be corrected; say something wrong and not be corrected; or (perhaps worst of all) say something wrong, be told it is wrong, but be given an inappropriate diagnosis or correction. This may not be as critical as the confusion of 'descend to a hundred feet' with 'descend 200 feet' in an air traffic control system (see page 32), but it could be harmful to a new language learner's progress – and it would almost certainly give the present language machine a reputation as an unhelpful, or even **damaging*** technology.

**Language teachers often need to highlight 'errors' – a word used to include a learner's choice of an inappropriate word or phrase for the context, a confusing pronunciation or a problematic use of grammar. Whether a learner makes an 'error' can be an issue for human judgement – beyond the scope of the present-day language machine.*

**But not all dangerous – there are many successful applications of speech and language technology in language teaching and education in general. See pages 48–49.*

The problem of the popular: Hal 9000 Series

Computers that can understand us when we talk to them, and answer back, have often featured in science-fiction, though their portrayal raises questions about our relationship with technology.

In the early days of Utopian science-fiction and the heyday of H.G. Wells, machines often provided, if not solutions, then some means by which humans could take control of situations; expanding their powers and abilities. Yet our faith in a technological future may have waned – especially when the machines began talking to us, thereby taking on a most human characteristic. The most obvious illustration is the red-lensed Hal in Arthur C. Clarke's 2001: a Space Odyssey – a softly-spoken malevolent presence with an agenda of its own – and perhaps reflecting a human fear of technology growing too powerful?

In more recent science fiction the vocal computer is a figure of fun: the humans have taken back the control. The computers in Red Dwarf or Douglas Adam's Hitchhiker's Guide to the Universe are comic creations that nanny the crews of their spacecraft – although not as ineptly as a Nutramatics drinks machine. The most positive example of speech technology in these scripts is credited not to a computer but to a fish – the babelfish which, when inserted into an ear, freely translates between the universe's languages in real time (before going on to inspire the naming of the first Web-based translation service introduced by Digital and Systran in 1997). What role for talking computers in science fiction when they're science fact?

determination

to succeed

**One scenario for the ELT business is a radical one: what if the language machine brought about the large-scale abandonment of language learning? Many people learn English so they can use it to communicate with native English speakers or as a lingua franca; but what if, in the future, they find that it is possible to communicate via machine translation devices?*

Is this scenario a realistic one? Implications for language professionals are explored on pages 50–57.

**Searching for Extra Terrestrials? 'Language detection is being applied to radio-telescope signals, to ascertain if language-like features exist within ... if ever such a structure were to be detected it would be from an intelligent source.' (Elliott, 1999)*

Hal – the
talking computer
made evil
See page 35

Can problems be overcome?

Taking a typical application – language teaching and learning – can we imagine that the problems described in this section will be overcome?

Clearly, full integration of the language machine is not foreseeable in the immediate future. At the very least, learners and teachers need to use computers confidently in everyday learning contexts. And, in its current state, the language machine is not yet appropriate as a replacement for some of the tasks undertaken by the language teacher.*

Yet as language technology improves, the boundary between tasks that can be delegated to the computer and those that should be done only with a human language teacher will shift. The advent of widescale personal computing may also have an impact in terms of allowing students to use language machines to learn some aspects of language outside the classroom. It is likely that new ways of teaching using computers will also spill into language teaching from other areas of education.

However the boundaries are redrawn, teachers will need time to learn to use, integrate and exploit the language machine. Schools will need time to see what's available, make funding plans against staff availability and projected student numbers, and then plan for maintenance and upgrades. In developing countries all this may be an unrealistic luxury: teachers in the state sector may not have updated equipment or software easily available to them or even offer the basic computer know-how.

Yet it is possibly easier for the state sector to provide computers than it is to provide smaller classes and trained teachers with the resources they need. For example this may be the case in Mexico, where the state ministry – while working with the British Council to upgrade the skills of its teachers – is keen to introduce new technology.

Interestingly for education, in applications such as machine translation (the focus of the box opposite) and for people working in **other professions**,* it is clear that the language machine is making clear inroads.

But do inroads lead to a revolution?

Machine translation? Speech recognition? Do they work?

The ease with which it is possible to find fault with any new technology is a significant factor in its take-up. Students of machine translation, for example, are often given examples of mistranslation, such as: 'The spirit is willing but the flesh is weak' translated first into Russian and then into English: 'The vodka is good but the meat is bad.' Many such examples can be found, which perhaps distort the successes of language technologies. But what do the professionals say?

Because quality is so uneven, monolingual users are advised to resort to MT only for 'inbound' work – texts being translated into their own language. In this case, they can see what they are getting, warts and all, and guess at the rest without offending potential clients' linguistic sensibilities. Unedited MT is simply not good enough for 'outbound' work such as promotional materials, where input from professional human translators is essential.

– Bulletin of the Institute of Translation and Interpreting, Feb-March 1999, page 29.

It becomes increasingly transparent that not only does machine translation not do away with the services of a professional human translator, it positively demands such services, because MT can, at its least sophisticated, put a 'spin' on the original meaning which makes this meaning even more difficult to tease out than from the foreign text source.

– John Thorogood, Centre for Information on Language Teaching and Research.

In reply to the question, "What effect will access to speech recognition technologies have on people's willingness and ability to learn new languages?" The main thing I have noticed about programs incorporating automatic speech recognition (ASR) – even though it is far from perfect at present – is that they are highly motivating. Programs such as TriplePlay Plus, Talk to Me and Tell Me More have been observed to encourage learners to try harder. The advantages of being able to practise in private rather than getting tongue-tied in front of a whole class probably outweigh the shortcomings of currently available ASR systems.

– Graham Davies, President of EUROCALL.

It's not wrong to have reservations about the relevance of computers to language-learning. After all, people learned foreign languages for thousands of years and got by just fine. It would indeed be a shame if a sharp new focus on computers threw up yet another intimidating roadblock to language-study for the poor or elderly.

– Neil Langdon Inglis, Senior Translator, International Monetary Fund.

Machine Translation is a wonderful tool. People will be able to get the gist of the material in order to find out if it is worth while translating. It will also make the translator's job a lot more interesting and better paid as they will be able to do a lot more work in a normal eight-hour working day. This will eventually weed out the translator who produces poor quality material.'

– Lucinda Mobaraki, T & I Services

Can the UK become 'Europe's digital laboratory' for such products? See pages 42–43.

UK advantage?

the key players

What achievements can the UK celebrate? What developments are happening elsewhere in Europe?

A global interest

Like other IT systems and products, research and development involves both universities and industry: while academic researchers have generally explored fundamental theory, industry has normally focused on turning theory into products for markets and consumers.

In the field of speech and language technology, the main companies developing and selling applications are transnational, rather than solely European or British: although many have offices or facilities in Britain, most have either American or Japanese parent companies. Much significant research and development takes place in the parent countries, with UK and European offices concerned with localisation of software, such as converting on-screen prompts and help-files to local languages, or providing local support and consultancy.

When assessing a country's productivity, examining export volume alone can therefore be misleading: although Ireland and the USA, for example, export roughly equal total volumes of software, Irish software exports focus on local, European versions of American products.

The EU is now trying to foster more homegrown research and development, in IT generally and language engineering specifically. Clearly, this is an appropriate strategy: there is massive global interest in language machines to merit the commercial investments in time, staffing and funding, and the mainstream IT big names – including IBM, Microsoft, Apple, and DEC-Compaq – will all have an interest in seeing the technology used and accepted. A number of electronics and telecommunications companies (such as Philips, Canon, Sharp, BT, Cellnet) are developing and selling related systems.

A third category of language machine players are smaller 'niche' companies specialising in speech and language technology, such as Dragon, Kurzweil, Entropic, Lernaut and Hauspie, Speech Machines, SRI, and Systran. Although small in comparison to the likes of IBM and Microsoft, these specialist language machine companies have a major stake in the future, as noted by Hunt (1988).

What's happening in the UK?

As well as basing research and development laboratories in the UK, the companies involved in developing speech and language technology can also collaborate with the many university researchers to mutual benefit: indeed UK researchers in universities and industry research centres are involved in all the component technologies and sub-fields listed on pages 20–25.

The UK community is also brought into regular contact via a *SALT Club* managed by the EPSRC (Engineering and Physical Science Research Council) and the Department for Trade and Industry, which organises meetings, newsletters, an email bulletin board and Web site (see page 63), designed for information dissemination and the encouragement of commercial applications for speech and language technology.

The EPSRC also funds a number of research projects at [universities](#)* across the UK; these are reviewed in EPSRC publications (EPSRC, 1996) which include summaries of EPSRC-supported research projects. Most of these projects involve collaboration with industrial partners, aiming to develop [commercial](#)* applications.

**The Universities of Bath, Belfast (Queen's), Brighton, Cambridge, Edinburgh, Essex, Imperial, Keele, Leeds, Liverpool, London (SOAS, UCL), Manchester, Sheffield, Southampton, Sussex, and York each research into language machines. (EPSRC, 1996)*

**'Until last June, Dragon Systems had not sold its products through retail channels in the US. Just six months later, the value of its monthly sales ranked number 13 in the list of all companies selling retail business software products of any kind in the US. It looks as though the chain reaction needed for the general acceptance of a new mode of communication with PCs may have started.' (Hunt, 1998)*

What do EU projects aim to do?

The EU is a major force in research, development, funding, collaboration and project support, yet many people remain unaware of project aims and achievements. Varile has identified three major focus areas:

1. adding multilinguality to information and communication systems, at all stages of the information cycle, including content generation and maintenance in multiple languages, content and software localisation, automated translation and interpretation, and computer assisted language training;
2. providing natural interactivity and accessibility of digital services through multimodal dialogues, understanding of messages and communicative acts, unconstrained language input-output and keyboard-less operation;
3. enabling active digital content for an optimal use and acquisition by all, through personalised language assistants supporting deep information analysis, knowledge extraction and summarisation, meaning classification and metadata generation.' (Varile, 1997)

UK-Europe

overview

Case Study: Lernout and Hauspie

The growing market for products based on speech and language technology is demonstrated by Belgium-based company Lernout and Hauspie, which has, since 1987, become a main European supplier, licensing technologies for clients in telecommunications, consumer electronics, automotive electronics, computers and multi-media.

Its customers include NEC, Deutsch Telekom, Pioneer, Unisys, Hitachi, and Samsung: Microsoft has bought into them to access the company's language technology, which it is likely to use in developing the next generation of voice-enabled computing for the Microsoft Windows platform.

Lernout and Hauspie's market position highlights several points: the trend to build in products; the likely mass-consumer markets ahead, supported by the major players – and the significance of multilingualism. The company has risen in a multilingual environment which may have assisted the development and take-up of products. This latter point may serve as a cautionary note to those companies operating in a monolingual context.

UK involvement across Europe

Many research groups also participate in European and international projects and networks as the following shows.

ELSNET: European research Network in Language and Speech This 'network of excellence' is funded by the EU, aiming to foster collaborative research and dissemination of results across Europe. The UK is a major participant, with 11 industrial and 12 academic sites (including the UK bases of several international companies, and many of the universities hosting EPSRC research). ELSNET organises European meetings, publishes a regular ELSNews newsletter, and runs an email list and Web site. It also commissions reports and surveys, including a survey of tools and techniques for machine translation teaching (Balkan et al, 1997).

ICAME: International Computer Archive of Modern and Medieval English Based at Bergen University, Norway, ICAME started in the early 1970s as an informal grouping of university English corpus linguistics researchers, English language teachers, and researchers who used English language corpora such as LOB, Brown, and London-Lund Corpora (Souter and Atwell, 1993). Although early conferences focused on academic applications in English language research and teaching, its corpus resources and researchers have moved on to other applications, including word processor error-detection and 'training corpora' for speech recognisers. ICAME holds annual international conferences, publishes a journal, and runs the *corpora* email discussion group and Web site. It also distributes English language corpora and related resources.

ELRA: European Language Resource Association ELRA was set up in the 1990s to archive and redistribute computer corpora, dictionaries, and other speech and language technology resources for a wide range of European languages (and some non-European language resources). ELRA publishes a newsletter and Web site.

For translation, security, medicine . . .

The following projects have **Europe-wide*** significance – of interest for UK researchers, customers and **markets***

MABLE (Multilingual Authoring of Business Letters): a system to allow standard business correspondence to be authored in a foreign language.

OTELO (Common Access to Translation Systems) and **TRANSROUTER** (Translation Router): tools and techniques to help translators make best use of machine translation.

LINGUANET (Communicating Through the Language Barrier), **AVENTINUS** (Advanced Information System for Multilingual Drug Enforcement) and **SENSUS** (LE for police and emergency service communications and information systems): language engineering systems to help European police forces to co-operate.

MIETTA (Multilingual Information Extraction for Tourism and Travel Assistance) for tourists and travel agents to query Web sites and databases in their own language.

EUROSEARCH: building a pan-European federation of national Web search engines and categorisation services, using machine translation so that people can search for and browse through Web-based documents in their native

**The EU has funded a range of programmes since the 1980s; the current Telematics Applications programme includes several projects linked to the theme of language engineering. See pages 62–63 for contact details.*

**One developing line of business is market analysis for clients looking to use language engineering. A Web search will reveal several companies now offering research, surveys and consultancy for policy-makers and publishers.*

Looking for Web site information on UK and EU projects?
See pages 62 & 63

The ten EU rules for language machines in telematics

The EU Telematics Applications Programme has produced 'Ten Commandments' for telematics research:

- 1 Focus not only on industrial competitiveness but also on support for other European Union policies.
- 2 Be user-oriented and cost-effective rather than technology-driven.
- 3 Start with market research and an analysis of users' needs.
- 4 Associate users' representatives at each stage of a project.
- 5 Focus on multimedia telematics rather than data telematics.
- 6 Concentrate on fewer projects and treat interoperability as a key issue.
- 7 Devote adequate resources to validation in user environments.
- 8 Maximise the generic content of telematics applications.
- 9 Exploit results, including standards, procurement and implementation recommendations.
- 10 Cooperate not only within a research sector (vertical coordination) but also between sectors (horizontal coordination).

(EU Telematics Applications programme (1996): <http://www.echo.lu/telematics/off-docs/introduction.html>)

supply & demand

the markets

**Who supplies speech and language systems? Niche suppliers may sell language machines not directly to the end-user, but as components of larger products and services. The channels of supply identified by EUROMAP are software vendors; telecommunications providers; value added network suppliers (including Internet Service Providers, or ISPs); communications equipment suppliers; information service providers; electronic publishers (including the news media); and systems integrators. (EUROMAP Web site, 1999)*

**The EUROMAP project (European Opportunity Mapping) has reported on a thriving European research and development community in speech and language technology: 'A total of some 10,000 language technology players in Europe had been identified, including research staff, vendors, and professional and corporate users, of which 7,000 can be assumed to be dealing with core language technology issues. Of these, more than 1,000 research sites, academic and commercial, were contacted, and around 300 provided information to the survey.' (EUROMAP Web site, 1999)*

language; MULINEX aims to complement this with multilingual tools for message extraction and text summarisation, to allow multilingual Web access, navigation, browsing, and filtering.

DocSTEP is developing an authoring system for creating technical documents written by several authors in a range of European sites.

MAY (Multilingual Access to Yellow Pages) will allow users to find business telephone numbers in a different European country in their own language.

GALEN is using language machine technology to assist in cross-European harmonisation of medical terminology.

ADVISER offers better access to European research results, by extracting such information from a range of Web and other resources across Europe and translating into a standardised format.

How are people encouraged to buy?

Much of the work now being completed in speech and language projects across Europe emerges as products for large organisations – pan-European administrations, transnational commercial organisations, or security and law enforcement agencies working across borders – but what can encourage the smaller organisation to **invest*** in speech and language technology?

Elsewhere in this book issues are raised – of cost; scale; training; dialogue between end-user and researcher – but one key ingredient in bringing a product to market is that of **information**.

Several projects now aim to develop general language resources and infrastructure, to support and promote the take-up of language engineering, and to raise awareness of markets, customer needs, standards and potential uses:

EUROMAP* (European Opportunity Mapping), EUROMAP II, and LINGLINK (Promotion and Support for Language Engineering in Europe) are surveying and promoting language engineering applications, helping to bring the technology to markets and users.

EAGLES (Expert Advisory Group on Language

Engineering Standards), MATE (Multi-level Annotation Tools Engineering), ELSE (Evaluation in Language and Speech Engineering) and DIET (Diagnostic and Evaluation Tools for Natural Language Applications) are developing standards and guidelines for the development and evaluation of language engineering resources and systems.

PAROLE (Preparatory Action for linguistic Resources Organisation for Language Engineering), EUROWORDNET (Building a multilingual wordnet database with semantic relations between words), SIMPLE (Semantic Information for Multifunctional Plurilingual Lexica) and SPEECHDAT (Speech Databases for Creation of Voice Driven Teleservices) are gathering corpora of text, speech, and dictionaries across the range of EU languages, to be re-used in European language engineering research and products.

One of the key questions many will ask, is how will these language machine projects benefit [people in the UK](#)?*

**'My vision for the UK is quite simple: to be Europe's digital laboratory. I want Britain to be the test bed for digital products and services in Europe, so that UK consumers have access to these first and so that British business can lead the world.'*

Peter Mandelson, speaking as Secretary of State for Trade and Industry, at 5th Annual CEO Summit on Converging Technologies, September 1998

What are the benefits to the UK?

Pursuing the language machine brings benefits and opportunities to the UK at every level – from competitiveness in markets across the world and collaboration with European-based agencies and researchers, to the practical, everyday situations in the home, service industry, office, and school.

EUROMAP suggests several potential markets – authoring tools as part of office automation and document management systems; speaker identification and authentication, speech recognition and synthesis embedded in call centres and built into the telecommunications infrastructure for voice messaging and voice delivery of email; information extraction embedded in knowledge management systems; semantic analysis in information retrieval services (including Web search engines); machine translation services supporting international messaging and multilingual information retrieval (EUROMAP, 1998).

UK involvement – in product manufacture, whether for end-user or third-party, support or supply – in any of the markets identified by EUROMAP would help position Britain as a key international player:

As a major participant in the development of speech and language technology.

As a provider of opportunities for cooperative research.

As a source of commercial experts and suppliers of technologically sophisticated equipment.

As a creator of contexts for diplomatic, security and business exchange across multilingual boundaries.

As a base for creative talent – and the adoption of the language machine for cultural activities, leisure, entertainment: for use by artists and visitors to galleries, libraries and museums.

Real-life

situations

Who's using the language machine now? Which contexts and customers are likely to benefit from the next wave of commercial products?

This section highlights how speech and language technology is used in diverse contexts – from helping children tell stories to conflict resolution.

Is there a typical 'user profile'?

The range and variety of potential uses of the language machine are as broad as the uses of language itself: any and every aspect of language can be studied and modelled. One way to survey the field is therefore to look at a selection of 'typical' past and current language machine applications.

Industrial interest in the language machine has tended to focus on broad application categories; in a market survey for IT managers, Johnson (1985) saw six main categories of use (see box opposite). Language machine researchers, on the other hand, have historically focused on their own academic interests and perhaps have failed to communicate potential uses and benefits to people outside of the field. An EPSRC report tried to reallocate the blame:

In general, UK industry seemed to consider that too much effort had been spent in the past on 'esoteric' academic problems ... Another problem was the lack of awareness of the applicability of speech and language technology amongst potential users. It was also felt that there was a lack of useful, analysable feedback.

(EPSRC, 1998)

The EU's Telematics Applications programme is seeking – like other EU-funded projects – to redress this imbalance, by requiring projects to consult and involve users at all stages and ensure a greater focus on applicability.

In defence of researchers, the breadth of potential application of SALT makes it difficult to characterise 'typical' uses – and hence consult 'typical' users – to a degree significantly finer than the broad application areas identified by Johnson.

Indeed, the EAGLES Report on evaluation of natural language processing systems (EAGLES, 1995) concluded that there is no typical user profile; the main determinants are an organisation's resources and its policy towards translation and languages. In size, users range from international organisations like the United Nations, EU, and IBM, down to individuals using any or all of the above categories

of applications on home or office PCs. Possible policies include the use of in-house translators, outsourcing, and use of general Internet-based language services.

Labs, factories, outdoors, offices, homes

'Safe-zones' for hands-free use of dangerous equipment or when making hazardous manoeuvres? Cross-border **policing**?* Securing sensitive areas by voice-operated access points? Providing distance education and self-paced learning for children and adults? New ways of using everyday domestic equipment? Most situations lend themselves to speech and language – below are just seven examples of language machine systems to illustrate the breadth of potential applications. Given the provisos that should be attached to any attempt at technological forecasting and the breadth of range of potential applications, deciding which of these – at this stage in language machine development – are 'typical' remains open for debate.

Peace-keeping 'Soldiers in Bosnia ... wear a small computer on their chests and say to it "Hands up" or "Get out of the car" or other things that soldiers have cause to order Bosnian civilians to do. The computer speaks the command in the local language. (Soldiers can get it to speak back to them in English if they want to be sure it is going to interpret the correct command.) ... Dragon Systems developed the Bosnian tool.' (Tyler, 1998)

Speech and language technology in the classroom

See pages 48 & 49

**Language technology in action for policing across Europe: national police forces across Europe are already pooling their language resources via the LinguaNet project, set up in 1995 and now linking over 30 police sites in seven countries. The system – an EC backed venture – uses machine translation for vehicle checks, missing persons, bank card enquiries and accident reports. The system may soon be extended by speech output and speech recognition facilities.*

How are people using language machines?

Database interfaces: front-ends to computer databases that allow people to ask for information by posing questions in ordinary English. **Dialogue interfaces:** more sophisticated front-ends, allowing people to use computing resources by means of a natural English dialogue with an apparently intelligent program (a Knowledge-Based System). **Content-scanning:** scanning electronic text (such as email, newswire services, or World Wide Web documents) to extract and/or summarise key information. **Text editing:** 'smart tools' to check grammar, idioms, and style are now options available in many word processors. **Machine translation:** translation of text from one language to another; more ambitious projects such as the German 'Verbmobil' aim to translate from spoken input (German) to spoken output (Japanese) to allow a 'virtual dialogue'. **Talkwriter:** transcription of spoken input into a word processor – this is seen as one of the main markets for language technologies. (Source: Johnson, 1985)

employment

society, leisure

**Some researchers argue all computer information should be available to us in graphical form: 'We are very pictographic animals, able to absorb animated, 3D colour images at a phenomenal rate. It is unnatural for us to read and write, or interpret spreadsheets; we are soon overloaded by information in such formats. Moreover, many of our species have a natural ability mentally to translate 2D drawings (plans) into a 3D world (models). And yet, if information is presented in a natural 3D form, we are all inherently able to absorb and understand the equivalent of a 20-volume encyclopedia in about 15 seconds.'* (Cochrane, 1997)

**Creating a better society? The EPSRC comments: 'In social terms, work on non-indigenous minority languages could have great benefits. ... It is estimated that an imperfect understanding of English can be a contributory cause of ill health amongst immigrant populations, as a lack of understanding means the speakers of non-indigenous minority languages are unaware of sources of advice, or have difficulties in following the instructions on medicines. There is also a need to provide aids for officials (and) doctors.'* (EPSRC, 1998)

Educational games 'I used a program called Storybook Maker with my daughter ... We would create a story together: she would dictate it to me, and I would key it in. She would add **pictures*** using the built-in picture maker and record her voice to create a multimedia presentation. She could get her story read back to her by clicking on a button, and the robotic voice was great fun.' (Private correspondence.)

Call centres 'Lufthansa has ALF, a friendly flight information service which holds conversations with callers at some 300 airports. The caller can speak naturally, he/she is not restricted to single words or touch-tone responses to those laboriously recorded sets of instructions about endless options in which one is not interested. ALF is based on the Speech Mania system from Philips.' (Tyler, 1998)

World Wide Web translation 'AltaVista, which is owned by computer giant Digital, launched a free machine translation service on the Internet at the beginning of 1998. The service is able to translate Web pages from English to French, German, Spanish, Italian and Portuguese or vice versa almost instantaneously. The results are far from perfect, but good enough to give the gist of a Web page. In May, AltaVista followed this coup up with the launch of a single world-wide index, which allows even Chinese, Japanese and Korean users to search across the Web in their own language from one central location ... The index translates the characters within documents into a standard encoding system, Unicode, so that it can identify words in languages as diverse as Chinese, Japanese, Korean, Cyrillic, Greek, Hebrew, Arabic and **Turkish,*** as well as western European languages.' (Pringle, 1998) (See pages 31, 37)

Email translation 'The most dramatic change of all has probably been the use of MT for electronic mail ... Usage has rocketed ... The use is not simple curiosity, although that is how it often begins. CompuServe records a high percentage of repeat large-volume users for its service: about 85% for unedited MT – a much higher percentage

than might have been expected. It seems that most output is used for assimilation of information, where poorer quality is acceptable. The crucial point is that customers are prepared to pay for the product – and that CompuServe is inundated with complaints if the MT service goes down!’ (Hutchins, 1996)

Information monitoring and retrieval ‘The SRI International Highlight system uses natural language processing techniques that help users find relevant information in large volumes of text and present it in a structured fashion. For example, it can extract information from newswire reports for a specific topic area – such as global banking, or the oil industry – as well as current and historical financial and other data. Although its accuracy will never match the decision-making skills of a trained human expert, Highlight can process large amounts of text quickly, allowing users to discover more information than even the most trained professional would have time to look for.’ (SRI, 1998)

Working while driving ‘IBM, Delco Electronics, Netscape Communications and Sun Microsystems ... collaboration has resulted in ... a demonstration-stage application in which car and lorry drivers* and passengers use voice commands to activate normal telephone services but also to get e-mail messages converted to listen to them on the move, to dictate replies, find a nearby restaurant or hotel,

**In-car speech recognition is now developing rapidly: several manufacturers offer ‘mobile multimedia systems’, combining computer and audio with speech recognisers responsive to simple voice commands. A driver may record memos, tune the radio, select a CD track or retrieve contact data. Information can be fed back via speech-synthesis. Microsoft, Clarion, Mecel, Delphi Automotive Systems, and Saab are all involved in development.*

**‘Research showed that drivers using a mobile phone – whether hand-held or hands-free – were four times more likely to have an accident during, and up to five minutes after, a conversation ... At least six road deaths have been linked to mobile phones in Great Britain.’ (The Royal Society for the Prevention of Accidents, Press Release, 25 February 1999)*

The art of translation. Or the translation of art?

Each time a technology comes along, designers, artists, theatre producers – indeed anyone working in the creative arts – will use it. The Internet, for example, has spawned the Degrees Feet and Inches Web site, dedicated to supporting artists and curators seeking to work specifically and exclusively online. Mixing human and machine translation, is also emerging as an art form – and perhaps an origin for new creativities:

‘Muntadas’ initial plan for On Translation was to create a chain of translations of a short text through 23 different countries/languages, an idea which is based in part on children’s games such as “telephone”, “Stille Post” or “Chinese whispers”. The work also bases itself on the translation process while utilising the Internet as a system and network. In its finished state, On Translation can be visited and used as an archive, with the work itself an artefact.’ (Rhizome Digest, 27 May 1999)

educational

resources

Have a go? Language technology by email

Add a twist to a language class: email English text to:

amalgam-tagger@scs.leeds.ac.uk

You'll receive back the same text, with each word annotated with the 'parts of speech'. The following sentence was emailed, and the parts of speech returned:

Remote access over the Internet, via WWW or email is an increasingly popular alternative.

Remote –adjective

access –noun, singular, common

over –preposition

the –article, singular or plural

Internet –noun, singular, proper

, –comma

via –preposition

WWW –noun, singular, proper

or –conjunction, coordinating

email –noun, singular, common

is –verb 'to be',
present tense
3rd person singular

an –article, singular

increasingly –adverb

popular –adjective

alternative –noun
singular, common

. –full stop

The text is annotated according to the LOB Corpus tagging scheme.

get route-finding help projected on the windscreen ... and do pretty much everything except actually drive the vehicle.' (Tyler, 1998)

Focus: language education & teaching

How do these broad categories of language machine use relate to language teaching? The following demonstrators are available for teacher use over the Internet.

Laureate BT's text-to-speech synthesis system: type in English text and hear it synthesised (with an artificial British English voice). Students may assess the 'naturalness' of pronunciation and intonation.

The Interactive Syntax Tutorial teaches you how to write formal phrase structure grammars for a fragment of English, and then use them in a parsing program.

The Internet Grammar of English A Web-based tutorial system aimed at English language students and teachers; students can experiment writing some part of this Internet grammar of English in the formalism used in the Interactive Syntax Tutorial.

The AMALGAM Part-of-Speech tagging service, developed with EPSRC funding, accepts English text emailed to amalgam-tagger@scs.leeds.ac.uk and adds a part-of-speech word-tag to each word, then emails the tagged text back to the original sender. Students can email a piece of English text, and then check the reply for part-of-speech tagging. (See box, this page.)

The World of Language This planned Exhibition Centre focuses on language and in particular English; initial proposals include ideas for more sophisticated SALT demonstrators (Atwell, 1997) of potential interest to English language learners and teachers:

touch a map to show where you come from and hear yourself greeted in your own language or dialect ... pick up a phone and talk to Chaucer /Shakespeare /Churchill.

Educational projects in progress

The following projects focus on language engineering applied to language education and training – and may indeed have longer-term educational implications.

RECALL: Repairing Errors in Computer-Aided Language Learning developed a multimedia environment for language learning, with graphical, written, and speech output.

ISLE: Interactive Spoken Language Education is integrating speech recognition and multimedia to develop an English language learner's pronunciation tutor. (The ISLE project is referenced throughout this book: see page 62 for a review.)

LETRAC: Language Engineering for Translator Curricula is developing teaching resources and components to include in university courses for translators.

SELECT: Strategies for European LE-Enhanced Communication Training is building a Language Learning Workbench that combines a wide range of existing language learner's aids within a user-friendly interface.

SPEAK: Supported Prototype Easy-access Authoring Keys developed a multimedia development environment to simplify the authoring of language learning materials and adaptation of materials to specific learner groups.

Academic resources for free? The researcher, teacher or student now finding out what language machine technology is available will discover that many resources have been developed primarily for 'in-house' use, and may be acquired from the originators. Many academic resources may be 'free' (i.e. they can be downloaded over the Internet without charge), but installing and developing an understanding of them can be time-consuming, expensive or – where resources are poorly documented – **unproductive**.* Remote access over the Internet, via WWW or email is an increasingly popular alternative, however, and may allow people who might otherwise never experiment with speech and language technology to do so (see box opposite).

**Unproductive hours* trying to make machines do what you want? The best way to integrate the language machine into teaching may be not for individual teachers to struggle to devise ways of using these demonstrators. Speech and language researchers and education must work together, as is now happening in the case of some European Telematics Applications projects that are aiming to develop systems for practical use. For those teachers eager to use speech and language technologies, demonstrators are available on the Web. Practitioners will need skills beyond IT knowledge: the ability to discriminate and evaluate products and uses; management strategies for individual and class access; the confidence to manipulate and control software tools; and the ability to integrate teaching with language machines within a pedagogic framework.

For Web-based information on projects, see pages 62 & 63.

Magazines – the route to updated information
See page 61

Future scenarios

the impact

What will happen from here? What will be the next stage of development?

Will the dreams become nightmares?

The next 20 years

Over the next 20 years, we can assume that the drawbacks of speech and language technology are likely to be overcome or surmounted, allowing the language machine to make a more significant impact. Bearing in mind the difficulty inherent in predicting future technological development and adoption, this section explores a number of possible future scenarios.

Scenario 1

IT and language machines will be all-pervasive

IT industry **futurologists*** agree that personal computing technology will continue to drop in price and increase in power, making it accessible to an ever-widening community. Personal computing will become as everyday as wristwatches and telephones; indeed, it will be *integrated* into wristwatches and telephones.

The language machine will be a core component in this new technology: wristwatch or telephone computing will not be able to rely on keyboard input and screen-display output, and speech input/output seems the most likely replacement. Today's students may each have their own wristwatches; in ten or 20 years' time, they may carry their personal computing around with them as an everyday tool or fashion accessory. Software for this consumer market will be made available: prime suppliers may be the language teaching business alongside IT industries.

Scenario 2

The language machine will get better

At the moment, it is easy to ridicule the language machine by showing examples of its fallibility. Most applications have a threshold of acceptability: a system that makes too many mistakes will not be used, although users will put up with it if the error-rate is low enough. Many computer-literate professionals have learnt to touch-type at speeds of 100 words per minute or more; speech recognition software has little appeal to them, and might be difficult to use in

noisy environments such as open-plan offices. On the other hand, many computer users struggle by typing with two or four fingers, risking repetitive strain injuries. As speech-recognition error-rates fall, there will inevitably come a time when personal thresholds of acceptability will be reached for language machines.

And indeed, the technology can only get better: personal computing software generally, and language machine technology in particular, has improved beyond recognition over the past decade, and there is no reason to suspect any imminent slow-down in this rate of development. This poses interesting questions for those currently teaching the upcoming generation of computer users. Should parents buy their children a typing tutor program? Or can they assume that they will simply talk to their PC in the future?

Scenario 3

Customised language machines will be used by individuals for particular tasks

One contemporary problem is already being addressed: researchers and developers have started to build specialised applications in consultation with people in **practical contexts***. For example, the ISLE project (see page 62) combines speech recognition technology with exercises from a popular series of English language textbooks and related learning resources (Klett's 'Bridges' package).

There is also a general trend in software engineering towards user-friendly software development toolkits and environments (such as Microsoft's Visual Basic, or Borland's Delphi), which allow people to customise or even build their own software systems from components. These systems – often referred to as visual programming – extend back to real building bricks in the programming environment provided with Lego Mindstorm, with which children can build their own simple programs for controlling Lego robots.

Language teachers, who often prefer to customise materials for the classroom – combining materials from a

**Predictions from many IT futurologists (Cochrane, 1997; Pearson, 1997) target Western society: middle-class and business consumers. Yet the status symbols of development – watches and televisions – quickly become available worldwide. Thus there is no reason to think that language technology will remain exclusive to developed countries. Present inequalities do, however, slow take-up: salary differences are highlighted on page 29.*

**e-talking? At Comdex 98, Motorola's Lexicus Division showed its Message Connect system, a speech-recognition combination of email and voice messaging which reads email over the phone; a listener can respond by dictating email back. (ZDTV Web site)*

Professionals will adapt to make use of the language machine. How?
See pages 52 & 53

implications

the professions

Rethinking our assumptions – and our self-image?

One of the implications of living in a digital age – where speech and language technology is fast becoming a daily reality – is that, as consumers and citizens, we need to adapt: remaining flexible to new working methods and open-minded to innovation and change.

Today, for example, we might be using the keyboard – tomorrow we need not to be self-conscious when talking via a microphone to the computer on the desk in front of us:

'Omnitel Pronto Italia, Europe's second-largest mobile phone operator, has launched the first Internet portal or gateway that can be navigated by talking to it using natural speech.

Omnitel 2000 provides access to up to 300 databases using TrueDialog and SpeechMania speech-recognition software developed in Austria by Philips, the Dutch consumer electronics giant.

The Omnitel portal can also be accessed via a web browser or a smart-phone using WAP (Wireless Application Protocol) software.'

(The Guardian, 17 June 1999)

range of sources rather than following a single textbook – may, in ten or 20 years from now, just as naturally select and combine speech and language technology materials into class material to create their own personal 'look and feel'.

Scenario 4

No aspect of language professions is immune

All areas of the language professions are likely to be affected to a greater or lesser degree. Language teachers have a particular right to be wary of any error-prone technology: although language teaching may not be as immediately safety-critical as air traffic control, a language machine that cannot fully discriminate when dealing with language will probably be widely judged as an unhelpful teaching aid.

This issue is crucial in speech and language systems that aim to detect and diagnose errors in learners' language production. One standard solution is for the system to assess its 'confidence' in an error: a pronunciation error is pointed out to the learner only if the system has a high degree of confidence that an error has actually occurred. Although this strategy tends to stop the error-diagnosis system giving incorrect advice, it also means that some genuine errors go uncorrected.

But is this so different from real – human – teaching practice? Teachers do not correct every imperfection in learners' English, but instead focus on specific weaknesses and/or language features of the day's lesson. Viewed in this way, it is hard to identify any aspect of language teaching where the language machine would be entirely inappropriate.

Scenario 5

Language professions will adapt to make use of the language machine

Language teachers, like many professionals, have a natural wariness of change for change's sake, and of new teaching

resources that call for radical changes in teaching methods. The relatively cumbersome and expensive PC of today may look to some like a successor to the cumbersome and expensive audio-visual technology of the language laboratory – an approach to language teaching that did not win universal popularity.

Unlike the language laboratory, however, the language machine is a technology that can permeate **everyday*** life; its acceptance into the classroom should be, therefore, much more natural. By analogy, many language learners in 20 years' time may be accustomed to using their own everyday personal computing resources and expect language learning to conform to other aspects of everyday life in making appropriate use of such technology.

The potential shift of balance between aspects of learning 'taught' in a classroom environment and those 'learnt' autonomously by students equipped with their own personal language machines may also prove significant to curricula and educational decisions. The teaching professions need to ensure that they are in a position to exploit the language machine to best advantage.

**Ownership of personal computers continues to rise in the UK – but will access to new technology become a new form of discrimination, social stratification, or inequality? Questions to help forward the debate on the impact of the language machine are listed on page 64.*

People with speech, hearing or sight difficulties? People needing English urgently for health and welfare?
See page 54

Scenario 1: the must-haves – handbag, lippy, wristwatch translator

Already, thanks to discreet ear-pieces that allow mobile-phone users to converse on the phone hands-free – often with the phone out of sight – we are becoming used to seeing solitary speaking. The future begins:

'In the next millennium, we will find that we are talking as much or more with machines than we are with humans. What seems to trouble people most is their own self-consciousness about talking to inanimate objects. We are perfectly comfortable talking to dogs and canaries, but not doorknobs or lamp-posts ... One thing that will make this ubiquity of speech move more rapidly today than in the past is miniaturization ...

Once you abandon the constraint of the natural spread of your fingers, which determines what makes for a comfortable keyboard, a computer's size is driven more by the size of pockets, wallets, wristwatches, ballpoint pens, and the like ... For all these reasons, the trend of increasing miniaturization is bound to drive the improvement of speech production and recognition as the dominant human-computer interface with small objects. The idea that twenty years from now you will be talking to a group of eight inch-high holographic assistants walking across your desk is not farfetched. What is certain is that voice will be your primary channel of communication between you and your interface agents.' (Negroponte, 1995)

questioning

the future

The digital helping hand

Language machines can help disabled people to more easily access technology: the following indicates some present activities.

The Royal National Institute for the Blind (www.rnib.org.uk) offers extensive information on speech synthesis packages and suppliers.

The EU's ISEAUS project aims to use language engineering in speech education and rehabilitation for adults with lifelong or acquired deafness (see page 63).

The funding body EPSRC actively encourages researchers to work on visual and pictorial languages for the hearing impaired and for use in noisy environments: 'Such studies could benefit, for example, hearing impaired users where a mixture of gesture, language and speech is needed.' (EPSRC, 1998)

A key UK researcher in this area is Masoud Yazdani, whose work on computer-based iconic language can be explored at his Web site: <http://www.media.uwe.ac.uk/~masoud/>

Some social projects are also in progress: Lernout and Hauspie is involved with a clinical reporting and transcription product with MedQuist. (Case study: page 40)

Having misgivings?

Clearly, there is great potential for the language machine to become a normal component in the personal computing 'kit' that is now so easily and widely adopted – and expected – by new generations.

But along with such easy adoption of new technology lie old fears – will machines replace humans? With the language machine to assist them, who will want to learn languages or use the services of a human interpreter or translator? Even if large numbers of people want to learn English for example, will their human teachers be replaced by language machine based self-tutoring software?

Learners prefer the human touch Even if we can afford the language machines of the future, that does not mean we will necessarily choose to use them. Language students, for example, may prefer human interaction: even when language machine-based software offers comprehensive English language self-tutoring facilities, many students will prefer a class led by a human teacher: teachers, after all, can offer what machines cannot – wit, warmth and empathy – as part of a learning experience.

Likewise, human dependency on machines can appear cold in social situations. In international communication, for example, particularly in face-to-face meetings, a reliance on a machine shows an unwillingness to learn to converse even at a basic level in more than one language – and may be judged adversely in terms of an individual's socialisation, diplomatic skills and hospitality.

People offer quantity – and quality As stated elsewhere in this book, the language machine is far from perfect: it is difficult to believe that its overall command of English will match (let alone exceed) that of an advanced learner of English, even in 20 years' time. Translators and teachers – capable of judging social language, multiple meanings, nuance, subtle forms, and language complexities – do not need therefore to fear they will be easily replaced!

More probably, machine translation will be used alongside good teaching and translation. It is used already

where a perfect translation is not essential – for example, in scanning and searching multilingual World Wide Web documents to track down information (Pringle, 1998). This kind of use of the language machine is bound to grow, but it is an addition to – not a replacement for – the current international flow of information.

The market for language professionals will expand As the language machine improves in accuracy and acceptability, more users will flock to it; but how many of these will be new users of multilingual resources who might otherwise have been put off by the effort of learning a new language?

The Internet has led to a huge growth in transnational access to multilingual documents, with CompuServe (now taken over by former rival AOL) translating 25 million words in the first year of its machine translation service (Flanagan, 1997). Yet little – if any – of this use replaced human translation. Without the possibility of machine translation, these documents probably would have been accessed only by those fluent in the language in which they were written.

The language machine holds out the promise of dramatic growth in multilingual **information interchange*** and communication. Machine translation services meanwhile will be used to produce ‘gist’ versions of documents that would not otherwise be translated – leaving a healthy market share for the practitioners of a high-

**A seminal paper on ‘The proper place of Men and Machines in language translation’ (Kay, 1980) argued that there were aspects of high-quality translation that called for uniquely human knowledge and experience. Nearly two decades on, the Machine Translation journal revisited Kay’s arguments, but was unable to refute them completely.*

While the visions may remain, changes in social attitude are vital if technology is to be accepted. Both men and women are now assumed to be creating, controlling and using new technologies.

The questions
we need to ask?
See page 57

Good times and fortunes

‘The very familiarity of MT systems will alert a much wider public to translation as a major and crucial feature of global communication, and probably to a degree never before experienced. Inevitably, translation will itself receive a much higher profile than in the past. People using the crude output of MT systems will come to realise the added value (i.e. higher quality) of professionally produced translations. As a result, the demand for human produced translation will rise, and the translation profession will be busier than ever. Fortunately, professional translators will have the support of a wide range of computer-based translation tools, enabling them to increase productivity and to improve consistency and quality. In brief, automation and MT will not be a threat to the livelihood of the translator, but will be the source of even greater business and will be the means of achieving considerably improved working conditions.’ (Hutchins, 1998)

leading the way

the next steps

Verdicts on
machine translation
See page 37

**Future development of language machine technology must include people who actively use the products for ordinary tasks – as highlighted here by Derek Child, Assistant Director Equal Opportunities, The Open University.*

'For the visually impaired, independent access to information has always been a barrier, and ICT offers the greatest liberation since Braille and audio-cassettes. But compared to older text-based systems, the visually-oriented Graphic User Interfaces can raise new barriers.

The conventions set out by RNIB might inform future software development and spare blind users the need for expert help in assessing products for their particular use.'

For RNIB Web site see page 54.

quality, human-centred language industry.

Language offers cultural resources It is very possible that there will remain some aspects of language use and multilingualism where the intervention of the language machine will be viewed as inappropriate or intrusive. Many people will want access, for example, to the rich 'cultural resources' of English, and hence prefer to learn English for personal and cultural fulfilment even though machine translation is available. Although they would not consider it an issue for technical or trade documents, they may choose not to read translated works of literature, believing that translation compromises the integrity of the original text.

Not only art and culture may lie beyond the reach of the language machine: quality is a watchword in business, and many customers will continue to equate care and service with social interaction, as will businesses who use human interaction as a selling feature of quality customer care.

Language is used for gate-keeping Latin was once a requirement for access to particular institutions or courses – veterinary science degree courses for example – demonstrating how proficiency in English (and indeed any language) can be used by educational systems and employers as a criteria for access to better prospects, even when individuals don't need the language itself for routine professional purposes. This is likely to continue in various subtle ways: it may even be that access to such technology becomes a new means of discrimination.

Ways ahead

An earlier book arising from the British Council English 2000 project, *The Future of English?* (Graddol, 1997) 'aims to provide thought-provoking ideas.'

Similarly, this book, *The Language Machine*, aims to provide background information and fuel for ideas. How might we develop the language machine? Can it be exploited for commercial benefit? What practical uses will it have for **people*** whoever they are? Here we identify ways forward to continue the debate.

Supporting a debate on the language machine and language professions Stakeholders in the language professions – including the British Council and language practitioners – should support ongoing debate of the issues around language machines, through presentations and seminars, further publications, email lists, and Internet dissemination and discussion.

Building better forecasting models The forecasting in this book is based on personal perspectives, on speech and language technology research, and IT futurology. Future debate should be able to draw on formal forecasting models and techniques used by professional futurologists, such as the UK Futurology Group.

Scenario building Uptake of the language machine is likely to vary substantially around the world. *The Future of English?* advocated ‘building scenarios for English in different parts of the world ... to explore further the impact of the complex interaction of global economic and technological trends.’ Scenarios for the impact of the language machine on local and regional language industries would similarly provide further assessment of likely needs, markets and potential impact.

Brand management *The Future of English?* advocates the co-ordinated promotion of ‘Britain as a leading-edge provider of cultural and knowledge-based products’. As we’ve already seen, many major providers of speech and language technology are European or transnational; the UK makes a major and significant contribution to all areas including research, development, and information dissemination.

We have before us now an opportunity to widen the debate in an emerging technology; to promote the British ‘language machine’; bringing together a potentially powerful and creative combination of language and technology, ‘positioning Britain as one of the 21st century’s forward-thinking nations.’

Wild child – or wired child?

Technology and youth culture grow ever closer. Portability has made technology something that can be taken – and posed with – anywhere.

Since the launch of the Sony Walkman in 1979, there has been a tide of high-spec, youth-targetted products: the portable CD-player; the mini-disc player/recorder; MP3 players that allow CD-quality music to be downloaded from the Internet; and MIDI – a computer protocol for musical instruments with which teenagers can record studio-quality albums in their bedrooms, before using a CD-writer and the Net to press and distribute them.

They can then use their mobile phone – voice-activated, capable of surfing the Net, sending email and faxes – with which they’ll launch the marketing operation for their first CD, compiling, of course – like techno, house, jungle – a technology-dependent form of dance music.

And speech? Easy.

Radiohead’s OK Computer (EMI, 1997) featured the track – fitter, happier – with lyrics ‘sung’ by an Apple Macintosh.

glossary

the jargon

AI – Artificial Intelligence, first conceived by Alan Turing in 1950 as a term for computer or machine intelligence: machines that can think and learn like people. AI usually involves the use of knowledge bases in which large numbers of rules are generated to govern the application's behaviour.

algorithm – a formula or set of ordered steps for solving a (typically complex) problem: the instructions in a computer program.

Ascii – American Standard Code for Information Interchange. Any roman character, from A to Z, plus numbers and punctuation marks, can be translated into Ascii code to be sent and received by computers.

browser – a computer program that can read Web pages. Many browsers are now multifunctional and can retrieve files, send mail and read newsgroups.

CALL – Computer Assisted Language Learning. Software that presents activities designed to help language learners. Many of the activities are based on conventional approaches, such as cloze.

computational linguistics – an area of applied linguistics concerned with the processing of natural language by computers.

controlled language – language designed to make language processing easier by restricting the number of words and structure of language used; commonly used in areas where precision and speed of response is critical, such as the emergency services or air traffic control.

corpora – a corpus is a body of language, either text or speech, used as a basis for analysing language to establish its characteristics. As well as national corpora of hundreds of millions of words, corpora are also constructed for particular purposes, such as recordings of car drivers speaking to simulated voice-activated control systems to help establish their marketability.

discourse – a continuous stretch of language comprising more than one sentence.

discourse modelling – the analysis of linguistic phenomena that range over more than one utterance or sentence.

GUI – Graphical User Interface. A visual interface for computers in which use of text commands and function keys is superseded by the use of icons and windows.

handwriting recognition – see Intelligent Character Recognition.

Hidden Markov Model – probabilistic pattern-recognition model used in speech recognition systems to help to determine which words are represented by the sounds that the computer has captured.

icon – a small on-screen image that represents an application or document in a graphical user interface (GUI).

Intelligent Character Recognition – involves word recognition techniques that use language models, such as lexicons or statistical information about word sequences to overcome the difficulty that OCR techniques encounter when attempting to recognise handwriting, unknown or decorative fonts or poor print quality.

language engineering – the application of knowledge of language to the development of computer systems that can recognise, understand, interpret and generate human language in all its forms.

lemmatise – to break an inflected word into its root (i.e. its base form) and ending components e.g. translating = translate + ing

localise – to adapt software to local requirements in terms of language and culture (including, for example, legal practice and business conventions).

machine learning – the area of AI concerned with algorithms for finding rules and patterns that describe sets of data: building abstract representations of what has been experienced.

machine translation – often abbreviated to MT, the process of automatically translating from one language to another by a computer.

MAT – machine aided translation: the process of assisting a human in translating from one language to another using computer software tools.

multilingual – refers either to something existing in a form that can handle several languages or to something that exists in several languages.

natural language – English or any other human language when used in interaction with a computer; a natural language query application, for example, allows users to type 'how much is that doggie in the window' rather than 'dog AND window AND (cost OR price)'.

natural language learning – machine learning applied to natural language corpora to extract computable language models.

neural network – an important area of Artificial Intelligence research, neural networks are statistical pattern-recognition systems built to emulate biological and neural systems.

Optical Character Recognition (OCR) – recognition of printed language using a symbolic representation derived from its graphical form: for alphabetic languages, this means recognising and transforming characters. OCR for a document using a single printed font can achieve a very high degree of accuracy, although problems can arise with unknown or decorative fonts or poor print quality. In these cases, and in the case of handwriting, good results can be achieved only by using Intelligent Character Recognition.

parse – to analyse language in order to establish its syntactic and/or semantic structure and relationships.

semantics – the study of meaning, and of the principles that govern the relationship between words and sentences and their meanings.

speaker independent – a speech recognition system that can recognise speech regardless of the speaker and does not need to be trained to recognise individuals.

speech generation – speech is generated from templates, either by playing recordings or joining units of speech (phonemes, words) together. To produce a continuous and realistic response, generated speech can handle aspects such as intensity, duration and stress. Dialogue, for example in the automated handling of telephone calls, can be established by combining speech recognition with speech generation.

speech recognition – computerised analysis of sounds of speech to identify units of sound (or phonemes) that make up words. Statistical models, derived from corpora (qv), are used to recognise discrete or continuous speech input. Other significant problems to be overcome if speech is to become a commonly-used medium for interacting with computers include the ability to recognise multiple speakers; eliminating noise; and dealing with accents, dialects, and the ungrammatical nature of much spoken language.

syntax – the study of the grammatical arrangement of words and morphemes in the sentences of a language or of languages in general.

tag – to annotate corpora by attaching information to words, to describe their grammatical context and/or associations with other words.

telematics – the application of informatics and/or telecommunications; the EU Telematics Applications programme is user-driven and focuses on the societal implications of information and communication technologies.

unicode: a character coding system that includes codes for non-Roman alphabets (compare with Ascii).

Corpora? Corpora are collections of texts and can be wide-ranging or narrowly focused; providing a 'bank' of language resources or specific examples.

Corpora are particularly useful in language machine development. Think of the words people might need to use, for example, a dishwasher, washing machine or vacuum cleaner – a corpus containing these context-specific words would be useful in creating domestic robots.

For a brief outline of corpora, see pages 26 & 27

Books? Magazines?
See pages 60 & 61

references

the books

- The Future of English?* (Graddol, 1997) published by the English 2000 project of The British Council, was commissioned to facilitate informed debate about the future use and learning of the English language worldwide. See: <http://www.britishcouncil.org> <http://www.english.co.uk>
- Interest in telematics?* For overviews of telematics programmes see Linglink, 1997; European Commission, 1998; EUROMAP, 1999. See pages 62 & 63 for information on Web sites.
- Thanks to The Nuffield Languages Inquiry Online for the responses from Neil Langdon Inglis, John Thorogood and Graham Davies on page 37. The Nuffield Languages Inquiry has been set up to estimate the UK's needs for languages in the next 20 years. For information about the Inquiry, see: <http://www.nuffield.org>
- NaturallySpeaking; ViaVoice; and Lego* are trademarks of Dragon Systems, IBM and Lego respectively.
- Atwell, E. (1993) *Knowledge at Work in Universities: the Higher Educational Funding Council's knowledge-based systems initiative*. Leeds: Leeds University Press.
- Atwell, E. (1996) Machine Learning from Corpus Resources for Speech And Handwriting Recognition. In J. Thomas and M. Short (eds) *Using Corpora for Language Research: Studies in honour of Geoffrey Leech*. Harlow: Longman.
- Atwell, E. (1997) *The World of Language: speech and language technology. Preliminary study for the British Council's World of Language project*.
- Atwell, E. (1998) What can SALT offer the English teaching professional? In *English Teaching Professional*, 7:46-47.
- Atwell, E., Herron, D., Howarth, P., Morton, R. and Wick, H. (1999) *Pronunciation Training: Requirements and Solutions*. ISLE project report D1.4, University of Leeds.
- Balkan, L., Arnold, D. and Sadler, S. (1997) *Tools and Techniques for Machine Translation Teaching: a survey*. Report to ELSNET, University of Essex.
- Bod, R. (1998) Spoken Dialogue Interpretation with the DOP Model. *Proceedings ACL-COLING '98*. Montreal, Canada.
- Bryson, B. (1994) *Made in America*. New York: Martin Secker and Warburg.
- BT Education Services (1998) *Factfile Number 11: Telecommunications 20/20 Vision*. British Telecom.
- Charniak, E. (1993) *Statistical Language Learning*. Cambridge, Massachusetts: Bradford / MIT Press.
- Churcher, G., Atwell, E. and Souter, C. (1997) Developing a Corpus-based Grammar Model within a Commercial Continuous Speech Recognition Package in WEB-SLS. *The European Student Journal of Language and Speech*, 97,03.
- Cochrane, P. (1997) *Tips for Time Travellers*. London: Orion Business Books.
- Cole, R., Mariani, J., Uszkoreit, H., Zaenen, Z. and Zue, V. (eds) (1995) *Survey of the State of the Art in Human Language Technology*. Report to US National Science Foundation and European Commission.
- Dertouzos, M. (1997) *What Will Be*. New York: HarperCollins Publishers Inc.
- Eastment, D. (1999) *The Internet and ELT: the impact of the Internet on English Language Teaching*. Oxford: Summertown Publishing.
- EAGLES (Expert Advisory Group on Language Engineering Standards). (1995) *Evaluation of Natural Language Processing Systems: Final Report*. EAGLES Document EAG-EWG-PR.2.
- Elliott, J. and Atwell, E. (1999) Is anybody out there?: the detection of intelligent and generic language-like features. *Journal of the British Interplanetary Society*, August 1999.
- EPSRC (1996) *Research Progress 1995-1996: Human Factors, IT and Computer Science*. Swindon: Engineering and Physical Sciences Research Council.
- EPSRC (1998) *Devising a Strategy for the Support of Human Factors*. Published online: <http://www.epsrc.ac.uk/documents/programmes/highlights/itcs/hfstrategypaper.htm>
- EUROMAP (1998) *The EUROMAP Report: challenge and opportunity for Europe's Information Society*. Luxembourg: Linglink
- European Commission DGXIII (1998) Important Issues in Today's Telematics Research. *Telematics Applications Conference Proceedings*. Barcelona.
- Flanagan, M. (1997) Emerging Roles, New Successes in Machine Translation, Special Issue on New Tools for Human Translators. *MT Today* 12,1-2: 25-27.
- Graddol, D. (1997) *The Future of English*. London: British Council.
- Haskel, L. (1999) The Accidental Artist. *Mute*, 13:71-73.
- Howard, T. (1998) Breeding like robots. *Personal Computer World*, October 1998, p.238.
- Hunt, M. (1998) Practical Automatic Dictation Systems. *ELRA Newsletter*, 3,1:4-7.
- Hutchins, J. (1996) Computer-based translation systems and tools. In *ELRA Newsletter*, 1,4:6-9.
- Hutchins, J. (1998) Translation Technology and the Translator. *Machine*

- Translation Review*, 7:7-14.
- Johnson, S.R. (1997) *Interface Culture: how new technology transforms the way we create and communicate*. New York: HarperEdge.
- Johnson, T. (1985) *Natural Language Computing: the commercial applications*. London: Ovum.
- Kay, M. (1980) *The Proper Place of Men and Machines in Language Translation*.
- Kurzweil, R (1998) *The Age of Spiritual Machines: when computers exceed human intelligence*. New York: Viking Penguin.
- Linglink (1997) *Language Engineering: a greater use of language in every sense*. Project Report to the European Union.
- Milner, R. and Wand, I. (eds) (1996) *Computing Tomorrow*. Cambridge: Cambridge University Press.
- Negroponte, N. (1995) *Being Digital*. New York: Alfred A Knopf.
- O'Donnell, J. J. (1998) *Avatars of the Word: from papyrus to cyberspace*. Harvard: Harvard University Press.
- Pearson, I. (1997) *BT Technology Calendar* (1997-2045). Published online: http://www.labs.bt.com/library/archive/btl_97_1/index.htm
- Pringle, D. (1998) Search warrant? The web view on the role of language processing. *LE Journal*, article #356. Published online: <http://www.linglink.lu/lejournal/article.asp?articleIndex=356>
- Reddy, R. (1995) Grand Challenges in Artificial Intelligence. *ACM Computing Surveys* (Special Issue on Artificial Intelligence), 7,3:301-303.
- Rothwell, A. (1998) Translating and the PC: the MAATS experience. *University of Leeds Computing Service Newsletter*, 228:10-11.
- Schillo, M., Atwell, A., Souter, C. and Denson, A. (1996) Language modelling for the in-car intelligent personal assistant. In C. Moghrabi (ed) *Proceedings of NLP+IA96 : International Conference on Natural Language Processing and Industrial Applications*.
- Souter, C. and Atwell, E. (1993) Preface in C. Souter and E. Atwell (eds) *Corpus-Based Computational Linguistics*. Rodopi: Amsterdam.
- SRI (1998) *Information Extraction with Highlight*. Cambridge Computer Science Research Centre. http://www.cam.sri.com/html/info_extraction.html
- Turing, A. (1950) Computing Machinery and Intelligence. *Mind*, 59:433-560.
- Tyler, G. (1998) The last word. *Management Services*, May 1998, 30-33.
- Varile, G. (1997) *Human Language Technologies: living and working together in the information society*. Discussion Document. Luxembourg: European Commission DG-XIII.
- Whyte, W. (1999) *Networked Futures: power and intelligence in the future*. Cambridge: J.Wiley and Sons.
- Wieners, B. and Pescovitz, D. (1996) *Reality Check*. San Francisco: Hardwired.

Magazines: the quickest route

Information on language machine developments often appears in popular computing magazines. The UK editions of *Computer Active* (July 1998), *PC Direct* (June 1998), and *Personal Computer World* (October 1998) all reviewed PC speech recognition systems. *Computer Active* recommended IBM ViaVoice Gold as the best (at the time). *PC Direct* did not nominate a favourite, but reviewed each package in detail. *Personal Computer World* preferred *Dragon Systems Naturally-Speaking*, and provided a useful summary of features to look for in voice recognition software.

Many magazines also include free CDs with sample software. The August 1998 *PC Plus* free CD included *Globalink's English-Spanish Power Translator Deluxe* (*Globalink* 1998); *Computer Buyer* May 1998 free CD included a trial version of *Dragon Naturally-Speaking* (*Dragon* 1998).

Reviews of language machines are also now appearing in magazines beyond the PC shelf, as specialist magazines recognise this technology is relevant to their readers: *Management Services* (May 1998) reviewed a number of speech processing systems for managers (Tyler, 1998).

Web information

cyber-sites

EU Telematics Applications programmes

General information: <http://www.echo.lu/telematics/>

Information Services Projects

EUROSEARCH (Multilingual European Federated Search Service):

<http://www.linglink.lu/le/projects/access>

MAY (Multilingual Access to Yellow Pages): <http://www.linglink.lu/le/projects/may>

MULINEX (Multilingual Indexing Navigation and Editing Extensions for the World Wide Web):

<http://www.linglink.lu/le/projects/mulindex>

Education and Training Projects

ISLE (Interactive Spoken Language Education): <http://www.linglink.lu/le/projects/isle>

LETRAC (Language Engineering for Translator Curricula): <http://www.linglink.lu/le/projects/letrac>

RECALL (Repairing Errors in Computer-Aided Language Learning):

<http://www.linglink.lu/le/projects/recall>

SELECT (Strategies for European LE-Enhanced Communication Training):

<http://www.linglink.lu/le/projects/select>

SPEAK (Supported Prototype Easy-Access Authoring Keys): <http://www.linglink.lu/le/projects/speak>

Public Interest Projects

AVENTINUS (Advanced Information System for Multilingual Drug Enforcement):

<http://www.linglink.lu/le/projects/aventinus>

LINGUANET (Communicating Through the Language Barrier):

<http://www.linglink.lu/le/projects/linguanet>

MIETTA (Multilingual Information Extraction for Tourism and Travel Assistance):

<http://www.linglink.lu/le/projects/mietta>

SENSUS (LE for police and emergency service communications and information systems):

<http://www.linglink.lu/le/projects/sensus>

International Business Support Projects

DOCSTEP (Product DOCumentation Creation and Management using STEP):

<http://www.linglink.lu/le/projects/docstep>

MABLE (Multilingual Authoring of Business Letters): <http://www.linglink.lu/le/projects/mable>

OTELO (Common Access to Translation Systems): <http://www.linglink.lu/le/projects/otelo>

TRANSROUTER (Translation Router): <http://www.linglink.lu/le/projects/transrouter>

Language Resources Projects

- EUROWORDNET (Building a multilingual wordnet database with semantic relations between words): <http://www.linglink.lu/le/projects/eurowordnet>
- PAROLE (Preparatory Action for linguistic Resources Organisation for Language Engineering): <http://www.linglink.lu/le/projects/le-parole>
- SIMPLE (Semantic Information for Multifunctional Plurilingual Lexica): <http://www.linglink.lu/le/projects/simple>
- SPEECHDAT (Speech Databases for Creation of Voice Driven Teleservices): <http://www.linglink.lu/le/projects/speechdat>

Infrastructure and Support Projects

- DIET (Diagnostic and Evaluation Tools for Natural Language Applications): <http://www.linglink.lu/le/projects/diet>
- EAGLES Expert Advisory Group on Language Engineering Standards): <http://www.linglink.lu/le/projects/eagles>
- ELRA (European Language Resources Association): <http://www.linglink.lu/le/projects/elra>
- ELSE (Evaluation in Language and Speech Engineering): <http://www.linglink.lu/le/projects/else/>
- ELSNET (European Speech Language and Speech Network of excellence): www.elsnet.org
- EUROMAP/EUROMAP II (European Opportunity Mapping): <http://www.linglink.lu/le/projects/euromap/>
- LINGLINK (Promotion and Support for Language Engineering in Europe): <http://www.linglink.lu/le/projects/linglink/>
- MATE (Multi-Level Annotation Tools Engineering): <http://www.linglink.lu/le/projects/mate/>

News, Technology Culture

- CNET – The Computer Network: <http://www.cnet.com>
- Degrees Feet and Inches: <http://www.dfi.org.uk>
- Feed magazine: <http://www.feedmag.com>
- Hotwired: <http://www.wired.com>
- Nua: <http://www.nua.ie>
- Rhizome Digest: <http://www.rhizome.org>
- Salon magazine: <http://www.salon.com>
- The Alan Turing Home Page: <http://www.turing.org.uk/turing/>
- Bletchley Park Trust: <http://www.bletchleypark.org.uk>
- The Loebner Prize: <http://www.loebner.net/Prizef/loebner-prize.html>
- ZDTV: <http://www2.zdnet.com>

Web information

... continued

Non-EU Projects, Organisations and Resources

AMALGAM (Automatic Mapping Among Lexico-Grammatical Annotation Models):
<http://www.scs.leeds.ac.uk/amalgam/amalgam/amalghome.htm>
 British Council: <http://www.britishcouncil.org>
 British National Corpus: <http://info.ox.ac.uk:80/bnc/>
 CAPITAL SIG (CALICO Computer Assisted Pronunciation Investigation Teaching and Learning Special Interest Group): http://tay.ac.uk/schools/sch_man/div_lang/capital/capital.html
 Centre for Speech Technology Research: <http://www.cstr.ed.ac.uk>
 Department of Trade and Industry Innovation Unit: <http://www.innovation.gov.uk/home.htm>
 Engineering and Physical Sciences Research Council (EPSRC): <http://www.epsrc.ac.uk>
 EUROCALL (European Association for Computer Assisted Language Learning):
<http://www.hull.ac.uk/cti/eurocall.htm>
 European Student Journal of Language and Speech: <http://web-sls.essex.ac.uk/web-sls/>
 GALEN (Generalised Architecture for Languages, Encyclopaedias and Nomenclatures in Medicine):
<http://www.galen-organisation.com>
 ICAME (International Computer Archive of Modern and Medieval English):
<http://www.hd.uib.no/icame.html>
 Information Communications and Technology for Language Teachers: <http://www.ict4lt.org>
 Institute of Translators and Interpreters: <http://www.iti.org.uk>
 Interactive Syntax Tutorial: <http://www.bangor.ac.uk/ling/java/lt/LingTutor.html>
 Internet Grammar of English: <http://www.ucl.ac.uk/internet-grammar/home.html>
 Language Learning and Technology (Internet journal): <http://polyglot.cal.msu.edu/llt>
 Leeds Electronic Text Centre: <http://www.leeds.ac.uk/acom/leedsetc/infostaf.htm>
 Speech and Language Technology (SALT) Club UK: <http://walt.essex.ac.uk/salt>
 SLIM (Multimedia Interactive Linguistic Software): <http://byron.cgm.unive.it>
 Wearable Computing, Essex University: <http://wearables.essex.ac.uk>
 World of Language: <http://www.worldoflanguage.com>

The ISLE Project?
 Referenced
 throughout this book.
 See pages
 17, 21, 29, 30,
 51

The ISLE Project aims to adapt and integrate speech technology in multimedia courseware environments: assessments are made to determine in what ways speech technology is helpful for learners who are trying to acquire foreign language skills. Many teachers and students are already involved in the project.
 See <http://www.linglink.lu/le/projects/isle>

Supply Information and Online Demonstrations

Altavista babelfish translation: <http://www.babelfish.altavista.digital.com/>
Aurolog TaLk to Me: <http://www.aurolog.com/eng/eng/alt/index.htm>
BT Laureate speech synthesiser demo: <http://innovate.bt.com/showcase/laureate/index.htm>
British Telecom SuperQuill: <http://innovate.bt.com/showcase/smartquill/index.htm>
Conter Software (1998) Storybook Maker: see <http://www.indelibleink.com/kded023.html>
Dragon Systems (1998) Dragon NaturallySpeaking (demo version). Published on free CD with Computer Buyer, May 1998
Dragon Systems Speech Recognition Products:
<http://www.dragonsys.com/frameset/product-frame.html>
Globalink (1998) Power Translator Deluxe, English-Spanish. Published on free CD with PC Plus, August 1998
IBM Home Page Reader: <http://www.austin.ibm.com/sns/hpr.html>
IBM Speech Recognition: <http://www.software.ibm.com/speech/index.html>
IBM ViaVoice98 – UK English: http://www.software.ibm.com/is/voicetype/uk_home98.html
Systran Professional Translation Software: <http://www.systran.com>

Disability

ISAEUS Project on Speech Therapy and Language Technology:
<http://www.echo.lu/telematics/disabl/isaesus.html>
Masoud Yazdani, work on computer-based iconic language: <http://www.media.uwe.ac.uk/~masoud/>
Royal National Institute for the Blind: <http://www.rnib.org.uk>

Future Technology

BT Technology Calendar: <http://www.labs.bt.com/library/on-line/calendar>
Philips Vision of the Future: <http://www.design.philips.com/vof/tocl/home.htm>
UK Futures Group: <http://www.futures.org.uk>

full contents

click to go

Preface	5
<i>Progress and Prospects</i>	5
The language machine	6
Surveying the territory	6
The impact?	6
<i>What's in a name?</i>	7
Describing words: the technology	8
What makes a language machine?	8
<i>Is it a language machine?</i>	8
<i>... So is it SALT?</i>	9
Telling histories: 1791–2020	10
Dreaming of speaking machines	10
<i>From Turing's child-machine to a</i>	
<i>25 thousand dollar race</i>	10
<i>When destiny calls, have the confidence to answer</i>	11
Seeing futures emerging	12
To 2045: the BT Technology Calendar	12
A widespread expectation	12
* the wildcards?	13
Research labs to desk-tops	14
An academic heritage	14
Where languages meet computer science	14
Creating popular markets	16
A use for everyone?	16
<i>The language of language processing</i>	16
Turning theory into practice	18
Why create language machines?	18
What's involved: input to output	20
Defining shared approaches	20
Listen, analyse, understand	22

<i>Medieval style or machine stylus?</i>	22
talk back and translate	24
<i>The European dilemma ...</i>	25
<i>... with no clear solution?</i>	25
Knowledge communication	26
<i>Language resources: hundreds, thousands, millions – and growing</i>	26
<i>Resources for developing the language machine</i>	27
The problems with machines	28
What slows development?	28
Problem 1: the language machine is expensive	28
Inappropriate technology	30
Problem 2: the language machine has been developed without consultation	30
Problem 3: the language machine doesn't match the expectations of customers	30
Problem 4: the language machine is not appropriate for some tasks	31
<i>I will compare the thee with the day of an été?</i>	31
Human-machine mismatched	32
Problem 5: to use the language machine successfully, we need to rethink how we approach tasks	32
<i>Know less, understand more</i>	32
Problem 6: newcomers need training and time to learn to use the language machine	33
<i>Expensive and useless?</i>	34
Problem 7: applications face all the above problems	34
<i>Tower of Babel? Or Cubicle of Babble?</i>	34
<i>The problem of the popular: Hal 9000 Series</i>	35
Determination to succeed	36
Can problems be overcome?	36

full contents

... continued

<i>Machine translation? Speech recognition? Do they work?</i>	37
UK advantage? The key-players	38
A global interest	38
What's happening in the UK?	39
<i>What do EU projects aim to do?</i>	39
UK-Europe overview	40
UK involvement across Europe	40
<i>Case Study: Lernout and Hauspie</i>	40
For translation, security, medicine ...	41
<i>The ten EU rules for language machines in telematics</i>	41
Supply & demand: the markets	42
How are people encouraged to buy?	42
<i>What are the benefits to the UK?</i>	43
Real-life situations	44
Is there a typical 'user profile'?	44
Labs, factories, outdoors, offices, homes	45
<i>How are people using language machines?</i>	45
Employment, society, leisure	46
<i>The art of translation. Or the translation of art?</i>	47
Educational resources	48
Focus: language education & teaching	48
<i>Have a go? Language technology by email</i>	48
Educational projects in progress	49
Future scenarios: the impact	50
The next 20 years	50
Scenario 1: IT and language machines will be all-pervasive	50
Scenario 2: the language machine will get better	50

Scenario 3: customised language machines will be used by individuals for particular tasks	51
Implications: the professions	52
Scenario 4: no aspect of language professions is immune	52
<i>Rethinking our assumptions – and our self-image?</i>	52
<i>Scenario 1: the must-haves – handbag, lippy, wristwatch translator</i>	53
Questioning the future	54
Having misgivings?	54
<i>The digital helping hand?</i>	54
<i>Good times and fortunes</i>	55
Leading the way: the next steps	56
Ways ahead	56
<i>Wild child – or wired child?</i>	57
Glossary: the jargon	58
References: the books	60
<i>Magazines: the quickest route</i>	61
Web information: cyber-sites	62
Questions for debate	70

questions

for debate

How will the language machine affect the language-based professions?

How can computer scientists and language professionals collaborate to mutual benefit?

How will practitioners find ways of influencing and informing – and of staying informed of – developments in technology and best practice? If discussion is needed, how will it be facilitated?

Will we stop learning languages? And, if we do, what will be the cultural and social impact?

Will endangered languages be saved by a lesser need to learn a lingua franca, or will software not be developed for languages with smaller, less profitable audiences?

Not all languages are similarly alphabetic, phonetic or tonal. Will some languages benefit more than others?

How equally will we benefit? Linguistic minorities and the hearing and sight-impaired stand to gain from these technologies, but will hardware and software costs or technophobia become new sources of division and inequality?

How will technology change our interactions? How will we adapt once technology shapes – and speaks – the actual words we send and receive? Will speech recognition re-emphasise oral skills?

How will a greater transparency of information in other languages affect economies?

Will customer services be automated further, with personal interaction becoming a premium service in the 21st century?

Will the UK be seen as a leader of language machine technology?

This digital book has been developed in PDF (Portable Document Format) for use with Adobe's Acrobat Reader. For general advice on using and configuring Adobe Acrobat – including how to set a Web browser to activate hyperlinks contained within this document – please refer to your Acrobat documentation.

What is a digital book?

A digital book is very like a conventional printed book, with the same content and page layout. The digital version, however, provides many benefits. You can, for example, move instantly to any location you want; navigate your way through the text using on-screen toolbars, in-built links and thumbnails; access Web sites simply by clicking on their addresses (provided you are already online); print hardcopies when you need; set up your own viewing preferences; – and use all these facilities across a range of computer systems.

Using hyperlinks

To make this book as useful as possible, this digital text contains 'hyperlinks'. You can see a hyperlink is present when the on-screen mouse pointer changes from an open hand to an index finger. Simply click on your mouse and you are taken immediately to the selected page.

You'll also find that addresses of Web sites in this document are 'active': when you click on a Web site link in the book you'll be taken directly to that site.

Remember that to use these links you need to tell your Acrobat Reader which browser you want to use and the directory where the browser application is stored. (If you need help, follow the instructions in the online guide of your Acrobat Reader.)

Hyperlink features in this digital book:

-  all references to Web sites, are active – simply click the link to activate your browser and access the site in question;
-  all page references in the *contents*, *full contents* and *glossary* sections are also active links: go to the page you want by clicking the link highlighted;
-  each main section of the book has been created as an article. Click on the main text to 'zoom' your window to the column width of the current article; simply continue clicking to follow
-  each section of the article in this expanded view (Acrobat Reader can display a list of the articles: please refer to your own documentation);
-  you can also navigate through this text using *Thumbnails*: refer to the documentation accompanying your copy of the Acrobat Reader, as navigation methods can vary between versions and operating systems.

You can print out a single page, or multiple pages, of this document. (If you do so, please ensure you conform to the licence conditions.) If you have the full Acrobat software running on your computer, then you can annotate this document to customise it for your own use. For example, you can add your favourite Web site addresses, make notes in the margins, or highlight text in the same way as you would with a marker pen on conventional paper-based books.

We suggest you explore the many different ways of viewing and using this digital book, to find those which are convenient on your own machine.

licence agreement

SITE LICENCE

Notice to user

This agreement is a contract between you and the publisher which stipulates how you may use this Electronic Document. By indicating your acceptance, you accept all the terms and conditions contained in this agreement.

Use of this Electronic Document

1. This Electronic Document is licensed for use on any number of computers on a single site. You are not entitled to loan, rent, lease or sell the document to others whose use is not explicitly covered by the terms of this licence.
2. This licence permits access to, distribution or use of the document over any networks to which the computer or computers may be connected (either at the time of installation or subsequently).
3. The purchaser is entitled to make a single back-up copy of the document for your own use.
4. Under the terms of this licence, you are entitled to print single copies or part copies for your own personal use. This entitlement does not cover printing of multiple copies, or the photocopying of copies that you have printed for your own use.
5. You may not show or display all or part of the document in public (for example, as part of a presentation).
6. Any fonts embedded in this document may not be extracted or disembedded and used for other purposes.
7. Please refer to the statement of copyright included in the document. The copyright holder reserves all rights. No part of the document may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, except as provided for in this agreement, without the prior permission of the copyright holder.
8. Limited Liability. The vendor, publisher and copyright holder exclude liability for any loss or damage caused to you as a result of this purchase or subsequent use of this electronic document.
9. This agreement governing the use of this electronic document is in addition to any contract of sale or warranty provided by the vendor.

This agreement will be governed by the laws of England.

Please read this agreement carefully. To accept the agreement and use this digital book simply click the button at the bottom of the page: [READ DOCUMENT](#).

If you click the [QUIT DOCUMENT](#) button, you will close the book and exit from Acrobat Reader.

READ DOCUMENT

QUIT DOCUMENT