

H. H. Rosenbrock (Ed.)

Designing Human-centred Technology

A Cross-disciplinary Project in
Computer-aided Manufacturing

设计以人为中心的技术 [英]

(B1)



Springer-Verlag

World Publishing Corp

H. H. Rosenbrock (Ed.)

Designing Human-centred Technology

A Cross-disciplinary Project in
Computer-aided Manufacturing

With 18 Figures



Springer-Verlag

World Publishing Corp

Howard Rosenbrock
Linden, Walford Road, Ross-on-Wye,
Herefordshire HR9 5PQ

ISBN 3-540-19567-X Springer-Verlag Berlin Heidelberg New York
ISBN 0-387-19567-X Springer-Verlag New York Berlin Heidelberg
ISBN 7 - 5062 - 1111 - 4

British Library Cataloguing in Publication Data

Rosenbrock, H. H. (Howard H, 1920-)

Designing human-centred technology.

1. Manufacture. Applications of computer systems

I. Title

670.42'7

ISBN 3-540-19567-X

Library of Congress Cataloging-in-Publication Data

Designing human-centred technology: a cross-disciplinary project in
computer-aided manufacturing/H. H. Rosenbrock, ed.

p. cm.—(The Springer series on Artificial Intelligence and Society)

Includes bibliographical references

ISBN 0-387-19567-X

1. Computer integrated manufacturing systems 2. Human engineering.

I. Rosenbrock, H. H. II. Series

TS 155.6.D474 1989

670'.285--dc20

89-21625

CIP

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

© Springer-Verlag London Limited 1989

The use of registered names, trademarks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant laws and regulations and therefore free for general use.

Reprinted by World Publishing Corporation, Beijing, 1991
for distribution and sale in The People's Republic of China only

Foreword

This second book in our series *Artificial Intelligence and Society* explores the issues involved in the design and application of human-centred systems in the manufacturing area. At first glance it may appear that a book on this topic is somewhat peripheral to the main concerns of the series. In fact, although starting from an engineering perspective, the book addresses some of the pivotal issues confronting those who apply new technology in general and artificial intelligence (AI) systems in particular. Above all, the book invites us to consider whether the present applications of technology are such as to make the best use of human skill and ingenuity and at the same time provide for realistic and economically sustainable systems design solutions. To do so it is necessary to provide systems which support the skill, and are amenable to the cultures, of the areas of application in question. In a philosophical sense it means providing tools to support skills rather than machines which replace them, to use Heidegger's distinction.

The book gives an authoritative account of the University of Manchester Institute of Science and Technology (UMIST) tradition of human-centredness and provides a participatory design approach which focuses on collaborative learning and enhancement and creation of new skills. It also argues that collaboration should be supported by institutions through the creation of supportive infrastructures and research environments. It emphasises the optimisation of practical knowledge with the help of scientific knowledge and rejects the alternative. The 'blank table debate' in the book shows how personal opinions, professional experiences, academic and practical perspectives of a multi-disciplinary team of designers can work together creatively, resolve conflicts through a constructive divergence and arrive at practical systems which are of benefit to the users. The book records and analyses the reasons for the research conflicts which are believed to be a product of the institutional splitting which an interdisciplinary project faces.

The book raises, both directly and indirectly, issues of central concern to the present series. Among the series' objectives are:

it will provide an authoritative treatment of the nature of knowledge, skill and practice and the nature of AI technology, and will examine issues

of the design, application and implications of AI systems within their broader context. The series will focus on the philosophical, social, economic and political contexts of AI development and will cover topics such as the nature of expert knowledge and the problems of its computerisation; AI technology and the transfer of knowledge and skills; human/computer integration in the management of industrial and military systems; social and cultural shaping of AI technology and the nature of work; knowledge-based society and the issues of control, responsibility, access and ethics. (*AI & Society* pamphlet.)

Given the growing significance of computerisation in all aspects of human endeavour, and the anticipated pervasive nature of AI systems, it behoves each professional group to explore creatively the issues and implications of technological change in their own areas. In the field of manufacturing technology it is becoming clear that many sections of industry are already precariously dependent upon machine-centred systems. Such systems are typically highly synchronised and coordinated, but frequently lack robustness, in the sense that they are inadequate in dealing with disturbances and uncertainties. Thus, if one part of the system goes down, the high level of synchronisation is suddenly turned into its opposite and becomes a high level of desynchronisation. If, however, the system is human-centred as suggested in this book, it is contended that this will provide for good human-machine symbiosis, in which there will be involved pro-active, creative human beings for whom the system is transparent, and who will be capable of acting in an informed way in the event of uncertainty.

We are at an interesting juncture in the development of AI which could prove to be a turning point for the design of intelligent systems. Dominant AI is rapidly finding its place in software engineering or cognitive science which, by nature and by adoption, have a machine-centred focus exemplified by the expert systems and laboratory-based human-computer interaction research. This book provides an alternative to the computational metaphor of AI and recognises the significance of the tacit dimension of human knowledge, and demonstrates the significance of working-life experiences and practices for the design of purposeful computer systems. In this context it is consonant with the increasing interest of industry, commerce and organisations which are beginning to recognise the limitations of machine-centred technology to assist in qualitative decision-making processes.

The introduction of AI systems into any field of skill or competence must reflect in part the nature of existing practices and technology. If the practices which precede the introduction of the AI systems are such as to diminish the role of the human beings and in many cases reduce them to abject machine appendages, it will tend to follow that AI systems will reflect these values, and an opportunity for a more creative and ultimately productive system will have been lost. The systems will tend to be "expert replace-

ment systems" rather than "expert systems". Present systems are introduced within a scientific tradition in which it is held that a system is only scientifically designed if it displays the three predominant characteristics of Western science and technology: predictability, repeatability and mathematical quantifiability. That by definition precludes intuition, subjective judgement and tacit knowledge. A human-centred approach questions, at a philosophical level, the wisdom of such an approach, and suggests rather one in which the best use is made of the human being and the system, paying due regard to the characteristics of both.

This book clarifies the debate on the tacit dimension of skills and practice through the UMIST project and thus lays a foundation for the design of practical systems in working-life environments where the fundamental issues of collaborative design will be common. It will be a challenge to AI researchers who may wish to transfer these methodologies to their own domains such as collaborative learning, knowledge transfer, office automation, human-computer interaction, decision support, natural language processing. The book will be of special interest to researchers and practitioners who are involved in the design of collaborative knowledge-based networks, multi-media learning systems, and collaborative training systems. The design challenge may be seen in dealing with the issues such as: sharing of knowledge space; sharing common experiences, distribution and diffusion of knowledge; open and distance learning; tacit knowledge; social and cultural determinants; human factors; information flows and knowledge flows; internationalisation of knowledge - transfer of knowledge across cultures; and mediation of knowledge. This book provides concepts, approaches and methodologies which may be transferred to the design of AI systems in the above areas and in dealing with the above issues. It shows how the UMIST tradition of human-machine symbiosis has evolved a methodology for human-centred design during the process of the project. This methodology became an important part of the ESPRIT project 1217. It will be reasonable to suggest that the concept of human-machine symbiosis can be applied equally to develop participatory design methodologies for domains other than manufacturing.

This book concludes with a brief outline of ESPRIT project 1217 - Human-centred CIM, which was greatly influenced by the UMIST research. Other significant influences included the social shaping philosophy of the German partners and the end user involvement techniques of the Danish partners. A subsequent book in this series will explore these influences and will describe the ESPRIT project in greater detail.

A great strength of the book is that it raises at a practical and a philosophical level sets of alternatives to the given orthodoxy. As such it challenges and invites other professions, such as those of law, medicine and banking, likewise to explore the underlying values of the systems being introduced in their areas, and to

question whether the new technologies, including expert systems and AI, might not be introduced in such a way as to enhance human skill rather than diminish it.

*Mike Cooley
Karamjit S. Gill*

Preface

The Industrial Revolution of the late eighteenth and early nineteenth centuries took place under social conditions very different from our own. One of its legacies is a technology relying upon assumptions which are strongly at variance with the needs and aspirations of men and women today. Production systems are designed with attention fixed upon the machines: their needs and their effective use are the main considerations. Men and women have to fit as best they can into the systems that result. Often they are subordinated to the machines, which determine human actions and rate of working. A specifically human input – by initiative, skill, control, and the response to unexpected events – is rejected.

This situation has been strongly criticised by social scientists for its damaging effects upon workers. It can equally be criticised for its wastefulness of a primary resource. It makes excellent use of machines, but very poor use of people. Systems designed to reject a human input, or to draw this only from a small centralised nucleus of people are rigid and inflexible. They cannot respond to unforeseen events or to rapidly changing market demands. So much is now becoming acknowledged.

If we wish to change the direction of technological development, so that the human input is accepted and valued, we shall have to intervene at the stage where new technology is being designed. The project at the University of Manchester Institute of Science and Technology (UMIST) which is described in this book was an attempt to do so in relation to new production systems which go under various names, but are usually known collectively as computer-integrated manufacturing (CIM). The technology, still fluid and developing, sets out to control and integrate the whole production process by means of a network of computers. As usually envisaged, the aim is to replace every human contribution by automation, with the ultimate goal of a 'workerless factory'. The dream is an old one, and it is well known that the result is not a factory without workers, but a factory in which workers must behave like machines.

In the project a different aim was proposed: computers and automation would not be used to reject human ability and skill,

but rather to cooperate with them to make them more productive. People should not be subordinate to machines; machines should be subordinate to people. Existing skills should be accepted, and room for their exercise should be provided. But room also should be provided for skills to change as technology changes: the evolution of skill was regarded as open-ended and unlimited, where the conventional view sees past skills as finite and doomed to extinction as they are replaced by automation. It was believed that, by following the proposed aim, a system could be designed that was at least as economic as conventional systems, but more flexible, more responsive, and better matched to the aspirations of workers.

One of the major obstacles to the proposed development is the burden of preconceptions brought to the design process by technologists – engineers, computer scientists and others. Their training and experience tend very strongly to concentrate their attention on machines, rather than people. On the other hand, there is a strong tradition in the social sciences of ‘socio-technical design’, in which the technology and the conditions of work are both studied together in the design stage. Unfortunately, opportunities to do this have usually been limited to small changes at the fringe of technology, rather than the deep reconsideration which is needed. A way of describing the aims of the project is that it sets out, by means of a collaboration between social scientists and technologists, to apply socio-technical design principles at a deep level in technology.

Such an aim poses great problems to those concerned, both technologists and social scientists. The technologists have no experience of designing with the human requirements in view. The social scientists have no deep understanding of the technology, or of the technological design process. So engineers and computer scientists cannot see how to balance human requirements against those of the technology. Social scientists cannot see when a stated technical requirement represents an engineering constraint, and when it is simply a preconception. On both sides it is necessary to acquire the knowledge and practical ability needed to carry out a new kind of design leading to a new kind of technology.

Equilibrium could probably have been achieved in two ways: with social scientists acting as advisers to technologists, but at arm’s length from the design process (a not uncommon arrangement), or conversely with social scientists controlling the project and making the major decisions. Neither extreme was thought likely to be fruitful. Instead a tension between the two sides was sought, out of which it was hoped that a joint understanding could arise. The account which follows shows what happened in this process, as seen from different points of view. Disagreements of opinion are as likely to be informative as agreements, and they have not been disguised: the account has the nature of ‘Faithful Contendings Displayed’.

With this aim in mind, no attempt has been made to achieve uniformity in the style or approach of the different contributions. To change an author's style is to make a subtle change in the impression conveyed, and to that extent to falsify the account. So, for example, the different practice of authors in regard to the masculine bias in the English language has been left unchanged. The majority have adhered to the traditional convention that 'he' includes 'she', whereas in other places 'he or she' or some other form has been used.

Some may conclude from a reading of what follows that my own role, as originator and grant-holder for the project, called for a more decisive direction of the work. Certainly this would have resulted in greater progress in the technological design, but it would have left many things less clearly defined – as agreements or disagreements or as open questions – than they will be found to be here. My aim was to elicit the views of others; and to state my own clearly and forcibly, but to refrain either from imposing them or from abandoning them if I was unconvinced. The reader should therefore find much material here for study and further development.

What is most to be desired is that the small beginnings in this project should be followed by a multitude of others of increasing depth and richness. The ESPRIT project, described in Chapter 11, carries on and broadens the direct line of the UMIST project. Similar work is being done elsewhere, and extensions in many directions are possible: from batch production to mass production, to office work, to professional work in relation to 'expert systems', and so on. The intellectual basis of 'human-centred' systems needs further development, while as the early attempts to achieve these systems come into use their success in human and economic terms will need to be studied, analysed and learned from.

If those who engage in this work can achieve a change in the direction in which our technology develops, towards something which enriches the human spirit rather than impoverishing it, they will have the gratitude and thanks of those whose work is described here, but also and more importantly of a great multitude who will work with that more 'human-centred' technology.

Ross-on-Wye
1 November 1988

Howard Rosenbrock

Abbreviations

AI	artificial intelligence
BITZ	Bremer Innovations und Technologiezentrum
CAD	computer-aided design
CAITS	Centre for Alternative Industrial and Technological Systems
CAM	computer-aided manufacture
CAP	computer-aided planning
CIM	computer-integrated manufacture
CNC	computer numerically controlled
CV	curriculum vitae
DNC	direct numerical control
EEC	European Economic Community
FMS	flexible manufacturing system
MDI	manual data input
MIT	Massachusetts Institute of Technology
NC	numerical control
PR	public relations
R & D	research and development
RA	research assistant
RPM	revolutions per minute
SERC	Science and Engineering Research Council
SSRC	Social Science Research Council
SMC	shop-floor monitor and controller
UMIST	University of Manchester Institute of Science and Technology
USAF	United States Air Force

Further abbreviations will be found in the Glossary for Chapter 11, on page 168

Contents

Abbreviations	XVII
1. The Background to the Project	
<i>Howard Rosenbrock</i>	1
Introduction	1
Why is Taylorism so Widespread?	3
An Alternative Technology?	4
The Steering Committee	6
The Application for Funding	8
Steering Committee Discussions	9
Methodology	10
The Aims of the Project	12
2. The Technical Problem	
<i>Howard Rosenbrock</i>	17
Introduction	17
Some History	17
Numerical Control	22
Record-Playback	24
Cutting Technology	27
Design	28
Learning, but Learning What?	30
My Own View	32
Terminology	35
3. Towards a Design Methodology: A Psychologist's View	
<i>Martin Corbett</i>	37
Introduction	37
The Design Context	37
Design Criteria Related to the Human as a Component of the System	39
Design Criteria Related to the System as the Operator's Environment	39
Heuristics in Design Negotiation: The 'Blank Table Debate'...	41
Towards an Interactive Design Process	42
Conclusion	43
Appendix: Operator Control and Automation	
<i>Howard Rosenbrock and Martin Corbett</i>	44

4. The Social and Engineering Design of Computer Numerically Controlled Technology

<i>Paul Kidd</i>	51
Introduction	51
Numerical Control Technology	52
The Development of Cutting Technology Software	55
The 'Blank Table Debate' – A Divergence of Opinion	61
Conclusions and Personal Reflections – The Engineer as a Social Scientist	63

5. A Computer Science View

<i>Roger Holden</i>	65
Introduction	65
Software Structure	66
Aspects of the UMIST System	69
Software Development	72
Use of Alternative Programming Methods	75
Conclusion	78
Appendix: Skeleton Syntax for Part-programs	79

6. On the Collaboration Between Social Scientists and Engineers

<i>Lisl Klein</i>	81
Background – The Dynamics	81
Models of Science	85
Operational Issues	91
Making It Happen – Institutions and Infrastructure	95

7. (How) Can Technology be Redirected? A Scandinavian Perspective

<i>Håkon Finne</i>	99
Introduction	99
The Concept of Non-subordinating Machinery	100
Methodologies for Designing Redirected Technology	105
Was the Project Appropriate and Efficacious in a Redirection Effort?	108
Future Developments	112

8. A Works Director's View

<i>Allan Chatterton</i>	117
Introduction	117
Evolution of New Technology and Human Skills	118
The Operator, His Machine, and a Human-centred Approach to Technology	119
The UMIST Project	120
Conclusion	121

9. The Coordinator's View

<i>Harold Palmer</i>	123
Introduction	123
General Achievements	123
The Conduct of the Project	124
Cross-disciplinary Work by Engineers and Social Scientists..	125
The Importance of Management and of Technical Process	
Development in Human-centred Working	126
The Future	131
Final Comment	131

10. Human-centred Systems

<i>Mike Cooley</i>	133
Introduction	133
The Early Stages	133
Technology, Science and Ideology	134
Proposals for a New Approach	135
First Practical Moves	136
Issues of Status, Skill and Uncertainty	137
Human-centred Capabilities	140
Future Development of Human-centredness	141

11. The ESPRIT Project

<i>Shaun Murphy</i>	145
Introduction and Summary	145
The Meaning of Human-centredness	146
The Factory of the Future	147
ESPRIT Project 1217 (1199)	152
CIM Design Considerations	153
The BITZ Demonstration Site	153
The Shop-floor Monitor and Controller	155
The Sketching Module	156
The BICC Demonstration Site	157
The Rolls Royce Demonstration Site	164
The Human-centred Lathe Controller	165
Concluding Remarks	167
Appendix: Participating Organisations	168
Glossary	168

12. Postscript

<i>Howard Rosenbrock</i>	169
--------------------------------	-----

Appendices	173
1. Short CVs of Contributors	173
2. SERC Application	177
3. Transcript of a Part of the Steering Committee Meeting Held on 12 July 1982	185

The Background to the Project

Howard Rosenbrock

Introduction

The personal experience out of which this project grew can be described in very few words. From about 1968, in collaboration with colleagues, I had been developing a theory, and a large software package, which would allow the computer-aided design (CAD) of control systems. The aim was to give the designer a more powerful tool to aid him in his work: a computer with graphic display and interactive conversational working, rather than paper and pencil and slide-rule. Our approach at that time was not a question of principle, but arose from previous industrial experience and our intention to use the methods we developed for our own purposes.

By 1974 it was coming to be accepted that this approach was successful. It allowed the designer to investigate a tentative design, and to see whether it was stable, what was the margin of stability, the speed of response, and so on. If these were unsatisfactory, it suggested ways in which they could be improved. An existing skill in the designer was presupposed, and for the most effective use of the CAD system this had to develop in certain rather natural directions.

If this work was successful at one level, it was quite unsuccessful at another. For many years, the preferred lines of research in control engineering had been very different. They had been aimed not at accepting the designer's skill and trying to extend it and make it more productive but rather at replacing his skill by a large computer-algorithm. The designer, it was supposed, would specify the problem in great detail, including, say, the dynamic behaviour of the plant to be controlled, the constraints, and a cost function to be minimised. Then the computer would find the best solution satisfying these requirements.

There were some technical reasons for this tendency. In the previous decade, interactive computing was not widely available, and methods were adapted to off-line batch working. The aerospace problem, then heavily funded, was also better suited than most industrial control problems to an algorithmic approach. There was clearly more than this involved, however, because objections were raised to the need for skill in the user. What we had regarded as an advantage of our approach, its use of an existing skill and the opportunity it gave for this skill to develop further, was seen as a defect: a method which would work only by virtue of the skill of its user was seen as, to that extent, incomplete.

Engineering problems can seldom be reduced to two independent stages, of formulation followed by solution. Consequently, the actual practice of designers using algorithmic methods did not differ so greatly from ours as this contrast in aims would suggest. Skill was necessary to apply the algorithms, and many successive attempts might be needed to obtain a solution. Nevertheless, the difference in ultimate aims was clear, with the algorithmic methods looking towards a continual reduction in the need for skill, and our own aiming to foster skill and cooperate with it.

The contrast inevitably brings to mind Frederick Winslow Taylor's 'Scientific Management',¹ with its goal of eliminating skill and responsibility in the worker: 'Under our system the workman is told minutely just what he is to do and how he is to do it; and any improvement which he makes upon the orders given to him is fatal to success'. My unease was expressed in the 1975 paper on 'The future of control'² and more strongly in a number of later papers³.

The spirit in which future technology was to be developed had a particular importance at that time. It was clear to those involved in the developing fields of microelectronics, computers, and communications, that very great changes were impending throughout industry and commerce, and in many professional activities such as engineering and medicine. Public awareness came later, but already the outlines of the developments were visible.

A period of rapid change in technology offers great scope for experimenting with alternative routes of development. Once the new techniques have become established, it becomes much more difficult to examine alternatives: the 'entrance fee' for the alternative is the cost of matching the performance of the dominant technology, and it rapidly becomes so high that options which were available become closed off.

For this reason it seemed that the period of 15 to 20 years after 1975 would be one of critical importance. If new technology was developed in a Tayloristic spirit, the working conditions of very many people would be severely damaged. Their scope for initiative and the exercise of skill would be reduced, and work would become more mechanical, more routine, and less human. Because of the power of new technology, this process could be carried further in the traditional domain of Taylorism – production and some office work – and could be extended to new areas, such as design and medical diagnosis, which had been largely unaffected in the past.

In 1979, the generous award of a five-year Senior Fellowship by the Science and Engineering Research Council (SERC) freed me to work on this question, which has two interrelated aspects: 'Why is Taylorism so widespread?' and 'Is it possible to demonstrate a better alternative?'

Why is Taylorism so Widespread?

Explicit support for Taylor's system of 'Scientific Management' is now uncommon, though it can still be found in surprising places. On recently quoting Taylor's dictum, given above, to the eminent Production Manager of a large British company, I was told 'But I agree with that. It's right.'

More generally, Taylorism is seen as an early and crude attempt, which is now outdated, to bring system and order into production engineering. Later developments such as job enlargement, job enrichment, and autonomous groups are stressed, which are ways of alleviating some of the worst features of work organised on Tayloristic lines. Yet if one looks at the actual practice of designers of production systems, the underlying spirit is as strongly Taylorist as ever. The same tendencies can be seen in the development of computerised office systems and CAD systems, though computer scientists are perhaps more generally aware of the problems this brings than are production engineers.

Reactions to this situation tend to be strongly influenced by political persuasions. Condemnation is stronger from the left than from the right, though we can note the pronouncement of Pius XI, 'from the factory dead matter goes out improved, whereas men there are corrupted and degraded'.⁴ Braverman⁵ argues that Taylorism is an inevitable consequence of capitalism, from which it follows that political change must precede any change in work organisation or in machines. To the objection that political change in Russia and Eastern Europe has led to an entrenchment of Taylorism, he replies that the technology is after all a capitalist technology, adopted by Lenin⁶ as an expedient to escape from Russia's desperate and dangerous industrial backwardness, and entrenched by international competition.

This might have been acceptable at an earlier date, but 70 years after the Revolution one would expect to see at least the beginning of something different. The question was certainly raised in Russia, where in the early 1920s there was a vigorous debate⁷ on whether a different kind of technology should be sought. Those who thought it should be were defeated, and the question has never been re-opened.

My own views are strongly coloured by contacts with an East European country where new manufacturing technology is under development. Both industrial practice and the aims of research had strong Taylorist tendencies. Attempts to persuade those concerned that a better course was possible met with no success, and I was left with the conclusion that no project such as the one described in this book would have found any support there.

Noble⁸ emphasises another but allied causation for Taylorism, namely management's efforts to obtain and extend control over the workforce. Though he stresses the link with capitalism, one can readily extend the argument to any bureaucratic management system. There is the experience of socialist countries to support the extension to them, and my own contacts with large-scale workers' cooperatives, whether the capital is state-owned or worker-owned, indicate that it can be applied there also.

But to say that bureaucracies always attempt to extend their control and limit the initiative of those whom they direct is less than an explanation. Why should they do so, when it involves the loss of the contribution which