Open Research Online

The Open University's repository of research publications and other research outputs

Can a machine design?

Journal Item

How to cite:

Cross, Nigel (2001). Can a machine design? Design Issues, 17(4) pp. 44–50.

For guidance on citations see [FAQs](#).

© [\[not recorded\]](#)

Version: Accepted Manuscript

Link(s) to article on publisher's website:

<http://dx.doi.org/doi:10.1162/07479360152681083>

Copyright and Moral Rights for the articles on this site are retained by the individual authors and/or other copyright owners. For more information on Open Research Online's data [policy](#) on reuse of materials please consult the policies page.

oro.open.ac.uk

CAN A MACHINE DESIGN?

NIGEL CROSS

*Department of Design and Innovation
Faculty of Technology, The Open University
Milton Keynes MK7 6AA, UK*

Abstract: One strand of my research has been concerned with the computer as a design tool; but a second strand has been concerned with design computing as a research tool for improving our understanding of the design process. Some of this latter research is based on the simulation of computer behavior by human beings - a reversal of the more usual approach - and some is based on comparisons of computational models with human design behavior. Despite recent doubts expressed by some authors, I suggest that the question, 'Can a machine design?' is still a useful question to ask.

Introduction

Asking 'Can a machine design?' is similar to asking 'Can a machine think?' The answer to the latter question seems to be, 'It all depends on what you mean by "think".' Alan Turing (1950) attempted to resolve the question by his 'Turing Test' for artificial intelligence - if you could not distinguish, in a blind test, between answers to your questions provided by either a human being or a machine, then the machine could be said to be exhibiting intelligent behavior, i.e. 'thinking'.

In some of my research related to computers in design, I have used something like the Turing Test in reverse - getting human beings to respond to design tasks as though they were machines. There have been various intentions behind this strategy. One intention has been to simulate computer systems that do not yet exist; another has been to try to shed light on what it is that human designers do, by interpreting their behavior as though they were computers. My assumption throughout has been that asking 'Can a machine design?' is an appropriate research strategy, not simply for trying to replace human design by machine design, but for better understanding the cognitive processes of human design activity. However, this assumption has been challenged recently. In this paper I will first review some of my research, and then return to this challenge.

Using humans to simulate computers

My first research project began when I completed my undergraduate course in architecture in the mid-sixties and went on to study in the new field of design research, at the Design Research Laboratory at UMIST, Manchester, run by John Christopher Jones. My MSc research project was in 'Simulation of

Computer Aided Design' (Cross, 1967) - a novel but strange idea that we might get some insights into what CAD might be like, and what the design requirements for CAD systems might be, by attempting to simulate the use of CAD facilities which at that time were mostly hypotheses and suggestions for future systems that hardly anyone really knew how to begin to develop. The strangeness about this idea was that we would effect these simulations through getting human beings to pretend to be the computers! This was the reverse application of the 'Turing Test'.

The project was based on getting designers (architects) to attempt a small design project in experimental conditions (like the protocol studies and similar studies that have grown up since that time). They were given the design brief, and asked to produce a sketch concept. As well as conventional drawing materials, they had a simulated computer system to help them: they could write questions on cards located in front of a closed-circuit TV camera, and would receive answers on a TV screen in front of them. In another room, at the other end of the CCTV link, was a small team of architects and building engineers who attempted to answer the designer's questions. Thus we had a very crude simulation of some features of what might actually now be parts of a modern-day CAD system, such as expert systems and databases.

The designers who participated in these experiments were not told what to expect from the 'computer', nor given any constraints on the kinds of facilities they might choose to ask of it. Thus I hoped to discover what kinds of facilities and features might be required of future CAD systems, and gain some insights into the 'systemic behavioral patterns' that might emerge in these future human-computer systems.

I conducted ten such experiments, each of which lasted about one hour. The messages between designer and 'computer' were recorded, and one of the analyses made was to classify them into the topics to which they referred, from the client's brief to construction details. This kind of data gave some insight into the designers' patterns of activity, such as a cyclical pattern of topics over time, from requirements to details and back again. The number of messages sent in each experiment was quite low, with normally several minutes elapsing between requests from the designer. Of course, the response time from the 'computer' could also be quite long, typically of the order of 30 seconds. Despite this apparently easy pace of interaction, all the designers reported that they found the experiments hard work and stressful. They reported the main benefit of using the 'computer' as being an increased speed of work, principally by reducing uncertainty (i.e. they relatively quickly received answers to queries, which they accepted as reliable information).

I also tried a few variations from my standard experiments. The most interesting was to reverse the normal set of expectations of the functions of the designer and the 'computer'. The 'computer' was given the job of having to produce a design, to the satisfaction of the observing designer. It was immediately apparent that in this situation there was no stress on the designer - in fact, it became quite fun - and it was the 'computer' that found the experience to be hard work. This led me to suggest that CAD system designers should aim for

CAN A MACHINE DESIGN?

‘a much more active role for the computer, tantamount to a virtual inversion of the present designer/computer roles. The computer should be asking questions of the designer, seeking from him those decisions which it is not competent to handle itself. The computer could be doing all the drawing work, with the designer instructing amendments. Drawings presented by the computer on a graphic interface would be gradually completed as the designer made more decisions . . . Programmed to proceed as far as possible without human intervention at each step, the computer would ask for decisions as required . . . We should be moving towards giving the machine a sufficient degree of intelligent behavior, and a corresponding increase in participation in the design process, to liberate the designer from routine procedures and to enhance his decision-making role.’ (Cross, 1967)

This vision of the intelligent computer was based on an assumption that a machine *can* design - that it can be programmed to do a lot of the design work, but under the supervision of a human designer. I still think that there is something relevant in this vision of the computer as designer - it still offers a more satisfactory basis for the human-machine relationship in computer aided design than current CAD systems. Why isn't using a CAD system a more enjoyable, and perhaps also a more intellectually demanding experience than it has turned out to be?

Comparing human and machine performances

I continued research on this question of human and machine roles in computer aided design for my PhD (Cross, 1974). My earlier studies had suggested that using computers in design might have adverse effects, such as inducing stress, whilst not having any beneficial effects on the quality of the resulting designs. The only ‘positive’ effect that CAD appeared to have was to speed up the design process. The potential negative effects of CAD that I identified were an intensification of the designer's work rate and a concomitant reduction in the person-power required in design offices. But on the other hand I suggested that CAD in architecture might lead to better communication between members of the design team, and to the inclusion of a wider range of participants, such as the new building's users. (Cross, 1972.)

However, I still seemed to have a belief that a machine *can* design - that it can produce designs that are somehow better - more efficient, or more elegant, or something - than designs produced by humans. Drawing on research in problem solving (of the ‘travelling salesman’ route-layout type) at the pioneering artificial intelligence centre at Edinburgh University, I expected that human-machine interaction (rather than wholly-human or wholly-machine problem solving) would efficiently produce design solutions that were better than either a human or a machine could produce alone.

So I set out to test that hypothesis, using the problem of efficient room layouts in a building plan. (There had been some early attempts at producing ‘optimum’ room-plan layouts. The idea was that if you had some data for the numbers of journeys that typically would be made by the future building's users between the different elements of accommodation, then you could get a computer to optimise the layout so as to minimise the ‘circulation cost’ - i.e. the number of journeys multiplied by the lengths of journeys. Rooms that would have a high number of journeys between them would be placed close together,

and so on.) I devised experiments in which fully-automatic computer programs, un-aided humans and human designers aided by interactive layout programs tackled the same layout problems.

I fully expected to replicate the Edinburgh results, and was genuinely surprised to find that (a) there were no significant differences between the performances (i.e. the efficiency of the layouts) of un-aided humans and automatic computer programs, and (b) human-machine interaction produced worse results than either un-aided humans or automatic machines! There were some strong mitigating circumstances arising from the crude nature of the human-machine interaction that was possible at that time (teletype terminals and storage-tube displays!), but nevertheless it was a surprising result, that shook my confidence in CAD developments at that time, and led me to the conclusion that machines *cannot* design very well at all, and actually make design results worse, rather than better. In writing up my thesis for publication in book form (Cross, 1977a), I concluded that CAD would be of very limited positive effectiveness as a design aid, but could have profound negative effects on design activity and the job of being a designer. In an article in the *RIBA Journal*, I confessed that 'I have seen the future; and it doesn't work!' (Cross, 1977b.)

Eliciting computable rules from human behavior

It was a long time before I returned to similar kinds of research. The developments in computing and CAD in the 1980s made me realise that, for good or bad, using computers in design practice was inevitable (indeed was already ubiquitous).

A project I was involved with in the mid-1990s was based on a sub-question of 'Can a machine design?' It was 'Can a machine make aesthetic judgements?' The aesthetic aspects of design are often assumed (by designers, if not by some of the CAD researchers) to be some of the most intractable aspects for computers to attempt. But my colleagues and I thought that there might be some implicit rule-based behavior in aesthetic judgements, which might be modelled in a computer system.

The design domain we were working in was that of graphic design, where designers normally guard their aesthetic freedom very jealously. We agreed with them that it might not be possible to construct rules of aesthetically 'good' design, but we thought that it might be possible to establish rules of 'bad' design. If so, then a rule-based expert system could be used to evaluate graphic designs, pointing out the 'bad' features. Users of such a system, even if they could not produce 'good' designs, might at least be able to produce designs that were 'not bad'. We had in mind users of wordprocessors and simple desktop publishing systems, producing amateur graphic designs such as in-house notices, newsletters, etc.

We collected examples of such amateur designs (A4 'posters') from around our own University Department noticeboards, and submitted them to critiques by two expert graphic designers. We then converted the experts' comments on the 'bad' design features into rules, and tested these rules by using ourselves as

CAN A MACHINE DESIGN?

‘human computers’ - strictly following the instructions in a machine-like way, until ambiguity was eliminated. (In a way, this was also following Turing’s early theoretical argument that problem-solving programs might in principle run on any kind of ‘machine’ - thus separating the program from the computer.) Then we applied the rules to a new sample of posters and compared the ‘machine’ results with those of the human experts’ critiques of the new posters.

We found that a relatively small number of rules (less than twenty) could be used to eliminate common ‘bad’ design features. Some of our rules were very simple, such as ‘Left and right margins should be equal’, or ‘If more than 70% of text is centred, then all text should be centred.’ But applying such simple rules does lead to designs that are ‘not bad’. We also found that the human experts were frustratingly inconsistent in applying their own ‘rules’; when we pointed this out to them, they were quite happy to accept that the rules were indeed valid, but need not always be applied rigourously in every case! This seems to be something like allowing the judge some leniency in passing sentence. This work is reported in Glaze et al. (1996).

This was not a demonstration that a machine can design. It was a demonstration that, in principle, a machine can do some things that many human beings regard as a uniquely human attribute - in this case, making aesthetic judgements. To me it was also a confirmation of the value of asking ‘Can a machine design?’ as a research strategy for investigating design. We had learned something about a relatively difficult area of design activity, and also something about designers and their ways of thinking.

Natural versus artificial intelligence

We might not necessarily want machines to do everything that human beings do, but setting challenges for machines to do some of the cognitively hard things that people do should give us insight into those things and into the broader nature of human cognitive abilities. I had always assumed that this argument was one of the validations for research in artificial intelligence. Thus we would learn more about ourselves. For example, the research programme in computer chess playing has presumably not had the ultimate aim of making it unnecessary for humans ever to ‘need’ to play chess again. Rather, it has been to gain understanding of the nature of the ‘problem’ of the chess game itself, and of the nature of the human cognitive processes which are brought to bear in chess playing and in the resolution of chess problems.

That has always been my assumption about the value of trying to get machines to do things that human beings do, whether that is playing chess or designing. But John Casti, of the Santa Fe Institute, came to a rather disturbing conclusion about the lessons that may have been learned from chess-playing machines. In his book, *The Cambridge Quintet*, Casti (1998) imagines a debate on computation and artificial intelligence between Turing, Wittgenstein, Schrödinger, and Haldane, chaired by C. P. Snow. In a postscript, Casti refers to the 1997 defeat of the world chess champion, Garry Kasparov, by the computer program Deep Blue II, and he quotes Kasparov as saying, ‘I sensed an alien intelligence in the program.’

Casti then goes on to come to the rather surprising, and depressing conclusion that 'we have learned almost nothing about human cognitive capabilities and methods from the construction of chess-playing programs'. So, in computer-design research, will we be forced to come to the same conclusion, that 'we have learned almost nothing about human cognitive capabilities and methods from the construction of designing programs'? Will designers rather nervously contemplate the 'alien intelligence' of the designing programs? Will we have built machines that can design, but also have to bring ourselves to Casti's view of the 'success' of chess-playing machines: 'the operation was a success - but the patient died!'?

Perhaps Casti is being unduly pessimistic. One thing that we have learned from chess-playing programs is that the brute force of computation actually can achieve performances that outmatch human performance in a significant area of human cognitive endeavour. And surely researchers of computer chess-playing have learned something of the cognitive strategies of human chess players, even though their programs do not 'think' like humans? Certainly I believe that, on a much smaller scale, our research on aesthetic judgements in design had that kind of value.

In more recent research I have also found that computational models of design activity can be useful descriptive or explanatory models of human design behavior. This has been particularly so in the field of creative design, where attempts to build computational models have provided some useful paradigms for the nature of creative design activity (Cross, 1997, 1998). I think that many of the attributes of design cognition that we now regard as essential features of the natural intelligence of design (Cross, 1999) have been identified as a result of attempts to simulate design activity in artificial intelligence.

It seems to me that research in artificial intelligence should always address the question, 'What are we learning from this research about how people think?' Similarly, our computer-design research should attempt to tell us something about how designers think. I believe we can learn some important things about the nature of human design cognition through looking at design from the computational perspective (although 'the computationalist paradigm in design research' has also been challenged by Liddament, 1999). For me, the value of asking the question 'Can a machine design?' is that it begs the corollary question, 'How do people design?'

References

- Casti, J (1998) *The Cambridge Quintet*, Little, Brown, London
- Cross, N (1967) 'Simulation of Computer Aided Design', *MSc Dissertation*, Design Research Laboratory, Department of Building, University of Manchester Institute of Science and Technology
- Cross, N (1972) 'Impact of Computers on the Architectural Design Process', *The Architects' Journal*, 22 March, pp. 623-628
- Cross, N (1974) 'Human and Machine Roles in Computer Aided Design', *PhD Thesis*, Design Research Laboratory, Department of Building, University of Manchester Institute of Science and Technology
- Cross, N (1977a) *The Automated Architect*, Pion Ltd, London

CAN A MACHINE DESIGN?

- Cross, N (1977b) 'Problems and Threats of Computer Aided Design', *RIBA Journal*, October, pp. 439-440
- Cross, N (1997) "Creativity in Design: analyzing and modeling the creative leap", *Leonardo*, Vol. 30, No. 4, pp. 311-317
- Cross, N and Dorst, K (1998) 'Co-evolution of Problem and Solution Spaces in Creative Design', in Gero, J S and Maher, M L (eds.) *Conference on Computational Models of Creative Design: HI98*, Key Centre of Design Computing, University of Sydney
- Cross, N (1999) 'Natural Intelligence in Design', *Design Studies*, Vol 20, No.1, pp. 25-39
- Glaze, G, Johnson, J and Cross, N (1996) 'Elicitation of Rules for Graphic Design Evaluation', in Gero, J S and Sudweeks, F (eds.) *Artificial Intelligence in Design: AID96*, Kluwer, Dordrecht
- Liddament, T (1999) 'The Computationalist Paradigm in Design Research', *Design Studies*, Vol. 20, No. 1, pp. 41-56
- Turing, A (1950) 'Computing Machinery and Intelligence', *Mind*, No. 2236 (October) pp. 433-460