Understanding Instruction Books and Programs

W. Hugh Chatfield I.S.P.
CyberSpace Industries 2000 Inc.
Nepean, ON, Canada

Abstract - This paper takes the form of a dialog between two 'super-observers' SO1 and SO2 discussing a modified Chinese Room experiment to conduct an examination of seldom-discussed artefacts, the instruction manual and/or the computer program in the room.

SO2 – So what's all this about?
SO1 – This is an experiment to measure whether or not it is possible to build an 'artificial intelligence'. What do you see?
SO2 – I see three rooms. A person is sitting in front of the rooms with a terminal. He has a switch that he sets to room 1,2, or 3. When he types on the terminal, eventually some response appears on his screen from whichever room he selects.
SO1 – What is this text?
SO2 – It is Chinese text.
SO1 – Do you see anything different about the responses from the three rooms?
SO2 – They all seem to be reasonable answers so far. Room 2 answers more slowly than the other two rooms. What has this got to do with 'artificial intelligence' and what do we mean by that phrase?
SO1 – Let's start with Alan Turing's 1950 paper [1] . He addressed the question – "Can machines think?" Turing proposed a quantitative test to help us determine the answer to this.
SO2 – Yes, I know this. He proposed an 'imitation game' based on the game where a man and a woman try to fool an observer in an other room. The man pretends he is a woman, and the observer tries to guess which one is really the woman. Then we replace one of these with a computer. If we then try this with many interrogators, will the interrogator decide correctly as often when the game is played like this as he does when the game is played between a man and a woman?
SO1 – And you do notice that Turing did not originally ask the observer to identify the computer. He only asked you to identify the woman. If there was no statistical difference in the number of wrong responses when a computer replaces one of the participants, Turing suggests that we must conclude the machine does 'think'. He devised this quantitative test because we can't really answer the question – "does the computer think?", any more than we can answer – "does the man or woman think?" Hence if we can't tell the difference based on measurement, then all of them must be 'thinking'.

I believe Turing's notion of measurement is important. Bohr proclaimed that "physics in general, and quantum mechanics in particular, do not describe the world itself, only what we are able to say about it". Information plays a mediating role between the objective material world and what we know about it. "We never see a chair, Bohr would say: We receive sense impressions that give us information which our brains somehow process into the idea of chair." [8]. Measurement is extracting information from a material world system. Further, Churchland [9] indicated that a condition for "conscious intelligence" is that the system must not only exploit the information it already has, but also, via an "energy flux" (including through its sense organs), must increase the information it contains.
SO2 – So this is a Turing Test apparatus then?
SO1 – Actually it is a version of Searle's Chinese Room experiment [4]. Searle devised this test to demonstrate that mere computation cannot give rise to 'understanding', or if you prefer he was arguing against what is commonly called 'strong AI' [5]. I like Penrose's description of this thought experiment. If there exists a program that is said to "understand" Chinese – we give it as a set of instructions to a person in a room who has absolutely no understanding of Chinese. He is presented with a set of Chinese characters. He takes this set and processes them based on instructions he has been given. The instructions ultimately get the person in the room to copy other Chinese characters into an output. What he doesn't know is that the symbols coming in represents a question in Chinese, and the symbols going out represent an answer to that question written in Chinese.
SO2 – What is the test?
SO1 – Searle suggests you to put yourself in the place of the person in the room, or even be that person, and asks you whether you understand Chinese? Searle asserts that the person in the room has no understanding of Chinese, and no matter how you configure the room or the person, will never understand Chinese… hence manipulation of these characters can never give rise to 'understanding'. From this you can conclude that a computer, which is merely a manipulator of these patterns can never 'understand' Chinese either – or for that matter - anything. Searle's argument is based on what he believes is two truths: *brains cause minds*, and *syntax doesn't suffice for semantics*. This particular version of the experiment will assume Searle's conditions and work backward to determine the characteristics of the instruction manual and program.
SO2 – We have been using the words 'thinking' and 'understanding' a little loosely here haven't we?

SO1 – Good observation. I like Edward de Bono's definition. In his text - Practical Thinking [2] he describes 3 ways a human 'knows what to do'.

1) *Instinct* - hard wired reaction to a situation (developed via evolution over long periods of time and inherited via the information encoded in the genes)
2) *Learning*
    a. First hand - trial and error response to situation
    b. Second hand - acquire immediate response based on others first hand learning.
3) *Understanding* - changing an unfamiliar process into one that is familiar and one that you know how to handle.

De Bono says: "the change from unknown to known is 'understanding', and the way this change comes about is 'thinking'. It may be a matter of understanding what something is or it may be a matter of understanding how to bring about some effect. Understanding is finding out how to bring about some effect. Understanding is finding out what to do. This finding out is thinking."
SO2 – So with this definition, what is doing the 'thinking' in Searle's room? It seems the person in the room may be 'learning' by repeating the instructions he has been given over and over.
SO1 – Perhaps, but Searle says even if the person has the instructions memorized, he still will have no understanding of Chinese – and neither will a computer.
SO2 – Yet the room responds to questions in Chinese. Surely there has to be 'understanding' somewhere in the room – else how can it answer the questions.
SO1 – That is what this experiment is designed to explore. Let's go into the rooms to see what is going on.

~ Room 1~

SO1 – Here is room 1. What do you see?
SO2 – There is a man at a computer terminal. He receives the text typed in from the outside on a screen, and he types an answer back on the keyboard.
SO1 – Let's chat with him and see what we can determine. Hello.
Man in Room 1 – Hello.
SO1 – I see you are typing answers to Chinese questions typed on the screen. You must understand Chinese even though you speak English.
Man in Room 1 – Yes, well enough. However, I do have to translate the Chinese to English before I think about it. I still think in English. I formulate my answers in English and translate back to Chinese as I type the answers.
SO1 – Do you know how you manage to generate an answer?
Man in Room 1 – I took Chinese lessons for many years when I had a computing work assignment in China. I felt it would be beneficial to understand the local language.

SO1 – So what does this tell us?
SO2 – Certainly he is exhibiting De Bono's second hand learning here. He is applying knowledge and rules acquired from others to process the Chinese.
SO1 – Does he understand Chinese?
SO2 – By his own admission he does?
SO1 – Ah, but how do you know he does?
SO2 – Based on *my* first hand learning experience, I would say if I had attended a Chinese language course for several years, I would have at least a rudimentary understanding of Chinese similar to the Man in Room 1. Further he exhibits 'understanding' by changing the unfamiliar language – Chinese – into a more familiar one – English - then reasons in English, develops an answer in English – finally translating back into Chinese. By de Bono's definition – he is exhibiting 'understanding' through a 'thinking' process.

~ Room 2 ~

SO2 – Lets try room 2. What do you see?
SO2 – There is a man who appears to be a twin of the man in room 1, sitting in front of a terminal. When he receives the text, he consults a large instruction manual, makes a lot of notes and then types out an answer in Chinese on the terminal.
SO1 – Now Searle asks you to put yourself in the place of this man. Do you understand Chinese?
SO2 – I think you may be trying to fool me here. As Turing pointed out – there is no way of answering that question unless you are that person. Turing proposed to answer the question by looking at the output from the room. There is no question that he is typing the answers in Chinese.
SO1 – Where did you get that manual?
Man in Room 2 – I got it from my brother in Room 1.

SO1 – What do you think is going on here?
SO2 – I think this is a case of de Bono's second hand learning. I suspect the manual he consults contains knowledge and rules for processing Chinese. Further, since he got it from his brother in Room 1, I suspect it contains instructions about how to answer questions posed in Chinese, in Chinese – most probably by translating into English.
SO1 – Why do you say that?
SO2 – Because that's all the man in room 1 knows how to do.

Room 3

SO1 – OK - lets take a look at room 3. What do you see?
SO2 – There is no terminal. There is a black box. I assume this is some sort of computing device since this room produced answers in Chinese as well.
SO1 – And what is making this happen?
SO2 – I would expect some sort of program being executed in the computer.
SO1 – I don't think you could put yourself in the place of this computer – but do you think this computer understands Chinese?
SO2 – Again, since the program answers back in Chinese, I don't know how you could say for sure it doesn't.
SO1 – Searle would say that neither the man in room 2 nor the computer could have any understanding of Chinese because both are merely manipulating the Chinese characters, and this manipulation does not give rise to 'semantics' and 'understanding'.

SO1 – Are we any further ahead in understanding what is going on?
SO2 – I think we need to take a closer look at the instruction manual and the program. This has to be key to understanding what is going on here. Certainly all three rooms are responding in Chinese to questions posed in Chinese, hence the manual and the program must somehow represent a great deal of understanding and knowledge of Chinese. The man in room 2 said his brother in room 1, who does understand Chinese, wrote it. Since his brother also is a programmer, he probably wrote the program in room 3. I can think of two ways these instructions could have been written.

> *Type 1:* The man in room 1 developed detailed instructions on how to gain a sufficient knowledge of Chinese to translate Chinese into English. This is most likely since this is certainly what he knows how to do. He then relies on the English language processing capabilities of the man in room 2 to reason and develop an answer in English. He will not need detailed instructions on how to do this part. The instructions then must instruct on how to convert from English back to Chinese. This would be an reasonable model of what the man in room 1 is doing, and again is what the man in room 1 knows how to do.
>
> *Type 2:* The man in room 1 develops even more detailed instructions that not only deals with the translation of Chinese to English – but also embodied all his first and second hand learning that deals with his understanding, reasoning and responding in English – followed by translation to Chinese. He also has to do this without any transfer of knowledge about Chinese to the man in room 2. In other words, a type 2 set of instructions have to be very carefully designed if Searle's premise, that the man in room 2 doesn't understand and will never understand Chinese, is to remain true. This will be an extremely difficult task to undertake.

Of course if the man in room 1 were a native speaking Chinese person who had knowledge of how to write instructions in English, the instructions probably wouldn't need to contain anything about translation. However, let us continue this train of thought for while to see where it leads.

I actually don't think the man in room 1 has enough knowledge to be able to accomplish this. I might even go further and say no person currently has this capability.

SO1 – Can you explain more about what you mean by Type 2 and why you say that?
SO2 – Sure - lets compare Type 1 with Type 2 to see the difference. Suppose the man in room 2 saw the following pattern on the screen.

Figure 1 - Chinese character

Type 1 instructions might ask him to look this up in the manual. The manual would then instruct him that this pattern was the Chinese symbol for 'cat'. Having done this the man in room 2 would have the learned capability of coming up with the concept of 'cat' every time he saw this pattern. He would further connect this pattern with everything he has ever learned about 'cats'.

Further if he saw a pattern –

Figure 2 – A 2cd Chinese character

He might reason (without further instruction) that since this pattern includes the pattern for 'cat' this might be a variation of the concept of a cat – and indeed if he looks it up in his manual he sees it means 'kitten'. This is de Bono's first hand learning at work. Already he can reason that the first part of the pattern must be a modifier to deal with a subset of 'cat' or a type of 'cat'.

Type 1 instructions rely completely on the first and second hand learning of the English language and the isomorphism to the real world developed over a lifetime by the man in room 2. This is what the man in room 1 does when he processes the Chinese. He doesn't have a second way of doing this.

If the man in room 2 puts other patterns together, he might find it represents a question about a cat. For example it might ask what colour of cat is considered unlucky? The answer would be 'black cat'. Without knowing anything more about Chinese he might deduce that you would need some sort of modifier representing colour along with the 'cat' pattern.

With Type 1 instructions, the man in room 2 is learning Chinese. In order to comply with Searle's condition for the experiment that he knows nothing about Chinese, and learns nothing about Chinese – the instructions cannot be of Type 1 – they must be of Type 2.

Further, the instructions must be abstract enough to never introduce any real world concept of what these patterns mean. The instructions must not rely on any first or second hand learning of the man in room 2, except for how to recognize patterns, read instructions in English and carry them out to the letter. Any other learning must be completely replaced by instructions that carry out exactly what is going on in the head of the man in room 1…. without revealing to the man in room 2 anything about the Chinese language.

A similar argument would apply to the room 3. The only thing you could rely on is the computers ability to read and follow instructions written in binary. These instructions are usually derived by the computer from other instructions written in a higher-level language – but this is merely a convenience to the program writer. In room 3, the computer instructions also have to carry out what is going on in the mind of the man in room 1.

If you think about how this might be accomplished, I'm sure you would agree that Type 2 has to be several orders of magnitude more complex than Type 1, and would require the man in Room 1 to be sufficiently aware of how his brain works, and have enough smarts to translate these into a set of instructions in English, and a set of instructions in some computer language.

Think for a moment about how to deal with the question being asked. What colour cat is considered unlucky? Imagine you know nothing about this particular language (English). What are some of the things that would have to be handled in the instructions?

- How to recognize the various patterns ('c', 'a', 't', etc.)
- How to recognize 'words' or collections of patterns ('cat')
- How to connect words with concepts (what does the word 'cat' mean in the real world?)
- How to connect concepts with facts related to those concepts (what is currently known about 'cat' – size, shape, colour, other names, relationships with witchcraft, etc.)
- How to determine what the question is and what it means?
- How to process the known facts about 'cat' to see if an answer is known.
- How to form up the answer 'black cat'

Of course I am using English to illustrate, but the same holds for Chinese, although the conventions of working with these patterns might be totally different. Since the man in Room 1 obtained his 'learning' and 'thinking capabilities' in the English language over a lifetime using his hard wired capabilities, as well as first and second hand learning related to Chinese, I think he will be very hard pressed to write down how he does any of this.

SO1 – Are Rooms 2 and 3 substantially different from Room 1.
SO2 – I don't think so. In Rooms 2 and 3, you have mechanisms (a brain and a computer processor) each 'executing' a set of instructions representing the complete set of first and second hand learning of the man in Room 1 as well as all his thinking processes (wired-in or learned). The instruction manual and program do nothing by themselves. It is the dynamic properties of that set of instructions that exhibits the behaviour of answering questions posed in Chinese, in Chinese.

In Room 1, you have a mechanism (a brain) that is using a set of "instructions" obtained through a lifetime of first and second hand learning (processing English), that allows the execution of another set of instructions obtained through first and second hand learning over the past few years (processing Chinese). It is the dynamic properties of these two sets of instructions that exhibit the behaviour of answering questions posed in Chinese, in Chinese.
SO1 – And the question of what 'understands'?
SO2 – I refer you to de Bono's definition again -- "the change from unknown to known is 'understanding', and the way this change comes about is 'thinking' ". In all three cases, it appears to me that an arrangement of patterns get converted to another set of patterns on the screen. In between is a dynamic process that encompasses both a set of learning, and a set of processes about how to change the unknown to a known. The three cases before you may differ in details and platforms executing the instructions, but all three cases exhibit 'understanding' and by de Bono's definition are 'thinking'. Turing would probably agree with this since we cannot distinguish any difference based on outputs alone.
SO1 – But doesn't Searle then argue that this merely makes a case for 'weak AI' – Rooms 2 and 3 are only 'simulating understanding'. He constructed this experiment to demonstrate that 'strong AI' can't be so. Strong AI proposes that the appropriately programmed computer really *is* a mind in the sense that computers given the right programs can be literally said to *understand* and have other cognitive states.
SO2 – This is hard to prove or disprove. Both men use identical processing to come up with an answer. The man in room 1 probably made the following (highly simplified) connections to come up with an answer.

Input ->Pattern – symbol for 'cat' / Colour black – considered – unlucky / Unlucky cat – means – black cat / Write symbols for black cat -> *Output*

He would have a lifetime of learning about what 'cat', black', unlucky' mean in the real world. He would then have to write the instructions as to how all these connections were made in his head, and then write the whole thing so that he is NOT relying on equivalent learning in the head of the man in room 2. The resulting instructions would likely be isomorphic to the instructions above – but with every name, concept and label disguised so that the man in room 2 cannot use any of his learning to process the instructions. Every piece of information and all processes that are required have to come from the man in room 1 - via the instructions. We must make sure that the man in room 2 cannot map what he is doing to anything he knows in the real world. Only if this holds will the man in room 2 be truly 'only manipulating patterns', and be equivalent to what is happening in room 3.

It is interesting to note that in order to comply with the conditions that the man in room 2 not understand Chinese, I have to assume the construction of a set of instructions and data that completely replaces the data, knowledge and processing that go on in the head of the man in room 1. All the semantics required for the room to operate is contained within this instruction set – and will never be found in the brain of the man – because this is how the room is constructed. The final question then is, not whether the man in room 2 understands what is going on – it is, ` if the construction of such a set of instructions was possible – would the dynamic properties of that set of instructions exhibit any evidence of "mind" or "cognitive states".

To me, the question of cognitive states and mind seems separate from the notion of understanding and thinking. To finally answer this we need to understand how 'mind' arises from brain. There is indeed a possibility that 'mind' arises out of particular unknown structures and processes in the brain, and that these are necessary for mind to arise. We don't know yet. For example, Penrose [6, 7] argues his position on this - that "appropriate physical action of the brain evokes awareness, but this physical action cannot be simulated computationally". If Penrose is correct, then room 3 could not exhibit self awareness, consciousness, and other things we attribute to humans – unless of course we find a way to construct a non biological structure equivalent to whatever in the brain gives rise to "mind".
SO1 – None the less, you seem to be arguing what Searle calls 'the systems reply'. Searle contends, it's just ridiculous to say 'that while [the] person doesn't understand Chinese, somehow the conjunction of that person and bits of paper [the instructions] might'.
SO2 – Absolutely, but I don't contend the bits of paper have understanding, I claim the information on those bits of paper, running as a dynamic process, either in the human brain or the computer, has to contain the understanding of Chinese, in order for them to respond in Chinese.

The understanding and knowledge of the man in Room 1, has been completely transferred into both rooms 2 and 3. Both rooms 2 and 3 run through the identical processes that the man in room 1 goes through to transform the input Chinese into the output Chinese. Again, it can't be anything else, since this is all the man in room 1 knows.

SO1 – You also said that you believed no one could currently develop such an instruction set. The argument seems to rest on the ability to construct the manual in room 2 and the computer instructions in room 3.
SO2 – It does, and right now I believe that the manual or program can't be built for the simple reason that there is incomplete knowledge of how the brain works. Without this knowledge we could not write the instruction manual

/program. We also have no idea how 'mind' arises from brain, so it is premature to conclude that it can't be done in a non-biological brain.
SO1 – Doesn't Kurzweil [3] predict otherwise in 'The Singularity is Near'?
SO2 – Quite so. He predicts that over the very near future, due to an exponentially increasing discovery process, we will be able to reverse engineer the brain and understand in some detail how it works. He further predicts that by 2045 we will be able to construct a "non-biological intelligence … one billion times more powerful than all human intelligence today" based on this newfound knowledge. He calls this event 'The Singularity'. The computing engines we implement this on might be quite different from what we know today.
SO1 – Does that computer in Room 3 approach this target?
SO2 – Lets go take a closer look

~ Room 3 ~

SO1 – What do you see?
SO2 – First I can see that this computer cannot be executing a Type 1 set of instructions.
SO1 – Because the Type 1 set of instructions relied on the first and second hand acquired learning of the man in room 2 and this is not built into this system?
SO2 – Exactly so. It has to be executing a Type 2 set of instructions. Everything has to be in that program. Now assuming that such a set of instructions could be written for the man in room 2, and further that a program could be written based on this, it is a valid question as to whether or not today's computers are sufficient to run this program. Both Turing [1] and Kurzweil [3] believed that the current computers of the time were insufficient to build an intelligent machine.
SO1 – Surely, at least in principle this is so. This still leaves the question as to whether or not the computer understands Chinese. What we have is still a 'machine' – how could we possibly believe it understands Chinese.
SO2 – For that answer, I think we better visit room 1 again.

~ Room 1 ~

SO2 – First lets stop time so we can look at what is happening in this man's brain. What do you see?
SO1 – Not much is happening. I see a lot of organic structures all frozen in time.
SO2 – What you are seeing is the state of the brain at a single point in time, or if you prefer, a vanishing small interval of time. This is equivalent to Searle's 'bits of paper'. Nothing is happening because thinking and understanding are dynamic processes. De Bono says 'understanding is finding out how to bring about some effect' or 'change from unknown to known'. Clearly this can only happen during a dynamic process.
SO1 – I grant you this.
SO2 – If we let time run over the period of time where he is looking at the Chinese pattern for cat on the screen, what do we see?
SO1 - I see photons leaving the face of the screen, getting focussed by the lens of his eye onto a retina at the back of his eyeball. The retina consists of a number of structures that send a collection of electrical signals into his brain. The brain is able to decode these signals and deduce that what the man is looking at is the Chinese character for 'cat'.
SO2 – How did he do this?
SO1 – It's very difficult to tell. There is a lot of chemical and electrical activity in the brain. There appears to be some synchronization occurring between parts of the brain that allows him to link the observed pattern with a stored 'mental picture' of a cat, and a whole collection of facts known about 'cats'. Exactly how this took place, or even how or where any of this is stored, I can't say.
SO2 – Could you write a program to do the same thing.
SO1 – Not unless I can understand what is going on a whole lot better.
SO2 – This why I said you were fooling with me when you showed me the computer answering questions. What part of the brain understands Chinese?
SO1 – I don't think any single part 'understands' Chinese. Certainly single parts bring about local and non-local effects, but none of the micro level effects seem to have anything to do with 'understanding Chinese'.
SO2 – Would you then say the brain, as a collection of parts understands Chinese.
SO1 – I would have to say yes, at least in the Turing sense. The entire system gives statistical evidence of this being true… but unless I am that person I cannot know for sure.
SO2 – Thus, in the case of the computer, even though no individual part of the computer 'understands' Chinese, there is evidence that the system as a whole understands Chinese, provided that the instructions are Type 2, and running as a dynamic process.
SO1 – Again, only if you could construct such a program running on this processor, and I think we both agree that at the moment we have no capability to construct such a program. So – where have we got?
SO2 – I think we can conclude the following:

- In order for this experiment to run as proposed, the man in room 1 must encode his entire acquired learning of Chinese into English instructions for the man in room 2 and computer instructions in some language for the computer in room 3.
- He must further encode his first and second hand learning of the English language and the isomorphism to the real world developed over his lifetime (i.e. what the string 'cat' represents)
- He must further encode everything he has learned over a lifetime into those instructions, since the instructions will have to make use of this knowledge to make decisions.
- He must further encode this knowledge in such a way as to avoid the man in room 2 using whatever learning and knowledge he has to assist in processing these instructions. What the instructions are doing must be deliberately hidden from the man in room 2. Without this condition, the man in room 2 will eventually learn and understand Chinese.
- All the above is also a necessary condition to ensure the entire set of knowledge required for the computer to do its job is given to it as input.
- Understanding of Chinese is 'inherent' in the properties of the instructions/program, that by themselves, viewed statically, are only ink on paper, or binary code stored in memory. Our perception that "understanding of Chinese" is taking place - arises from the dynamic properties of the instructions/program.
- All three rooms contain all the knowledge and understanding of the man in room 1 and also step through the same process that the man in room 1 goes through when he answers the questions.
- In both rooms 2 and 3, there is no awareness that Chinese is being processed, since that is how we constructed the rooms, yet all three rooms understand Chinese (by de Bono's definition).
- In order for room 3 functionality to be equivalent to the room 2 functionality it must first be provided with the capabilities of recognizing patterns, following the instructions provided, and most importantly – all the "wired-in" capabilities of the man in room 1 such as the ability to convert an "unknown" into a "known'.

SO1 – I can now clearly see the difficulty you allude to. What seems to be required is that the man in room 1 must download the entire contents of his brain, including some if not all of the "wired-in" properties of the brain.. It would seem that Searle's conditions can only be met if we could somehow do this brain dump, and we both agree that we don't know enough about how our brains work in order to transfer key portions. Ignoring the issue of whether this is possible, the remaining unanswered question is, if we could do this, would "mind" be a property of the system?

References:

[1] Computing Machinery and Intelligence. Mind 49 - pp. 433-460.

[2] Practical Thinking – De Bono

[3] The Singularity is Near –Kurzweil

[4] Searle, J. R. (1980a), `Minds, Brains, and Programs', Behavioral and Brain Sciences 3, pp. 417-424.

[5] The Chinese Room Argument – The Internet Encyclopedia of Philosophy
http://www.iep.utm.edu/c/chineser.htm#SH4a

[6] The Emperors New Mind – Concerning Minds and the Laws of Physics,
Roger Penrose, Oxford University Press, 1989.

[7] Shadows of the Mind - A Search for the Missing Science of Consciousness
Roger Penrose, Oxford University Press, 1994.

[8] Information – The New Language of Science – Hans Christian von Bayer pp. 227

[9] Matter and Consciousness – Paul. M Churchland (Revised Edition) pp.173

The author, W. Hugh Chatfield, although he holds an honours degree in Physics has worked his entire life in Information Technology. Since 1993, he has been involved with general markup technology including SGML, XML and the progress towards the Semantic Web. He currently is the principal of CyberSpace Industries 2000 Inc. specializing in XML Training and Consulting. Of particular personal interest is the advances being made in the intersection between the fields of Physics and Information Technology. [Further information can be found at his web site http://www.urbanmarket.com/csi2000]