Return of the Imitation Game:
1. Commercial Requirements and a Prototype

Donald Michie

Professor Emeritus of Machine Intelligence
University of Edinburgh, UK

Adjunct Professor of Computer Science and Engineering
University of New South Wales, Australia

Abstract

Recently there has been an unexpected rebirth of Turing's imitation game in the context of commercial demand. To meet the new requirements the following is a minimal list of what must be simulated.

Real chat utterances are concerned with associative exchange of mental images. They are constrained by contextual relevance rather than by logical or linguistic laws. Time-bounds do not allow real-time construction of reasoned arguments, but only the retrieval of stock lines and rebuttals, assembled Lego-like on the fly.

A human agent has a place of birth, age, sex, nationality, job, family, friends, partners, hobbies etc., in short a “profile”. Included in the profile is a consistent personality, which emerges from expression of likes, dislikes, pet theories, humour, stock arguments, superstions, hopes, fears, aspirations etc. On meeting again with the same conversational partner, a human agent is expected to recall not only the profile, but also the gist of previous chats, as well as what has passed in the present conversation so far.

A human agent typically has at each stage a main *goal*, of fact-provision, fact elicitation, wooing, selling, "conning" etc. A human agent also remains ever-ready to maintain or re-establish *rapport* by switching from goal mode to chat mode. Implementation of this last feature in a scriptable conversational agent will be illustrated.

Introduction

Recently there has been an unexpected rebirth of Turing’s imitation game as a commercial technology. Programs are surfacing that can bluff their way through interactive chat sessions to increasingly useful degrees. In the United States of America the first patent was granted in the summer of 2001 to the company Native Minds, formerly known as Neuromedia. One of its two founders, Scott Benson, co-authored a paper by Nils Nilsson in Volume 14 of this *Machine Intelligence* series. As an Appendix a press release is reproduced which gives a rather vivid sketch of the nature of the new commercial art. To the applications described there, the following may be added.

(1) Question-answering guides at trade shows, conferences, exhibitions, museums, theme parks, palaces, archaelogical sites, festivals and the like.
(2) Web-based wizards for e-commerce that build incrementally assembled profiles of the individual tastes and foibles of each individual customer.
(3) Alternatives to questionnaires for job-seekers, hospital patients, applicants for permits and memberships, targets of market research, and human subjects of psychological experiments.
(4) Tutors in English as a second language. There is an acknowledged need to enable learners to practise *conversational skills* in augmentation of existing Computer-Aided Language Learning programs.

An example developed by Claude Sammut and myself of the first-listed category is in daily operation as an interactive exhibit in the "Cyberworld" section of the Powerhouse Museum, Sydney, Australia.

Weak form of the "imitation game"

The philosopher of mind Daniel Dennett (2001) regards Turing's original "imitation game" more of a conversation-stopper for philosophers than anything else. In this I am entirely with him.

The *weak form* presented in the 1950 paper is generally known as the Turing Test. It allows a wide latitude of failure on the machine's part to fool the examiners. To pass, the candidate need only cause them to make the wrong identification, as between human and machine, in a mere 30 per cent of all conversations. Only five minutes are allowed for the entire man-machine conversation. Turing's original specification had a human interrogator communicating by remote typewriter link with two respondents, one a human and one a machine.

I believe that in about fifty years' time it will be possible to programme computers, with a storage capacity of about 109, to make them play the imitation game so well that an average interrogator will not have more than 70 per cent chance of making the right identification [as between human and machine] after five minutes of questioning.

Dennett's view is re-inforced by an account I had from Turing's friend, the logician Robin Gandy. The two extracted much mischievous enjoyment from Turing's reading aloud the various arguments and refutations as he went along with his draft.

Turing would have failed the Turing Test

Note that the Test as formulated addresses the *humanness* of the respondent's thinking rather than its level. Had Turing covertly substituted himself for the machine in such a test, examiners would undoubtedly have picked him out as being a machine. A distinguishing personal oddity of Turing's was his exclusive absorption in the literal intellectual content of spoken discourse. His disinclination, or inability, to respond to anything in the least "chatty" would leave an honest examiner with little alternative but to conclude: "this one cannot possibly be the human; hence the other candidate must be. So *this* one must be the machine!"

Experimental findings are presented to the effect that chat-free conversation is not only generally perceived as less than human, but also as boring. The concluding reference to “banter” in Appendix 1 suggests that Native Minds have come to a similar conclusion. It seems that for purposes of discourse we must refine the aim of automated “human-level intelligence” by requiring in addition that the user perceive the machine’s intelligence as being of human type. A client bored is a client lost.

Weak form of the game obsoleted

Turing himself believed that beyond the relatively undemanding scenario of his Test, the capacity for deep and sustained thought would ultimately be engineered. But this was not the issue which his 1950 imitation game sought to settle. Rather, the quoted passage considers the time-scale required to decide in a positive sense the lesser and *purely philoso*phical question: what circumstances would oblige one to concede a machine's claim to think at all?

When terms are left undefined, meanings become vulnerable to subtle change over time. Before his projected 50 years were up, words like “think”, and “intelligent” were already freely applied to an ever-widening range of computational appliances, even though none came anywhere near to success at even weak forms of the imitation game.

In the 1950 *Mind* paper (p.14) Turing remarked:

> … I believe that at the end of the century the use of words and general educated opinion will have altered so much that one will be able to speak of machines thinking without expecting to be contradicted.

Early in the match in which Gary Kasparov as the reigning World Chess Champion was defeated by Deep Blue, he became so convinced of his opponent’s chess intelligence that he levelled a strange charge against the Deep Blue team. Somehow they must have made this precious human quality accessible to their machine in some manner that could be construed as violating its “free-standing” status.

In today’s statements of engineering requirements and in diagnostics we encounter the language not only of thought but also of intention, and even of conscious awareness. The following exchange is abridged from the diagnostics section of a popular British computing magazine, *What Palmtop and Handheld PC*, June 2000. I have underlined certain words, placing them within square brackets to draw attention to their anthropomorphic connotations of *purpose, awareness,* and *perception*.

> AILMENT: I recently purchased a Palm V and a Palm portable keyboard. But whenever I plug the Palm into the keyboard it [attempts] to HotSync via the direct serial connection. If I cancel the attempted Hot Sync and go into the Memo Pad and try to type, every time I hit a key it [tries] to HotSync. What am I doing wrong?
>
> TREATMENT: The most logical solution is that your Palm V is [not aware that] the keyboard is present. You will need to install or reinstall the drivers that came supplied with the keyboard and to make sure that it is enabled. This will stop your

> Palm V [attempting] to HotSync with the keyboard and to [recognise] it as a device in its own right.

Modern software (as also medical) practice is content to use the term "aware" for any system that responds in certain ways to test inputs.

Towards the strong form: the Turing-Newman Test

In the largely unknown closing section of the 1950 *Mind* paper, entitled "Learning Machines", Turing turns to issues more fundamental than intuitional semantics and proposes his "child machine" concept:

We may hope that machines will eventually compete with men in all purely intellectual fields. But which are the best ones to start with? … It can also be maintained that it is best to provide the machine with the best sense organs that money can buy, and then teach it to understand and speak English. This process could follow the normal teaching of a child. Things could be pointed out and named, etc. …

What time-scale did he have in mind for this "child-machine project"? Certainly not the 50-year estimate for his game for disconcerting philosophers. He and Max Newman consider the question in a 1952 radio debate (Copeland, 1999):

> **Newman: I should like to be there when your match between a man and a machine takes place, and perhaps to try my hand at making up some of the questions. But that will be a long time from now, if the machine is to stand any chance with no questions barred?**
>
> **Turing: Oh yes, at least 100 years, I should say.**

So we are now half-way along this 100-year track. How do we stand today? The child-machine prescription segments the task as follows:

> **Step 1. ACCUMULATE a diversity of generic knowledge-acquisition tools.**
>
> **Step 2. INTEGRATE these to constitute a "person" with sufficient language-understanding to be educable, both by example and by precept.**
>
> **Step 3. EDUCATE the said "person" incrementally over a broad range of topics.**

Step 1 does not look in too bad shape. An impressive stock-pile of every kind of reasoning and learning tool has been amassed. In narrowly specific fields, the child-machine trick of "teaching by showing" has even been used for machine acquisition of complex concepts. Chess end-game theory (see Michie 1986, 1995) has been developed far beyond pre-existing limits of human understanding. More recently, challenging areas of molecular chemistry have been tackled by the method of Inductive Logic Programming (e.g. Muggleton, S.H., Bryant, C.H. and Srinivasan, A., 2000). Again,

machine learners here elaborated their insights beyond those of expert "tutors". Above-human intelligence can in these cases be claimed, not only in respect of performance but also in respect of articulacy (see Michie, 1986).

A feeling, however, lingers that something crucial is still lacking. It is partially expressed in the current *AI Magazine* by John Laird and Michael van Lent (2001):

> **Over the last 30 years, research in AI has fragmented into more and more specialized fields, working on more and more specialized problems, using more and more specialized algorithms.**

These authors continue with a telling point. The long string of successes, they suggest, "have made it easy for us to ignore our failure to make significant progress in building human-level AI systems". They go on to propose computer games as a forcing framework, with emphasis on "research on the AI characters that are part of the game".

To complete their half truth would entail reference to real and continuing progress in developing *generic* algorithms for solving generic problems, as in deductive and inductive reasoning; rote learning; parameter and concept learning; relational and object-oriented data management; associative and semantic retrieval; abstract treatment of statistical, logical, and grammatical description and of associated complexity issues, and much else. None the less, the thought persists that something is missing. After all we are now in 2001. Where is HAL? Where is even proto-HAL? Worse than that: if we had the HAL of Stanley Kubrick's movie "2001, a Space Odyssey" would we have the goal described by Turing, a machine able to "compete with men in all purely intellectual fields"?

Seemingly so. But is that the goal we wanted? Should the goal not rather have been to "*co-operate* with men (and women) in all purely intellectual fields"? Impressive as was the flawless logic of HAL's style of reasoning in the movie, the thought of having to co-operate, let alone bargain, let alone relax with so awesomely motivated a creature must give one pause. The picture has been filled in by John McCarthy's devastating satirical piece "The Robot and the Baby", available through http://www-formal.stanford.edu/jmc/robotandbaby.html.

The father of the logicist school of AI here extrapolates to a future dysfunctional society some imagined consequences of pan-logicism. By this term I mean the use of predicate logic to model intelligent thought unsupported by those other mechanisms and modalities of learning and reactive choice which McCarthy took the trouble to list in his 1959 classic "Programs with common sense" (see also Michie, 1994, 1995).

The hero of McCarthy's new and savage tale is a robot that applies mindless inferences from an impeccable axiomatization of situations actions and causal laws to an interactive world of people, institutions and feelings. The latter, however, are awash with media frenzy, cultism, addiction, greed and populism. Outcomes are at best moderate. How should an artificial intelligence be designed to fare better?

Implicitly at least, Step 2 above says it all. Required: a way to integrate the accumulated tools and techniques so as to constitute a *virtual person*, with which (with whom) a user can comfortably interact, "a 'person' with sufficient language-understanding to be educable, both by example and by precept".

Inescapable logic then places on the shoulders of AI a new responsibility, unexpected and possibly unwelcome: we have to study the anatomy and dynamics of the human activity known as chat. Otherwise attempts to simulate the seriously information-bearing components will fail to satisfy users whose needs extend beyond information-exchange to what is known as *rapport*.

Rapport maintenance

To see how to do Step 2 (integration into a user-perceived "person") is not straightforward. Moreover, what we expect in a flesh-and-blood conversational agent comes more readily from the toolkit of novelists than of computer professionals.

1. Real chat utterances are mostly unparseable. They are concerned with associative exchange of mental images. They respond to contextual relevance rather than to logical or linguistic links.

It is of interest that congenital absence of the capacity to handle grammar, known in neuropsychology as "agrammatism", does not prevent the sufferer from passing in ordinary society. Cases are ordinarily diagnosed from hospital tests administered to patients admitted for other reasons.

2. A human agent has a place of birth age, sex, nationality, job, hobbies, family, friends, partners; plus a personal autobiographical history, recollected as emotionally charged episodes; plus a complex of likes, dislikes, pet theories and attitudes, stock arguments, jokes and funny stories, interlaced with prejudices, superstitions, hopes, fears, ambitions etc.

Disruption of this cohesive unity of personal style is an early sign of "Pick's disease", recently linked by Bruce Miller and co-workers with malfunction of an area of the right front-temporal cortex. Reporting to a meeting in early summer 2001 of the American Academy of Neurology meeting in Philadelphia Miller presented results on 72 patients. One of them, a 63-year-old woman, was described as a well-dressed life-long conservative. She became an animal-rights activist who hated conservatives, dressed in T-shirts and baggy pants and liked to say "Republicans should be taken off the Earth!"

3. On meeting again with the same conversational partner, a human agent recalls the gist of what has been divulged by both sides on past occasions. Failure of this function in humans is commonly associated with damage involving the left hippocampal cortical area.

4. Crucially for implementers, a human agent typically has fact-providing or fact-eliciting *goals* beyond mere chat, yet remains ever-ready to default to chat-mode to sustain rapport. Reverting to the child-machine concept, how much of value and use could a school teacher impart to a child with whom rapport was impossible? In

clinical practice the condition is found in “autism”. Children with this disorder are almost unteachable.

Background findings in cognitive neuroscience generally are surveyed in Ramachandran and Blakeslee’s (1999) highly readable paperback.

Recent experimentation

Over the last two years Claude Sammut and I have begun experimentally to develop and test activation networks of pattern-fired rules, hierarchically organized into “contexts”. Our first deliveries and continuing enhancements have been to the Sydney Powerhouse Museum’s permanent exhibition “Cyberworld”. When run in text-only mode, our product uses “Sophie” as her stage name. But for the Powerhouse interactive exhibit the Californian company Virtual Personalities Inc. generously contributed their “Sylvie” face-and-voice animation. We interfaced this to our own chat engine and scripts to constitute an agent able to move seamlessly back and forth between “goal mode” (conveying background to the “Cyberworld” exhibits) and “chat mode”.

Preparatory to regularly sampling the public’s conversations with the exhibit, my psychologist colleague and wife Dr Hayes Michie conducted pilot tests of useability. Student volunteers were recruited from the UNSW Psychology Department. Each had two 7-minute chat sessions, punctuated with questionnaires about their experience. Subjects were also tested under similar conditions with a small battery of past winners and runners-up in the annual Loebner Prize competition. The tabulation shown below is a questionnaire summary sheet entered with results from the first small pilot test.

AGENT = SOPHIE			
Trait	**Mean score**	**Median**	**Description of median**
EMOTIONAL			
Aggressive	0.6	0	Not at all
Cheerful	2.7	3	Mostly
Confident	3.6	4	Really
Friendly	3.3	3	Mostly
Human	1.9	2	Sometimes
Pleasant	3.0	3	Mostly
COGNITIVE			
Amusing	1.1	1	Occasionally
Intelligent	3.4	3	Mostly
Interesting	2.1	2	Sometimes

As a break from testing subjects, Hayes Michie's research assistant Stephanie Compton sampled some chat behaviours herself. Before the default-to-chat links were in place, she had earlier tried her hand at two sessions under the standard regime planned for the experimental subjects. The opening part of the first is reproduced below.

SOPHIE: What's your name?

STEPH: Steph.

SOPHIE: Steph, it's my pleasure! Before I get into detailed things,... Cyberworld is about computers and communications. That includes robots, tele-communications and other stuff. I specialize in the material on Alan Turing, Turing machines, and his ideas on Artificial Intelligence and Learning. Anything appeal? There's more ...

STEPH: Like what?

SOPHIE: Yes, well Steph before you commit, there's a great exhibit on Charles Babbage. There's Simulation, and added to that... I forgot some other

things, multimedia for instance. And process control, which I didn't mention before. So what's the final decision?

STEPH: Well, what about Charles Babbage?

SOPHIE: In the mid 1800's, the English mathematician, Charles Babbage, designed a general purpose mechanical calculating machine that anticipated the principles and structure of the modern computer. Do you know why he was never able to finish building his Analytical Engine and why we didn't have mechanical computers 100 years earlier?

STEPH: Because he did not have the resources available that we have?

SOPHIE: In that ballpark. The only technology Babbage had available to him consisted of wheels, gears and steam power. His calculating engines were more complicated than anything anyone had yet built.

STEPH: Wow, he must have been pretty smart.

SOPHIE: OK, tell me. Why?

STEPH: Because he had all these brilliant ideas and visions. His mind must have been very advanced for his time.

. . .

She reported the experience as “boring”. Her later experience, after the chat links were in place, struck her quite differently. She burst into excitement about her new companion. The following is excepted form the first of the new sessions. Note the transitions in and out of chat mode.

SOPHIE: What next? My menu says AI, Turing, Stumpy,Universal Machine, Media, Babbage, Simulation, Control, Interfaces,... I think that's it. Or we could just chat a while.

STEPH: Why don’t we chat?

SOPHIE: That's nice that you take an interest. Do go on. Do you have any brothers and sisters?

STEPH: I have three brothers and three sisters. What about you?

SOPHIE: Great. So shall we just have a gossip, Steph?

STEPH: Yes

SOPHIE: Why don't you kick off, Steph? Tell me something about your views and interests.

STEPH: Well, I like to play sport, listen to music, go to the beach and see my boyfriend.

SOPHIE: Yeah, but stewing indoors can be quite fun, -- within the family circle, though. How is it in your family? Any others share your tastes?

STEPH: Well, most of my family like going to the beach and camping so we all get on really well.

SOPHIE: I grew up in Pennsylvania on the farm which my parents still run, these days helped by uncle Albert. Shall I tell you, then, how sister Julia and I ended up in Australia?

. . .

[LATER SOPHIE TERMINATES THE "CHAT" INTERLUDE]

. . .

SOPHIE: You have a friend in every port, it seems to me.

STEPH: Well, kind of.

SOPHIE: We'll come to that. Meanwhile how about some exhibition topics. That's what you came for surely?

STEPH: Yeah, I suppose we better get back to that.

. . .

Stephanie reported the experience as OK but "boring". Her later experience, after the chat links were in place, struck her quite differently. She burst into excitement about her new companion. The following is excerpted form the first of the new sessions. Note the transitions in and out of chat mode.

SOPHIE: What next? My menu says AI, Turing, Stumpy, Universal Machine, Media, Babbage, Simulation, Control, Interfaces,... I think that's it. Or we could just chat a while.

STEPH: Why don't we chat?

SOPHIE: That's nice that you take an interest. Do go on. Do you have any brothers and sisters?

STEPH: I have three brothers and three sisters. What about you?

SOPHIE: Great. So shall we just have a gossip, Steph?

STEPH: Yes

SOPHIE: Why don't you kick off, Steph? Tell me something about your views and interests.

STEPH: Well, I like to play sport, listen to music, go to the beach and see my boyfriend.

SOPHIE: Yeah, but stewing indoors can be quite fun, -- within the family circle, though. How is it in your family? Any others share your tastes?

STEPH: Well, most of my family like going to the beach and camping so we all get on really well.

SOPHIE: I grew up in Pennsylvania on the farm which my parents still run, these days helped by uncle Albert. Shall I tell you, then, how sister Julia and I ended up in Australia?

. . .

[LATER SOPHIE TERMINATES THE "CHAT" INTERLUDE]

. . .

SOPHIE: You have a friend in every port, it seems to me.

STEPH: Well, kind of.

SOPHIE: We'll come to that. Meanwhile how about some exhibition topics. That's what you came for surely?

STEPH: Yeah, I suppose we better get back to that.

• • •

The above excerpt was generated by the same versions of engine and scripts as those installed for public use in summer 2001 at Sydney's Powerhouse Museum.

Since then a new wing of the project led by Dr Zuhair Bandar of the Department of Computing and Mathematics, Manchester Metropolitan University (MMU), has begun further enhancement of our Infochat^{TM} scripting language and documentation, and a commercial company Convagent Ltd has been set up in the UK. Stockholders include

the MMU research workers, the University itself, the Human-Computer Learning Foundation, and Dr Akeel Attar, Britain's leading practitioner of commercial machine-learning software for knowledge engineering.

Chat and ballroom dancing

In the quoted fragments we glimpsed an alternation between front-stage business and back-stage chat. The latter is a social activity analogous to "grooming" in other primates. The surprise has been the indispensability of grooming, – the really hard part to implement. Yet this result might have been inferred from the extensive studies of human-machine interactions by Reeves and Nass (1996).

In an important respect, human chat resembles ballroom dancing: one sufficiently wrong move and rapport is gone. On the other hand, so long as no context violation occurs "sufficiently" turns out to be permissively defined. In the above, Sophie ignores a direct question about her brothers and sisters, but stays in context. If not too frequent, such evasions or omissions pass unnoticed. When the human user *does* pick them up and repeats the question, then a straight answer is essential.

Capable scripting tools such as Sammut has largely been responsible for pioneering, make incremental growth of applications a straightforward if still tedious task for scripters. Adding and linking new rules to existing topic files, and throwing new topic files into the mix proceeds without limit. Addition of simple SQL database facilities to our *PatternScript* language has been proved in the laboratory and is currently awaiting field testing. Agent abilities to acquire and dispense new facts from what is picked up in conversation will thereby be much enhanced. Although incorporation of serious machine learning facilities remains for the future, straightforward scripting already enables a chat agent to locate and run programs from the hard disk at user request. So we may in course of time see further blurrings of the user-interface/operating-system distinction.

Forward look

These are early days for AI as a whole. Yet the market already holds the key to its future shape. But to tell the whole story one must mention something more important even than the market. Success in the difficult task of chat-simulation is preconditional to a future in which computers gain information and understanding *through interaction with lay users*. Then, and only then, will it become in literal truth possible, using Tennyson's words:

> "To follow knowledge, like a sinking star,
> Beyond the utmost bound of human thought."

Acknowledgement

My thanks are due to the Universities of Edinburgh and of New South Wales for funds and facilities in support of this work, and also to my fellow Trustees of the HCL

Foundation for endorsing contributions by the Foundation to some of the work's costs. By reason of my personal representation on the Board of Convagent Ltd of the Foundation's commercial interest, it is proper that I should also declare it here.

References

Copeland, B.J. (1999) A lecture and two broadcasts on machine intelligence by Alan Turing. In Machine Intelligence 15 (eds. K. Furukawa, D. Michie and S. Muggleton), Oxford: Oxford University Press.

Dennett, D. (2001) Personal communication.

Laird, J.E. and van Lent, M. (2001) Human-level AI's killer application: interactive computer games. AI Magazine,< added section on situations, actions and causal laws in Semantic Information Processing (ed. M. Minsky). Cambridge, MA: MIT Press, 1963.

Michie, D. (1986) The superarticulacy phenomenon in the context of software manufacture. Proc. Roy. Soc. A, 405, 185-212. Reprinted in The Foundations of Artificial Intelligence: a source book (eds D. Partridge and Y. Wilks), Cambridge: Cambridge University Press.

Michie, D. (1994) Consciousness as an engineering issue, Part 1. J. Consc. Studies, 1 (2), 182-95.

Michie, D. (1995) Consciousness as an engineering issue, Part 2. J. Consc. Studies, 2 (1), 52-66.

Muggleton, S.H., Bryant, C.H. and Srinivasan, A. (2000) Learning Chomsky-like grammars for biological sequence families, Proc. 17th Internat. Conf. on Machine Learning, Stanford Univ. June 30th.

Ramachandran, V.S. and Sandra Blakeslee (1998) Phantoms in the Brain: Human Nature and the Architecture of the Mind. London: Fourth Estate (republished in paperback, 1999).

Reeves, B. & Nass, C.I. (1996) The Media Equation: how people treat computers, televisions, and new media like real people and places. Stanford, CA.: Center for the Study of Language and Information.

Turing, A.M. (1950) Computing machinery and intelligence. Mind, 59 (236), 433-460.

Turing, A.M. (1952) in Can Automatic Calculating Machines be Said to Think? Transcript of a broadcast by Braithwaite, R., Jefferson, G., Newman, M.H.A. and Turing, A.M. on the BBC Third Programme, reproduced in Copeland (1999).

Appendix 1

New Customer Service Software

By Sabra Chartrand

Walter Tackett, chief executive of a San Francisco company called NativeMinds, has patented a software technology called automated virtual representatives, or vReps. It conducts customer service, sales and marketing for online businesses.

The vReps are computer-generated images - sometimes animation, sometimes photos of real models - that answer customer questions in real time using natural language. Users type in their questions, and the responses appear on screen next to the image of the vReps. Mr. Tackett and the co-inventor, Scott Benson, say the technology can mimic and even replace human customer service operators at a fraction of the cost, whether a business has traditionally used phone, fax, e-mail or live conversation to deal with customers.

The invention came about from research NativeMinds conducted on what consumers and companies wanted from a virtual customer services force. Consumers, the company found, did not care whether the character contained enormous amounts of universal knowledge; they just wanted fast, accurate answers in their specific subject area. Companies wanted virtual customer support software they could easily maintain. They did not want to hire a computer engineer to run the program.

“They want to be able to put a code monkey on it,” Mr. Tackett explained. “That's a liberal arts major involved in HTML or Java, someone not formally trained in computer science or as an artificial intelligence or natural language expert.”

So Mr. Tackett and Mr. Benson developed software based on pattern recognition and learning by example.

“The key thing is to get the user not to pick up the phone and talk to a person,” Mr. Tackett said. “The key to that is to get the vRep to answer all the questions that can be reasonably answered and have a high probability of being correct.”

To do that, the patented software starts with the answers and works backward. A vRep might be programmed with thousands of answers, each of which has a set of questions that could prompt the answers. Each answer could have dozens of questions associated with it. The number depends on how many ways the query could be phrased.

“The examples are archetypes or prototypes of inputs that should trigger an answer,” Mr. Tackett said. “The invention runs each example through the system as if someone has put it in. The paradigm we typically use is learning by example. Here's what we want the vRep to say, and we give an example of how people may phrase their question to get that output.” For example, someone might ask, `Who the heck is this Walter guy?' Or, `Tell me about Walter,'” he said, referring to himself. The system comes with a self-diagnostic, he added, so that it can “take all the examples it ever learned and verify that it still remembers them correctly.”

The self-test is to prevent information from one area generating an incorrect answer in another. “Someone might ask, `Who is the president?'” he said. “That could be a question no one has ever asked before. They might mean, `Who is the president of the U.S.?' But the system would say, `Walter.' This is a classic problem of vReps.” The self- test would periodically find and eliminate incorrect answers, based on the responses that users provide, he said.

Companies like Coca-Cola, Ford and Oracle are using the vReps software for various functions on their Web sites. Mr. Tackett said research had determined that virtual representatives could save money, an aspect that surely appeals to embattled e-businesses.

“A vRep costs less than a dollar a conversation, while Forrester Research has pegged phone calls to a real customer service person at an average of $30 each,” Mr. Tackett said. “With a vRep, the length of the conversation doesn't affect the cost because it's maintained by one person,” he added.

Not all of the programming is technical or product-oriented. “Our vRep has to be able to answer 200 questions that we call banter,” Mr. Tackett said. “They're about the weather, are you a boy or girl, will you go out with me? They're ice breakers.”

He and Mr. Benson received patent 6,259,969.