The Turing Test

by Bruno W. Augenstein

Stunned silence.

That was the initial reaction of the Wightman Company's Advisory Board. Members had been called in for a special meeting on Wightman's android project. Now the advisors were convinced of the reality of that project - but what had convinced them was totally surprising. The board members took some moments to recover from the statements just made by the projects Principal Integrator, or PI.

The next hour in the board room was hectic. Responses to the PI's concluding comments, after the astonishment which flooded over the board was absorbed, ranged from the commonplace to the practical. A few of the members succumbed to the well-known psychological phenomenon of believing that now, knowing the object of their study to be an android, they recalled after the fact some "prior feelings" that this might be so - but couldn't quite put their finger on the reasons. Some members first felt tricked; reflection quickly dispatched that emotion.

More important actions also resulted. The young physicist who earlier had proposed the consciousness tests went over some technical details of the tests with me and the PI. The PI was increasingly intrigued. The tests, when performed, wound up creating more new questions - but isn't that itself a characteristic of the "human" experience?

While this was going on, Jon Wightman and the Executive Committee of his board huddled, drafting the announcements to be made the next morning. Two prime goals had to be accommodated. The technological revolution embodied in the android had to be made startling and clear to the general populace. The result Wightman hoped for was a substantial increase in the value of Wightman stock. At the same time, the announcements had to reflect the epochal social and cultural implications of the emergence of androids with so many human-like qualities. One of the outcomes of the huddle was prominently reflected in the next day's announcements. The android was given a name, and was not simply identified as "the Wightman android," or by any other neutral term.

You've all read these terse announcements. They didn't convey the high drama of the afternoon. Nor have the other parties to the board meeting ever publicly elaborated on their session. But I believe that two landmark events happened that afternoon which deserve special emphasis in the history of computing. One was the disclosure of a technological revolution; the other was the singular and elegant Turing test so essential to that disclosure.

As it happened, the announcements achieved both goals. Giving a name to the android, with the implications that act carried, was a prescient move. While we have by no means resolved all the detailed legal and policy issues of defining the status of androids, we have at least avoided some of the philosophical baggage which could have emerged to encumber us. The android is regarded as an entity holding some certain rights. The humans he is associated with are seen legally as guardians, with well-defined responsibilities, and not as holders of mechanical slaves. What will emerge ultimately from this relationship time only will tell, but a sensible and inspired start has been made.

Wightman's plan to increase the value of Wightman stock also succeeded nobly. After two weeks of trading, Wightman stock was the third highest in value on the exchanges, and an era of high growth and profitability followed for Wightman.

There were always those analysts who could not correlate the Wightman profits with the growth of Wightman product lines. The mystery too was very recently cleared up, by the announcement of the third Wightman revolution - the final stages of development of a star ship prototype. Here Wightman's single-minded intensity succeeded, where our government's will, planning scope, and bureaucratic suppleness were all inadequate.

But this story shall be told from the beginning. You can then sense for yourselves the uniqueness of the events as they unfolded.

I am Chris Chric, technical editor of the major professional journal (Computing Progress) concerned with new projects and products of the still very accomplished technologists of Silicon Valley. Of course, today as much biocomponentry as silicon, Josephson junctions, or newer quantum devices is used in the valley, so its name is a bit archaic now. I'd gravitated to that publishing niche of to an academic of industrial career. The built-in exposure to new developments as they were happening attracted me.

I had graduated from Cal Tech just after the turn of the century, with a degree in computer science. Cal Tech was (and remains) idiosyncratic in its approach to computer science - where else would you find some of the better computer courses given in the Biology department? That practice, started in Feynman's time, still produces a few special graduates every so often. Their prime interests lie in mid-brain issues and in the realities underlying the brain-computer metaphors. I was one of these graduates. Many people today think that the presence of component androids has blurred these questions. But from my perspective the real issues have simply gotten harder and more subtle.

One of my friends and classmates at Cal Tech had gone on to become VP for Research at Wightman. It was through him that I was nominated to the Wightman board. I guess he felt my interests would give Wightman a good sounding board. Besides, my nomination was well timed for the disclosure of the android work at Wightman. It was at my first board meeting that the events I'm going to relate happened.

I was present at a full board meeting in the afternoon at Wightman, when what was judged by the board to be a truly revolutionary development was unveiled for the first time. I knew from my journal activities that Wightman, for its most important and closely held developments, used a form of the Skunk Works organization so well exploited by Lockheed in the previous century. Such an organization could protect its activities and streamline management, while getting the best out of its people - an ideal match for bringing about revolutionary progress shielded from unwelcome attention. For Wightman to have used such a procedure to develop something unknown to all but a very few seemed completely in character.

About a week before, the Wightman board was asked to attend a special meeting, set for that afternoon and the next morning. Members were told only that a new research development with profound implications would be unveiled, and that they would be asked, by a form of testing, to access the development. Naturally, this piqued the interest of the board (just two members had been briefed all along). In the two decade history of Wightman, such an invitation had been made only once before, and the unveiling - this was before my time - had fully lived up to its claimed importance. So that afternoon the full 20-man board had convened in the Wightman board room, ready to be critical, but eager to see a new revolution. The Wightman board was an impressive group, containing ex-Secretaries of Defense, retired presidential Science Advisors, Nobelists, heads of prestigious universities, a few young academics already accumulating vast reputations, and others. The group formed a tough audience for proclaimed revolutions. I was part of the "show me" component of the board. I'd written about too many fads and fizzles at computing's forefronts to be otherwise.

At 1:30 in the afternoon the Wightman board and their senior executives - including the CEO Jon William Wightman himself, the legendary innovator - got the meeting underway. Half a dozen members of Wightman's research arm were present, and my friend the Research VP and two of his senior staff made presentations. The gist of these was simple. The VP said the new product was a very competent android capable of emulating many human actions. You can imagine the board's surprised response to this.

The VP then first introduced the Project Manager, who detailed the composition and administration of the large team involved in the android development, mentioning funding levels impressive even to those board members used to dealing with large sums. After the Project Manager's presentation, and a few questions from the board, the VP next introduced a young man identified as the Project Integrator, or PI, on whom the bulk of the technical discussion was to rest. Most of the afternoon was in fact taken up by the Project Integrator in describing both the nature of the android project, and the board's assessment functions to evaluate the success of the project.

Splitting a very large, challenging project in this way, the development load shared between a manager and an integrator, was standard at Wightman. In one of his rare public speeches, Jon Wightman had put it this way:

"I choose the PI to be the focus of all technical efforts of the project. I choose the Project Manager to shield the PI from everything but those technical issues. That way neither the manager nor the integrator has to dilute his special competence. There is very little that can upset a leading edge technology development as much as an administrator trying to make technical judgments, or a technologist trying to run the day-by-day operations of a very large team. Kelly Johnson learned this at Lockheed nearly a century ago; it's as relevant today as it was then."

From the moment the Integrator took charge of the meeting, his broad technical competence was apparent to all of us on the board. The Research VP in his opening comments, backed forcefully by Jon Wightman, emphasized the importance of the board being wholly familiar with the background of the project before the android was "previewed."

The PI, said the VP, had the familiarization task as his afternoon's prime responsibility. I'll give you just some of the highlights of the subsequent interplay between the PI and the board.

The young man first defined the goals of the android project. These were straightforward - the creation of a human-looking machine with cognitive, perceptual, and motor skills emulating human capabilities.

One of the board's Nobelists asked, "Why give the android a completely human appearance?"

The Project Integrator replied, "Almost all of the uses we see for such androids involve working with humans. We want humans to be comfortable with androids, and non-human features would be a distraction. Great lengths are gone to in this humanizing goal. One development was the adaptation of Plasti-skin, our artificial skin for burn victims, as the outer envelope of the android. As you know, during Plasti-skin growth in the lab, hair can be implanted, giving a completely natural appearance. I'll be interested later in your comments on how well we've succeeded in this naturalness goal."

The board's questions to the Integrator were searching and detailed. One issue of great interest arose - the intellectual level of the android.

The Integrator commented, "The project's goal was to make the android very competent. This meant emulating quite well human perceptual and motor skills. Cognitive abilities, such as reasoning, learning, and especially memory, were keyed to achieving capabilities of clever and thoughtful humans. The project team believes a cognitive level has been achieved in the android which allows it to participate in cooperative activities as a full partner to humans. On the other hand, little progress was available to make the android's faculties in abstract conceptualization and devising of wholly new ideas in any sense remarkable. The android is a very competent Everyman in behavior, not a new class of Genius. The project goal was to provide a partner with flexibility and initiative, attractive to humans. This artificial partner was not expected to provide new insights on forefront problems of the human mind. No new results on basic unresolved problems of mathematics were realistically anticipatable. Nor was the project expecting from the android any new discoveries on the frontiers of physics, where many disciplines were for the first time cooperating in approaches to the possibility of interstellar travel at effectively above light speed, by space-time shortcuts, called "wormholes."

The Project Integrator concluded, "The project is bound by the same constraints everyone faces. In spite of a century of effort, we still don't know enough about how the human brain makes fundamental advances in the highest cognitive functions to form any basis for emulation."

At this point, I sensed a certain relaxing in demeanor by a few of the board members.

As you would expect, the Wightman board room had built-in all the latest technology. There were display panels around most of the room, where spoken directions generated vital statistical graphics in real time; spoken directions could also generate very complex computations whose end results were promptly available. The latter was particularly useful in responding to one of the Nobelists' questions on how the neural networks - which were an essential part of the android's computational power - were organized, and what process of solution was built into the networks. (A major feature was stochastic Hopfield networks, a culmination of five decades of development, much of it by the Wightman laboratories.)

But there were also several highly personal aspects to the afternoon's discussions. Old habits being hard to lose, even in the 21st century earlier briefing techniques persisted. The same board room display panels with electronic outputs could also simulate, via simple wands, the old process of informal blackboard presentations. The Integrator's briefing style often involved walking around the board room towards the board member who might have raised a question. There the two of them could hold a close colloquy at the "blackboard," to convince each other that the responses provided were appropriate to the questions raised. During the afternoon just about every board member had such an interaction with the Integrator. Each time, the two participants seemed to be in intimate discourse with one another, addressing each other single mindedly as though no other participants were present.

There was extensive discussion of the test -- a "Turing test" - Wightman wanted the board to carry out. By this time in the 21st century, the term, originally defined by the famous logician and computer scientist, Alan Turing, about the middle of the previous century, as the process for distinguishing between "computers" and "people," was a standard part of the vocabulary. But, the Project Integrator noted, the test in this instance was to be very comprehensive, to make as sure as possible that no significant major interfaces remained showing separations between humans and the android. That is, observation of the android performing motor, perceptual, and mental functions was (hopefully) not to reveal traits or behavior which would allow humans to say immediately, "Aha! This is a machine!"

The young Integrator focused on several aspects of the proposed Turing test.

In response to one of the board members' questions, he replied, "One of the critical Project issues was providing for the android what is termed human common sense. Research starting in the '90s of the last century defined the enormity of this problem. The numbers of descriptors of common sense, and the ability to draw further inferences from these descriptors, require for success the use of very striking computer power - power levels available only relatively recently. Success here means that we emulate the performance of competent humans in real time, with no evidence of uncharacteristic halting, hesitation, or hunting for appropriate outcomes. One of the important aspects of your Turing test will be to determine if any such evidence is apparent in your interactions with the android."

Another one of the board members - an eminent computer scientist - asked how this "common sense" was built into the android.

The PI responded at some length. "First, the Project Team asked if the repertoire of android capabilities could be built in by specifying every detail of the common sense required. The team quickly determined that this approach was unproductive, for a number of reasons now clear to see, although a few decades ago this was a popular approach. Instead, the team decided on a new approach modeled after the way humans acquire common sense - by experiences in real environments. But our problem was considerably more difficult - we could not afford the two and more decades of experiences which shape human behavior, and we had in effect to build multiple competencies into one android. The Project Team decided that an almost universal android was critically important. That meant building in the backgrounds of many disciplines - science, engineering, linguistics, psychology, and so on - and integrating them all into a base which allows each discipline to be exercised. This was a special task for me, and I can tell you that many difficulties had to be overcome. In the end, the Project Team devised an approach which succeeded very well, we believe - training is done in a greatly sped-up way in virtual reality environments. This permits learning appropriate disciplines and the relevant common sense by experiences in contexts where the guiding rules are very numerous and hard to prescribe in advance. The success of this approach should again be determined by your Turing test queries of the android. I want individual board members to pose to the android the same kinds of questions, both disciplinary and general, that we're exploring here, to test the naturalness of the android's performance."

Some of the other board members then raised questions in their disciplinary areas for the PI to field. This phase of the afternoon's activities took almost three hours, and the PI wound up by observing, "The training methods used to develop the android capabilities are a very critical and important part of the Project Team's results. We now feel for the first time we have realistic hopes that routine extensions of human capabilities are also possible. But you, the Wightman board, must satisfy yourselves that these conclusions of the Project Team are well-founded. That's why we are placing so much importance on the Turing test you conduct."

The PI then discussed with the board the possible uses for androids with the capabilities claimed by the Project Team.

Summarizing the board's observations, he said, "The Project Team felt initially that androids of this class could serve high priority functions in very hazardous tasks, in specialized educational settings, and as partners in tasks humans now undertake where continual attention and vigilance are at a premium. You have elaborated on these functions, suggesting special classified uses as commando, intelligence, and anti-terrorist teams. Other proposals include uses in entertainment, such as new kinds of sports competitions. Finally, it has been suggested that, in the long run, these current androids are the beginning of a parallel evolutionary track whose form we can see only dimly now. What we can foresee is that such advanced technology can never be divorced from its applications and societal impacts. Androids have now been developed to the stage where Wightman will be happy to send them out to interact with humans, without concerns that special legal and safety issues will arise. The intellectual and cultural consequences can be profound; we need to step very carefully to see that these are kept in balance with the benefits of android use in special situations."

The board agreed with these cautionary remarks.

Incidentally, we of course know, some years after the events, that Wightman was being a bit disingenuous in these discussions of android uses. The prime objective of Wightman even then was to create very capable crews for what Jon Wightman considered his ultimate technical priority - development of star ships for interstellar travel. Human biology, physiology, and even psychology still are barriers to the practicalities of human interstellar flight. Intensive research on life prolongation, cryobiology, and hibernation or suspended animation has to date fallen far short of safe use by human crews, hence we have only been able to populate but a few planets and moons in our own solar system with colonies of self-contained industrial and tech-based cities, albeit they have been quite successful and a boon to the overall economy.

Use of android crews first as the sole interstellar explorers, then possibly as attendants for generations of stored fertilized ova or embryos, raised and taught by the androids, and perhaps finally as full partners for human crews, performing essential continuing functions to ensure the human's safety and the expedition's success - this was Wightman's vision. Today, we have reasonable hopes for powerful interstellar propulsion, melding arcane aspects of super-string theory and gravitational waves to draw power from and to navigate across the space between stars, where there are possibly many beautiful planets like earth to discover and to live on and thrive as a symbiotic race of humans and their human-like partners, the androids. But that's an altogether different story, yet to be told in detail.

The last part of the afternoon's session touched on a difficult question raised by one of the board members and then pursued by others: "Is the android conscious, in the sense we are? Is there an awareness of self, a free will, an introspective capability which tells us our internal mental states, and which can rehearse the various qualities of subjective felt experiences and thoughts? Do you know what the nature of the android's 'being' really is?"

Wightman himself joined in the discussion, saying, "There can be no firm answer to this question now, because humans cannot satisfactorily define what is meant by consciousness."

This view was concurred in by many present. The PI summarized the Team's position on these questions.

"First," he said, "the android senses that there is a world external to him which he can affect by voluntary actions. Second, the android does not sleep in any human sense, and has no comparable dream state - a different state of consciousness, presumably, for humans, in which some poorly understood brain activities are carried out while sensory input and motor output are almost absent. Third, the android does not have set responses to external events, but carries out some form of internal optimization to define his reactions. Behavior generally conforming to the old Campbell-Asimov 'laws' results from this optimization, not from special circuitry. In a nutshell, the android appears to have responsibility, self-awareness, intention, and volition. Whether this kind of behavior is comparable to exercise of free will," the PI added, "I cannot say conclusively. We're treading here on unknown, and perhaps unknowable territory, just as is the case when we reach barriers in trying to emulate the highest cognitive functions of humans."

This afternoon's discussion would have ended on this somewhat uncertain note, if one of the rising young academics in theoretical physics had not spoken up.

"I have a modest proposal to make," said this physicist. "After over one hundred years of quantum mechanics, the interpretation of the quantum mechanical rules is still obscure to many of us. There is one large school of thought - von Neumann and Wigner were the original strong proponents in the last century - which believes that the wave function is not reduced to some event until, and unless, the experimental phenomena finally impinge on the consciousness of some observer. That is, for me, quantum descriptions are affected by impressions which enter my consciousness. Now, why don't we perform some of these experiments using the android as the observer?"

The physicist then described how such experiments might be carried out.

I analyzed the actual experiments conducted a few years later in my journal, Computing Progress. The results got the physicist the Nobel Prize, but the outcomes of the experiments, run in several variants since, remain fascinating, puzzling, and bizarre, and actually disturbing to many. But again - that's another story.

By then it was quite late. Over four and a half hours of briefings and discussions had taken place, the great bulk of these falling on the shoulders of the Project Integrator.

The board chairman noted the lateness, and asked his final conductive question (I learned later that he was one of the two fully briefed board members):

"Jon, do you want to postpone our observation of your android to tomorrow morning, and hold off our testing until then?"

Jon Wightman got up and scanned the board members.

Finally, he said, "I don't think it will be necessary to pursue this meeting on observing and testing our android until tomorrow."

In the puzzled quiet which followed, he added, "Our PI has a few final comments he wants to make."

The eyes of the board members turned back to the PI, who had remained standing during this exchange. The PI's next words should, I believe, be imprinted on the history of technology, ranked along with Bell's first call on the telephone, and with the clever inversion proposed by the first true artificial intelligence in response to a bystander's query - "Sum, ergo cognito." The lineage from this first intelligence to the Wightman android is now clear to all.

And what were the PI's words to us?

"Gentlemen, you have already conducted your Turing test."

Pause.

"You see, I am the android."

Stunned silence...

THE END