

EIGHTEENTH INTERNATIONAL JOINT CONFERENCE ON ARTIFICIAL INTELLIGENCE



IJCAI-03

PROCEEDINGS

ACAPULCO, MEXICO • AUGUST 9-15, 2003

What is Artificial Intelligence? Psychometric AI as an Answer

Selmer Bringsjord & Bettina Schimanski

selmer@rpi.edu • schimb@rpi.edu

Department of Computer Science

Department of Cognitive Science

Rensselaer AI & Reasoning (RAIR) Lab

Rensselaer Polytechnic Institute (RPI)

Troy NY 12180 USA

Abstract

We propose an answer to the “What is AI?” question, namely, that AI is really (or at least really ought in significant part to be) *Psychometric AI* (PAI). Along the way, we: set out and rebut five objections to PAI; describe PERI, a robot in our lab who exemplifies PAI; and briefly treat the future of Psychometric AI, first by pointing toward some promising PAI-based applications, and then by raising some of the “big” philosophical questions the success of Psychometric AI will raise.

1 Introduction

What exactly *is* AI? We’d be willing to wager that many of you have been asked this question — by colleagues, reporters, friends and family, and others. Even if by some fluke you’ve dodged the question, perhaps you’ve asked it yourself, maybe even perhaps (in secret moments, if you’re a practitioner) *to* yourself, without an immediate answer coming to mind. At any rate, AI *itself* repeatedly asks the question — as the first chapter of many AI textbooks reveals. In this paper we want to propose an answer, namely, that AI is really (or at least really ought in significant part to be) *Psychometric AI* (sometimes just ‘PAI’ (rhymes with π) for short). We also want to tell you something about both PERI, a robot in our lab who exemplifies PAI, and the future we envision for PAI.

Our plan herein is as follows. In the next section, 2, we answer the “What is AI?” question from the standpoint of Psychometric AI, and introduce some of the tests at the heart of this brand of AI. In section 3 we rebut some objections that will inevitably be brought against Psychometric AI. The rebuttal to the first of these objections will reveal the foundation for PAI: the Turing Test (TT) and its more demanding cousin, the Total TT (TTT). In section 4 we introduce you to PERI. Our penultimate section briefly treats the future of Psychometric AI, first by pointing toward some promising PAI-based applications, and then by raising some of the “big” philosophical questions the success of Psychometric AI will raise. We end by addressing a second round of objections, formulated by those who read earlier versions of the present paper.

2 What is AI? Psychometric AI as an Answer

Presumably the ‘A’ part of ‘AI’ isn’t the challenge: We seem to have a fairly good handle on what it means to say that something is an artifact, or artificial. (We can ignore here conundrums arising from self-reproducing systems, systems that evolve without human oversight, etc.) It’s the ‘I’ part that seems to throw us for a bit of a loop. What’s intelligence? *This* is the big, and hard, question. Innumerable answers have been given, but most thinkers seem to forget that there is a particularly clear and straightforward answer available, courtesy of the field that has sought to operationalize the concept in question; that field is psychometrics. Psychometrics is devoted to systematically measuring psychological properties, usually via tests. These properties include the one most important in the present context: intelligence. In a nutshell, then, the initial version of our account of intelligence is this: *Some agent is intelligent if and only if it excels at all established, validated tests of intelligence.* (This account is inadequate, for reasons we explain below before supplanting it with a more sophisticated one.) AI then reduces to Psychometric AI: the field devoted to building a computational system able to score well on such tests. This may strike you as a preposterously narrow definition of AI. The first step (in a series taken as this paper unfolds) in diffusing this attitude is to take a look at some intelligence tests, some of which, we surmise, are a good deal richer than you might at present think.




Figure 1: Sample Problem Solved by Evan’s (1968) ANALOGY Program. *A is to B as C is to ...?*

In the early days of AI, Psychometric AI was at least implicitly entertained. After all, in the mid 1960s, the largest Lisp program on earth was Evans' 1968 ANALOGY program, which could solve problems like those shown in Figure 1. Evans himself predicted that systems able to solve such problems would "be of great practical importance in the near future," and he pointed out that performance on such tests is often regarded to be the "touchstone" of human intelligence. Unfortunately, ANALOGY simply hasn't turned out to be the first system in a longstanding, comprehensive research program: after all, we find ourselves, at present, trying to start that very program. What went wrong? Well, certainly Psychometric AI would be patently untenable if the tests upon which it is based consist solely of geometric analogies. This point is entailed by such observations as this one from Fischler & Firschein 1987:

If one were offered a machine purported to be intelligent, what would be an appropriate method of evaluating this claim? The most obvious approach might be to give the machine an IQ test. . . . However, [good performance on tasks seen in IQ tests would not] be completely satisfactory because the machine would have to be specially prepared for any specific task that it was asked to perform. The task could not be described to the machine in a normal conversation (verbal or written) if the specific nature of the task was not already programmed into the machine. Such considerations led many people to believe that the ability to communicate freely using some form of natural language is an essential attribute of an intelligent entity. ([Fischler & Firschein, 1987], p. 12)

Unfortunately, while this quote helps explain why ANALOGY in and of itself didn't ignite a research program to drive AI, Fischler & Firschein apparently are familiar with only what we call **narrow**, as opposed to **broad**, intelligence tests. Arguably, this distinction goes back to Descartes' (descartes.haldane.ross.vol1, p. 116) claim that while a machine could in the future pass any test for a particular mental power, no machine could pass a test for any mental power whatsoever. This rather speculative claim can be seen to be cashed out in two different and longstanding views of intelligence within psychology: Thurstone's 1938 and Spearman's 1927. In Thurstone's view (put barbarically), intelligence consists in the capacity to solve a *broad* range of problems, e.g., verbal analogies, geometric analogies, digit recall, story understanding, commonsense reasoning, arithmetical calculation, and so on. In Spearman's view (again, put roughly), intelligence is a specific, *narrow*, underlying capacity (notoriously) referred to as *g*, summoned up to the highest degree when solving highly focused and abstract problems like those ANALOGY solved. The most famous set of "*g*-relevant" problems is the tightly guarded and much-used Raven's 1962 Progressive Matrices, or just 'RPM.' An example of a problem from RPM is shown in Figure 2, which is taken from [Carpenter *et al.*, 1990]. As part of PERI, we have built theorem prover-based agents able to infallibly crack not only geometric analogies, but RPM items they have never seen before (Figure 2 shows part of an OTTER [Wos *et al.*, 1992] proof that serves to identify the solution). (The algorithms deployed by these agents were devised as part of contracted research for The Educational Testing Service, or ETS.) It is much harder

to build agents able to solve broad tests of intelligence ETS tests that include sub-tasks demanding the kinds of communicative capacities Fischler & Firschein have in mind.




Figure 2: Simple RPM Problem "Cracked" by RAIR Lab's PERI

Psychological Corporation's popular WAIS (Wechsler Adult Intelligent Scale) is a paradigmatic example of a broad intelligence test that includes the full array of "Thurstonean" sub-tests (the complete array is enumerated in Baron 2000). While we don't have enough space to explain all these sub-tests, we point out that Fischler & Firschein's criticism of simplistic versions of Psychometric AI certainly evaporates in the face of the WAIS. That this is so follows from the sub-test on the WAIS known as "Comprehension," in which, in ordinary conversation, subjects are asked fiendishly tricky "general knowledge" questions. For example, examinees might be asked to explain why the tires on automobiles are made of rubber, rather than, say, plastic. Perhaps you'll agree that such a question would make a nice challenge to any system purported to have commonsense intelligence. Were CYCORP (<http://www.cyc.com>) to herald a system having general, common-sense intelligence, the WAIS would be a vehicle for verification. There are other sub-tests on the WAIS that are just as challenging.

With help from additional researchers in the RAIR Lab, we are in the process of "cracking" the WAIS, by way of the design and construction of PERI. The sub-test we have cracked most recently is "Block Design." PERI, when given any configuration of blocks in the space of all possible ones, deduces the solution in under a second of CPU time, and proceeds to assemble this solution with its manipulator/gripper. In section 4 we provide some details about PERI's exploits. (Readers wanting to peek ahead should see Figures 4, 5, 6, and 7.)

At present we are tackling the much more difficult sub-test "Picture Arrangement," which requires of examinees that they arrange jumbled snapshots to form coherent stories. (For readers wanting to look ahead, a home-grown example is shown in Figure 3. For legal reasons, actual WAIS examples cannot be shown.) Here our research attempts to make use of prior work in story generation (e.g., Bringsjord & Ferrucci 2000).

We anticipate that some will insist that intelligence tests, even broad ones, are still just too narrow, when put in the context of the full array of cognitive capacities seen in *homo sapiens*. Well, we agree! But we are understanding intelligence, from the standpoint of psychometrics, to include many varied




Figure 3: Examinees must arrange to make a coherent story. tests of intellectual ability. Accordingly, we now move to a less naive definition of PAI:

Psychometric AI is the field devoted to building information-processing entities capable of at least solid performance on all established, validated tests of intelligence and mental ability, a class of tests that includes not just the rather restrictive IQ tests, but also tests of artistic and literary creativity, mechanical ability, and so on.

This definition replaces the old, provisional one: the new definition, when referring to tests of mental ability, is referring to much more than IQ tests. For example, following Sternberg, someone with much musical aptitude would count as brilliant even if their scores on tests of “academic” aptitude (e.g., on the SAT, GRE, LSAT, etc.) were low. But specifically what sorts of additional tests would be involved? We don’t have space to canvass the myriad tests that psychometricians have validated. To give a quick sense of how latitudinarian (and therefore challenging) Psychometric AI is intended to be, we mention The Torrance Tests of Creative Thinking [Torrance, 1990; 1988]. This test comes in both “visual” and “verbal” forms. In the visual form, test takers are asked to draw pictures (often by enriching existing sketches); in the verbal form, test takers are asked to write — creatively. For example, one of the activities subjects engage in on the verbal test is the following.

Most people throw their tin cans away, but they have thousands of interesting and unusual uses. In the spaces below and on the next page, list as many of these interesting and unusual uses as you can think of. Do not limit yourself to any one size of can. You may use as many cans as you like. Do not limit yourself to the uses you have seen or heard about; think about as many possible new uses as you can. (From the verbal version of [Torrance, 1990].)

Believe it or not, after the Torrance Test is administered, one can send it out to be professionally judged by fixed and reliable criteria. Building an intelligent agent capable of scoring well on the Torrance Test is a tall order, but some researchers hitherto unaware of Psychometric AI are apparently in the process of working specifically toward this aim, and variants thereof (e.g. see Bringsjord & Ferrucci 2000 and Bringsjord 1998).

3 Objections

Objection 1: “Can you be serious? PAI is so idiosyncratic!” Recall that we mentioned at the outset that AI textbooks tend to be self-reflective. Let’s look a bit at some of these volumes. Doing so will reveal that Psychometric AI is far from idiosyncratic, because it is (at least arguably) a generalization of a longstanding answer to the “What is AI?” question, namely, the answer that appeals to the Turing Test (TT) and its relatives. To see this answer in action, let’s turn to *AIMA* [Russell & Norvig, 1994], which tells us that there are four general, different ways to define intelligence (pp. 4–8): we can say that an entity is intelligent if and only if it “thinks like humans”, “acts like humans”, “thinks rationally”, or “acts rationally.”

Russell & Norvig 1994 opt for the fourth route, but we want to draw your attention to the first and third ones, which don’t seem exactly promising, because ‘thinking’ is probably no clearer than ‘intelligence.’ However, Turing came up with the test that now bears his name precisely because he found the concept of thinking hopelessly vague. As Russell and Norvig point out, TT and other more stringent tests, e.g., Stevan Harnad’s 1991 *Total Turing Test* (in which a passing robot must display not only human-level linguistic performance, but sensorimotor performance at this level as well), provide a way to clarify the first route in the quartet. Specifically, it can be said that AI is the field devoted to building artificial entities (or systems) able to pass TTT.

We could go on to present case after case in which TT or variants are used to define AI (e.g., see Ginsberg’s 1993 introductory AI text), but perhaps you will agree that whether or not you affirm the TT-based answer to the “What is AI?” question, you have to admit that Turing, by the lights of a good many, is definitely on to something. But what, exactly? Well, no doubt there is more than one reason for the apparent immortality of the TT. But surely one reason for its longevity is simply that it’s a clean, crisp *test*. Tests are attractive in no small part because they have determinate starts and ends, and yield concrete verdicts that can silence unproductive debate. (Of course, computability theory relies heavily on tests. E.g., when we say that a set A is **decidable**, we are among other things saying that we can successfully apply a test to some object o in order to determine whether or not $o \in A$.) Turing’s 1950 goal in his seminal “Computing Machinery and Intelligence” was indeed to supplant the maddeningly vague “Can a computing machine think?” with “Can a computing machine pass the Turing test?” We’re not concerned here with whether he reached his goal. Rather, the idea is that PAI extends and clarifies Turing’s approach.

Objection 2: “But Don’t TT and TTT Subsume Psychometric AI?” We offer a three-part rebuttal: (1) In an attempt to build a robot able to pass for a human, certainly “divide and conquer” is a prudent strategy, and Psychometric AI automatically enforces that methodology: tests are crafted to check for different capacities in focused fashion. Since all topics are fair game in TT and TTT, they have much less value as engineering goals. (2) There is another reason why PAI can be viewed as a helpful generalization of the Turing/Harnad approach. This reason is that, let’s face it, neither TT nor TTT is currently a meaningful objective: they are both gi-

gantly ambitious goals, so much so that no one has really made any progress toward reaching them. (At *Turing 2000*, the conference held at Dartmouth to commemorate both Turing's 1950 prediction that a TT-passing machine would be created before the new millennium and AI's inaugural 1956 conference at Dartmouth, no system better at conversation than a toddler appeared.) Psychometrics offers us an abundance of tests, many of which are *reasonable* challenges for an artificial agent. Psychometric AI appropriates these tests. (3) The tests in question haven't been picked out of thin air. These tests allow us to look into the heart of mind/brain. That's the beauty and power of tests, when they have been empirically and statistically validated. Tests have a gem-like character, and PAI piggybacks on this. Given this, if we build an agent able to pass established tests, we can be fairly confident that as a welcome side-effect we will have an agent capable of many significant higher-level feats.

Objection 3: "But AI has applications that need to be built!" Of course. And Turing wasn't saying all that anybody would work on was passing the TT. After all, few work directly on building an agent able to pass it. An agent able to pass the tests in question will have the capacity to provide the desired applications. Momentarily we present some applications we are working toward that stem directly from PAI.

Objection 4: "But PAI will only tap logicist AI!" Actually, the tasks in question will *unite* logicist and sub-symbolic approaches. For example, "Block Design" on the WAIS requires robotic manipulation, and therefore cognitive processing that is rapid and reactive. This processing is not reasoning-based in PERI. The same can be said for PERI's vision and speech systems.

Objection 5: "But plenty of AI researchers don't do PAI!" This objection can be fleshed out as follows: "I fail to see how you can be seriously proposing a foundation for AI, given that plenty of AI researchers don't deal with tests of any kind. Would we be wrong in regarding such people to be doing AI?"

This is an important objection. It raises the general question of whether we are promoting Psychometric AI as **descriptive** or **prescriptive**. Do we mean to suggest that AI is *in fact* PAI? Or do we mean to maintain that AI *ought* to be, or at least be viewed as, PAI?

The answer is this: Our overall claim about Psychometric AI is that it's *both* descriptive *and* prescriptive. If an AI application is sufficiently ambitious, we hold that Psychometric AI will automatically kick in: the theory is descriptive, in that the developers will inevitably modularize the challenge, and set about building sub-agents that successfully negotiate the tests associated with these modules. On the other hand, if the AI application is a "humble" one, it probably itself constitutes a miniature test: our claim is once again that PAI is descriptive. But how is it then that the theory is also prescriptive? Well, again, we submit that R&D dedicated to building test-taking agents will produce building blocks for accomplishing some very helpful systems. Some examples of such systems are presented in section 5. But first, as planned, we say a few things about PERI.

4 The Robot PERI

PERI, whose name stands for "Psychometric Experimental Robotic Intelligence," is a system capable of logic/reasoning, vision, physical manipulation, speech, and hearing. It is important to note that PERI was not designed to simulate how a human thinks. Our work in connection with this robot is AI, not cognitive modeling.

PERI interacts with the environment via its five-degree-of-freedom vertically articulated mechanical arm, a SCORBOT-ER IX model from the Intelitek Corporation, and a pneumatic two-finger parallel gripper which can either be completely open or closed around an object. Its vision is based on the output of a Sony Black-and-White XC55 Video Camera and Cognex MVS-8100M frame grabber. PERI's speech is transmitted through computer speakers and it hears through a microphone attached to the speaker's head while using the Dragon Naturally-Speaking Professional Solutions Version 6.0 software. At the core of PERI resides its brain and nervous system — a complex Lisp program and an associated Scrobot Advanced Control Language Library. Due to the lack of space, this is as much technical detail as we can prudently reveal about PERI.

The rest of this section will focus on the Block Design task and PERI's success with it. PERI can not only solve the particular Block Design problems in the WAIS, but any Block Design puzzle given to it. For legal reasons we are unable to disclose the Block Design task from the WAIS, therefore we discuss another similar yet even more challenging block puzzle (courtesy of the Binary Arts Corporation). (In the interests of space, and in keeping with the "Philosophical Foundations" category under which this paper falls, we leave aside the mathematization of the WAIS block puzzle and the harder one from Binary Arts. For ease of exposition, we refer to the space in question as *S*.)



Figure 4: One Puzzle Block Folded Out

In this particular puzzle there are a total of four blocks, each of which is different. There are only three colors (pink, purple, and yellow) used to make the design on each side; that design is either a combination of up to four triangles or one solid color. This is done merely to give a specific color to each edge of a block. In fact, all the sides of all four cubes are different from one another. This means there are a total of 24 unique sides. Refer to Figures 4 and 5 for a closer look.

The task is, after having been presented with the cubes for the first time, to place them together so that every edge that touches another is the same color. All cubes must be used, but obviously there are quite a few different solutions. One solution is shown in Figure 6. Does the task sound easy to you? If so, you are supremely confident. While PERI solves the hard-




Figure 5: Blocks (from Binary Arts' # 5766) Scattered

est configuration in a matter of seconds, after having visually examined the blocks, in our experience it can take a clever human several long minutes, if not a half hour, to conquer the entire task. Figure 7 shows PERI assembling a solution.




Figure 6: A Solution to Binary Arts Corp.'s Puzzle 5766

The same basic algorithm (described below) that PERI uses for the WAIS Block Design task can be used to solve any such puzzle in the overall mathematical space \mathcal{S} (\approx 3D regular solid with each side having a characteristic capturable in extensional first-order fashion). The first step is to encode the pieces as declarative information about geometric shapes in PERI's "mind." Before the test is administered to a human participant, he is given a chance to examine the blocks (or other shapes), as would PERI. What follows is a general algorithm which PERI can apply to any 3D physical shapes within a limit of size (i.e., which its gripper can properly hold and manipulate).

General Algorithm for "Cracking" any 3D Block Design in \mathcal{S}

1. Document original pieces by color, dimension, characteristics on each side, and total number.
2. Input goal configuration (a picture that will need to be deciphered).
3. Partition the goal into distinguishable pieces that match similar aspects of those that are available pieces in the original. Start first with the entire goal as one piece. Some aspects of the pieces may be ignored at this stage.
4. Once the goal has been partitioned, determine if original puzzle pieces match the partitioned ones. If not, go back to step 3 and partition it into two pieces, three pieces, etc. (An exceedingly large cutoff is imposed to handle cases where no partitioning is valid, otherwise non-halting is possible.) If there are matching original puzzle pieces to the goal partitioning, go on to step 5.
5. Start with a goal piece and match it to an original piece that has not yet been used. There will be a finite search for each matching piece since step 4 has been passed, indicating the goal is known to be solvable. When a match is found, the original piece is physically added to the solution "arena" by changing the $\langle x, y \rangle$ positioning of the original piece as well as the angle, side, or any other necessary aspect. Continue the present step until no more pieces in the goal exist that need a match.

In the case of the puzzle from the Binary Arts Corporation, the goal configuration is not specified ahead of time (as it is in the WAIS). Therefore, we assume that the goal is given ahead of time and are then able to use the above general algorithm without any modification. PERI can solve the original version of Binary Arts' puzzle; however, the original version doesn't correspond to the WAIS Block Design task, the cracking of which was our goal. The next challenge PERI faces is the more

difficult "Picture Arrangement" subtask of the WAIS, which is discussed in the next section.



Figure 7: PERI Solving a Block Design Puzzle

5 The Future of Psychometric AI

We very briefly offer some thoughts on the future of PAI, from the standpoints of both applications and philosophical foundations.

Much of our prior AI work in automated test generation has been supported by ETS, in connection with wide-scale "high stakes" tests. Our efforts to crack the WAIS should lead to systems able to generate both new items for other established tests, and new tests. In general, it seems to us that there is a growing symbiosis between AI and the test industry. It isn't just the generation side that holds promise: essays written by examinees taking the Educational Testing Service's GMAT (Graduate Management Admission Test) are now read and graded by machine (though there is human input as well), and this trend promises to accelerate. Another application is the use of tests as security devices. For example, our work in Psychometric AI yields a fairly rigorous account (presentation of which is beyond the scope of the present paper) of what tests a machine can solve, versus what tests positively stump machines, but are solvable by humans. This account can be used, for instance, to devise tests designed to weed out troublesome softbots from human users in situations where hackers attempt to build bots in order to overwhelm online systems (e.g., polling systems).

We believe that the "Picture Arrangement" task (see again Figure 3) provides a helpful microworld for another application: the problem of threat anticipation in the intelligence community. Currently, ETS is supporting research in our lab devoted to augmenting the predictive power of intelligence analysts, as that power is described, for example, in [Heuer, 1999]. However, we intend to pursue a new dimension in this work, one based on the notion that predicting what (say) a terrorist will do is in significant measure the completion of a story based on "snapshots" earlier in the "narrative" in which the terrorist is an actor. Our attempt to enable PERI to crack "Picture Arrangement" is the first step in the exploration of this new dimension. In our initial work on this problem, once a snapshot in a group like that shown in Figure 3 is selected, a search for a consistency proof that a particular successor is possible under narrative constraints is fired. If a proof is found, the successor is selected, and the process iterates. Narrative constraints are declarative formalizations of plots, themes, and characters.

We end by briefly discussing two philosophical issues raised by Psychometric AI.

It's perhaps not uninteresting to ponder the philosophical consequences of a future PERI able to excel on *all* established tests. Would we declare PERI in this case to be a bona fide genius? Or would we instead infer that since a "mere" machine is displaying such mastery, human intelligence must go well beyond what can be tested for? We don't have the answers to such questions, but we will be doing our best, through concrete engineering, to raise these questions. We suspect that the degree to which these questions are debated will relate directly to the degree to which our engineering (and those joining us) succeeds.

Finally, we point out that Psychometric AI seems to clarify the dividing line between so-called "Strong" and "Weak" AI. This is so because PAI is fundamentally driven by how things *appear*, not by whether or not some invisible property is in play. "Weak" AI strives for artifacts which, as a matter of directly observable fact, display certain behaviors. "Strong" AI, on the other hand, aims for machines having such mysterious properties as consciousness and qualia. The distinction can be traced back to the "Objection from Consciousness" in Turing's 1950 defense of TT. The objection is famously encapsulated in Professor Jefferson's Lister Oration of 1949:

Not until a machine can write a sonnet or compose a concerto because of thoughts and emotions felt, and not by the chance fall of symbols, could we agree that machine equals brain — that is, not only write it, but know that it had written it. No mechanism could feel (and not merely artificially signal, an easy contrivance) pleasure at its successes, grief when its valves fuse, be warmed by flattery, be made miserable by mistakes, be charmed by sex, be angry or depressed when it cannot get what it wants. (p. 17 of Turing 1950)

Turing responds: "This argument appears to be a denial of the validity of our test." Precisely! He goes on to stick to his guns: he proclaims that if the computer produces the kind of language normally taken to justify ascriptions of consciousness (and he gives examples of such language in connection with a deep understanding of Shakespeare), we ought to go ahead and make those ascriptions: we ought to hold that the computer *is* conscious. Psychometric AI and "Weak" AI are firmly in this camp. If there is some sensible, established test for consciousness, if *appearance* is both demanded and supplied, then members of this camp confidently declare the underlying properties to be in place. In light of this, we readily admit that some proponents of "Strong" AI will not be fans of Psychometric AI. Our readers must judge whether or not that is a virtue or a vice for the form of AI we have defended.

6 Objections — Second Round

Objection 6: "You argue that your approach is much in line with Turing's TT and Harnad's TTT, yet those two are concerned with machines behaving like humans, whereas Psychometric AI is concerned with doing *well* on IQ tests (and the like), which many humans obviously do not necessarily do! One crucial question which surprisingly you don't address at all is: What are the cognitive limitations that stop us from doing much better on such tests? Should similar limitations

be put into PAI systems just to make their performance more human-like and thus give them a chance to behave human-like (and pass TT and TTT), or should PAI systems be built to perform as well as possible (but without a chance of passing TT and TTT)? Should PAI, for example, be concerned with building mediocre chess players with human-like limitations instead of really good ones (assuming that chess could count as some kind of test). And if so, wouldn't PAI have to be cognitive modeling (including modeling of limitations), not "just" AI, contrary to what you say (at least about your work with PERI)?"

Rest assured that we are doing AI first and foremost; cognitive modeling is subsidiary, and can be derived from success on the AI front. This objection is plagued by a *non sequitur*: it doesn't follow from the fact that PAI aims at perfect-performing artificial test-takers that there is there no chance of such systems passing TT and TTT. As an engineering technique, it is wise to "crack" tests via algorithms that can enable us to build artificial agents able to perform flawlessly on them. Once that is accomplished, it will be easy to "de-smart" these agents to produce ones able to match any level of poor performance seen in the human sphere. As Turing (1950) pointed out, if a robotic interlocutor in the TT is sufficiently smart, a query to it like "What is $13,567,890,399 \times 23,456,899,221$?" will return a suitably incompetent (= humanesque) response (e.g., "I would need my trusty calculator for that!"), despite the fact that the machine has the power to produce the correct answer instantly.

Objection 7: "I think you are making a mistake in the end by identifying your approach with 'Weak AI', which according to Searle is just the claim that computers are (merely) useful tools in the study of mind. The position of 'Strong AI,' on the other hand, claims that computers (or computer programs) actually can have human-like mental states (if running the right program). You argue that ascriptions of consciousness (etc.) to computers (or computer-controlled robots) could and should be made, and that in my opinion puts you in the 'Strong' AI camp."

Unfortunately, our critic misunderstands the distinction between "Strong" and "Weak" AI. "Strong" AI is indeed the view that cognition (including subjective awareness or phenomenal consciousness) is computation, and that an appropriately programmed information processing machine operating at or below the Turing Limit can literally *be* subjectively aware. "Weak" AI holds that though the external behavior of such machines might convince external observers to ascribe subjective awareness to them, such machines can't possibly have the relevant mental states. Some opponents of "Strong" AI, such as Bringsjord himself, are actively seeking to engineer artifacts (such as PERI himself!) that would convince everyone (or at least *nearly* everyone) that such artificial creatures are literally subjectively aware! — so the objection obviously rests on a confusion (see Bringsjord 1992, Bringsjord & Ferrucci 2000, and the just-published Bringsjord & Zenzen 2003).

Some additional objections are inevitable. We have space here to but rebut some of them in rapid-fire fashion. More thorough rebuttals are of course forthcoming, in a more capacious venue.

- “The statement that PAI requires a system to perform well in ‘any possibly existing’ psychometric tests makes PAI ill-defined.”

- Not at all. The great challenge we’re willing to face up to is to engineer PERI with the capacity to crack tests he has never seen before. We ask only that the test be validated via the ordinary statistico-mathematical standards used in psychometrics.

- “Your presented example for PERI to tackle — block design — is too simple even in the general case to justify your deep claims.”

- With all due respect: Get serious. Herbert Simon, at the 1956 inception of AI, proclaimed that thinking computers were just around the corner, on the strength of LOGIC THEORIST’s ability to prove such marvelously subtle theorems as that $\neg q \rightarrow \neg p$ can be derived from $p \rightarrow q$. (For the verbatim sanguine proclamation, see Russell & Norvig 1994.) Today, AI has yet to produce a machine with the conversational power of a sharp toddler. PERI is a research program moving now onto tests that no current AI technology can crack, and we will avail ourselves not only of logic-based formalisms, tools, and techniques that are at the heart of the RAIR Lab, but also — to give PERI robust perception/action capability — ammunition from cognitive robotics that may well be outside the logicist paradigm.

- Turing (1950) certainly made some rather deep claims (e.g., that all of human cognition is at bottom computational, and that by the year 2000 not only the conversational power of a toddler would be matched by machines, but that of a novelist would be), and as far as we can tell, the only implementations he could point to for justification were painfully primitive.

Acknowledgments

We are indebted to the Educational Testing Service for sponsoring applied AI in the intersection of AI and testing (the eWRITER Project), that led us to ponder the philosophical foundations of AI in connection with this applied work.

References

- [Baron & Kalsher, 2000] R. Baron and M. Kalsher. *Psychology*. Allyn and Bacon, Boston, Massachusetts, 2000.
- [Bringsjord, 1992] S. Bringsjord *What Robots Can and Can't Be*. Kluwer, Dordrecht, The Netherlands, 1992.
- [Bringsjord, 1998] S. Bringsjord. Chess is Too Easy. *Technology Review*, 1012:23–28, 1998.
- [Bringsjord & Ferrucci, 2000] S. Bringsjord and D. Ferrucci. *Artificial Intelligence and Literary Creativity: Inside the Mind of Brutus, a Storytelling Machine*. Lawrence Erlbaum, Mahwah, New Jersey, 2000.
- [Bringsjord & Zenzen, 2003] S. Bringsjord and M. Zenzen. *Superminds: People Harness Hypercomputation, and More*. Kluwer, Dordrecht, The Netherlands, 2003.

- [Carpenter *et al.*, 1990] P. Carpenter, M. Just, and P. Shell. What One Intelligence Test Measures: A Theoretical Account of the Processing in the Raven Progressive Matrices Test. *Psychological Review*, 97:404–431, 1990.
- [Descartes, 1911] R. Descartes. *The Philosophical Works of Descartes, Volume 1*. Translated by Elizabeth S. Haldane and G.R.T. Ross. Cambridge University Press, Cambridge, United Kingdom, 1911.
- [Evans, 1968] G. Evans. A Program for the Solution of a Class of Geometric-Analogy Intelligence-Test Questions. In Minsky, M., ed. *Semantic Information Processing*, pages 271–353, Cambridge, Massachusetts, 1968. MIT Press
- [Fischler & Firschein, 1987] M. Fischler and O. Firschein. *Intelligence: The Eye, the Brain, and the Computer*. Addison-Wesley, Reading, Massachusetts, 1987.
- [Ginsberg, 1993] M. Ginsberg. *Essentials of Artificial Intelligence*. Morgan Kaufmann, New York, New York, 1993.
- [Harnad, 1991] S. Harnad. Other Bodies, Other Minds: A Machine Incarnation of an Old Philosophical Problem. *Minds and Machines*, 1(1):43–54, 1991.
- [Heuer, 1999] R. J. Heuer. *Psychology of Intelligence Analysis*. Central Intelligence Agency, Center for the Study of Intelligence, 1999.
- [Jensen, 1998] A. Jensen. *The g Factor: The Science of Mental Ability*. Praeger, New York, New York, 1998.
- [Raven, 1962] J. C. Raven. *Advanced Progressive Matrices Set II*. H. K. Lewis, London, United Kingdom, 1962.
- [Russell & Norvig, 1994] S. Russell and P. Norvig. *Artificial Intelligence: A Modern Approach*. Prentice Hall, Saddle River, New Jersey, 1994.
- [Spearman, 1927] C. Spearman. *The Abilities of Man*. MacMillan, New York, New York, 1927.
- [Sternberg, 1988] R. Sternberg. *The Triarchic Mind: A New Theory of Human Intelligence*. Viking, New York, New York, 1988.
- [Thurstone, 1938] L. Thurstone. *Primary Mental Abilities*. University of Chicago Press, Chicago, Illinois, 1938.
- [Torrance, 1988] E. P. Torrance. The Nature of Creativity as Manifest in its Testing. In Sternberg, R., ed. *The Nature of Creativity*, pages 43–75, Cambridge, United Kingdom, 1988. Cambridge University Press
- [Torrance, 1990] E. P. Torrance. *The Torrance Tests of Creative Thinking*. Scholastic Testing Service, Bensenville, Illinois, 1990.
- [Turing, 1950] A. Turing. Computing Machinery and Intelligence. *Mind*, LIX(236):433–460, 1950.
- [Wos *et al.*, 1992] L. Wos, R. Overbeek, E. Lusk, and J. Boyle. *Automated Reasoning: Introduction and Applications*. McGraw Hill, New York, New York, 1992.