

The Symbol Grounding Problem ... Remains Unsolved*

Selmer Bringsjord

Department of Cognitive Science

Department of Computer Science

Lally School of Management & Technology

Rensselaer AI & Reasoning (RAIR) Laboratory

Rensselaer Polytechnic Institute (RPI)

Troy NY 12180 USA

Selmer@rpi.edu

version of 123013NY

Contents

1 Introduction and Plan	1
2 What is the Symbol Grounding Problem?	1
3 T&F's Proposed Solution	5
3.1 Remembering Cog	7
4 The Refutation	9
4.1 Anticipated Rebuttals	11
Rebuttal 1	11
Rebuttal 2	12
5 Concluding Remarks	13
References	15

List of Figures

1 A Concretization of the SGP	5
2 Floridi's Ontology of Information	9

*I'm indebted to Stevan Harnad for sustained and unforgettable discussions both of SGP and Searle's Chinese Room Argument (CRA), and to Luciano Floridi for teaching me enough about information to realize that SGP centers around the level of *semantic* information at the heart of CRA.

1 Introduction and Plan

Taddeo & Floridi (2007) propose a solution to the Symbol Grounding Problem (SGP).¹ Unfortunately, their proposal, while certainly innovative, interesting, and — given the acute difficulty of SGP — brave, merely shows that a class of robots can in theory connect, in some sense, the symbols it manipulates with the external world it perceives, and can, on the strength of that connection, communicate in sub-human fashion.² Such a connection-communication combination is routinely forged (by e.g. the robots in my laboratory (by the robot PERI in Bringsjord & Schimanski 2003) and in countless others, and indeed by robots now on the market), but is rendered tiny and insignificant in the vast and towering face of SGP, which is the problem of building a robot that genuinely understands *semantic* information at the level of human (natural) language, and that acquires for itself the knowledge we gain from reading such language. This problem, which as we'll see can be expressed in concrete, real-world terms that contemporary roboticists painfully appreciate, remains, *contra* T&F, as I show herein, utterly unsolved.

The sequence about to unfold is as follows. I begin by quickly explaining that SGP is indeed the problem just described (§2), and then (§3) briefly recount T&F's proposal. Next, in section 4, I refute this proposal with a short deductive argument, and then (§4.1) anticipate and refute two rebuttals on behalf of the proposal's defenders. A brief conclusion then wraps up the paper — a paper that, I confess, isn't exactly encouraging for those who, like Pollock (1995), expect to sooner rather than later see robot persons.³

2 What is the Symbol Grounding Problem?

Fortunately, this question can be immediately and uncontroversially answered, at least from an informal point of view, by turning to the *locus classicus*, provided by Harnad (1990), wherein three examples of SGP are

¹Their paper is reprinted in (Floridi 2011) (as Chapter 7; Chapter 6 is a sustained, stage-setting argument that other approaches fail to solve SGP (and is also derived directly from a journal paper; viz., (Taddeo & Floridi 2005))). References herein are to the paper itself.

²In terms of the hierarchy of increasingly intelligent artificial agents introduced in (Bringsjord, Noel & Caporale 2000), T&F point the way to synthetic creatures quite dim on the scale. For commentary on a scale specifically for communication, which allows that birds communicate pretty much at the level of the hide-and-feed robots that T&F are impressed by (see §3), see (Bringsjord & Bringsjord 1993).

³Rather long ago I pointed out that this expectation is folly: (Bringsjord 1992).

crisply and succinctly provided.⁴ The first of these examples, and the only one I need herein, is Searle’s well-known Chinese Room Argument (CRA), which I have refined and defended at length (e.g., see Bringsjord 1992, Bringsjord & Noel 2002). It’s probable that the reader is familiar with CRA; regardless, there’s no need here to present CRA in its full and maximally powerful form; we don’t need anything even *close* to this form. It suffices to recall that CRA includes a thought-experiment in which a robot responds to squiggle-squoggles (Chinese, actually) by following a rulebook that pairs these squiggle-squoggles with other squiggle-squoggles (also Chinese) — the catch being that the robot’s activity is based on an homunculus (= Searle) inside the robot who doesn’t understand any Chinese, but who can deftly follow (à la a computer program) the rulebook and thereby give outside observers the *impression* that the robot understands Chinese.⁵ As Harnad puts it when summarizing the CRA for purposes of presenting SGP:

Searle’s simple demonstration . . . consists of imagining himself doing everything the computer does — receiving the Chinese input symbols, manipulating them purely on the basis of their shape . . . , and finally returning the Chinese output symbols. It is evident that Searle (who knows no Chinese) would not be understanding Chinese under those conditions — hence neither could the computer. The **symbols** and the **symbol** manipulation, being all based on shape rather than **meaning**, are systematically interpretable as having **meaning** — that, after all, is what it is to be a **symbol** system But the interpretation will not be intrinsic to the **symbol** system itself: It will be parasitic on the fact that the **symbols** have **meaning** for us, in exactly the same

⁴We read:

Before defining the symbol grounding problem I will give two examples of it. The first comes from Searle’s (1980) celebrated “Chinese Room Argument,” in which the symbol grounding problem is referred to as the problem of intrinsic meaning (or “intentionality”): Searle challenges the core assumption of symbolic AI that a symbol system able to generate behavior indistinguishable from that of a person must have a mind. More specifically, according to the symbolic theory of mind, if a computer could pass the Turing Test (Turing 1950) in Chinese — i.e., if it could respond to all Chinese symbol strings it receives as input with Chinese symbol strings that are indistinguishable from the replies a real Chinese speaker would make (even if we keep testing for a lifetime) — then the computer would understand the meaning of Chinese symbols in the same sense that I understand the meaning of English symbols. (Harnad 1990, p. 337)

⁵I have here skipped to the version of CRA designed by Searle to counter the so-called “Robot Reply” in (Searle 1980). My expansion and refinement of CRA in (Bringsjord & Noel 2002) is a (confessedly intricate) robot-reply version of the argument.

way that the **meanings** of the symbols in a book are not intrinsic, but derive from the **meanings** in our heads. Hence, if the **meanings** of **symbols** in a **symbol** system are extrinsic, rather than intrinsic like the **meanings** in our heads, then they are not a viable model for the **meanings** in our heads: Cognition cannot be just **symbol** manipulation. (Harnad 1990, p. 337; bolded emphasis mine)

I've availed myself here of bolded text to highlight the plain fact that CRA revolves firmly around symbols and their meaning. If CRA, when suitably rigorized and unpacked, is ultimately sound, and if robots are defined as physicalized symbol systems (or Turing machines, etc.), it follows, forever and unshakably, that SGP is unsolvable; hence T&F would of necessity fail, no matter what they say. While my own view is indeed that CRA is sound, the present paper is composed from the point of view of those, like T&F (and the thinkers they understand to have tried to solve SGP, and failed, e.g. Ron Sun; see Chapter 6 of (Floridi 2011)), see CRA is merely a *seemingly* sound argument. SGP then becomes the problem of showing that there is a way to escape Searle's reasoning.

Of course, it must be admitted that non-philosophers will regard Harnad and Searle and Bringsjord to have delivered dizzying disquisitions on the topic; it's thus fortunate that SGP is really and truly quite straightforward, and disturbingly concrete: You own a young robotics company, Robots R Us, and are rather keen on manufacturing and selling, specifically, housekeeping robots — robots that will clean floors, bathrooms, kitchens, dust; do laundry, declutter, cook, and so on. You know that if your robo-housekeepers perform in stellar fashion, you could become quite rich; but you *also* know that if your robots wreak havoc in the home of even one wealthy, well-connected family, the public-relations fallout could kill your dreams of becoming a modern-day Cornelius Vanderbilt. In this context, we can further and fast fix SGP by considering the following missive:

My Dearest George,

Thank you *so* very much for the V Day gifts you left for me! I know of late that we haven't exactly been in connubial bliss, but they, and the new protocols that can be extracted from them, are nothing short of aphrodisiacs, and when you return I'm sure that our relationship will be radically warmer. I commend you on your purchase of roses just starting to wilt, as I'm sure you saved a good deal of money; and the chocolates, despite the long-ago sell-by date on the box, were stale, but redeemed by being commendably cheap. Frugality has always been your strong suit, but now you've infused it with romantic thoughtfulness. Here's hoping that Rob will follow suit, and that from now on all our arrangements will be dreary from the start, and all our food old but, technically speaking, edible.

♡ Love Allison ♡

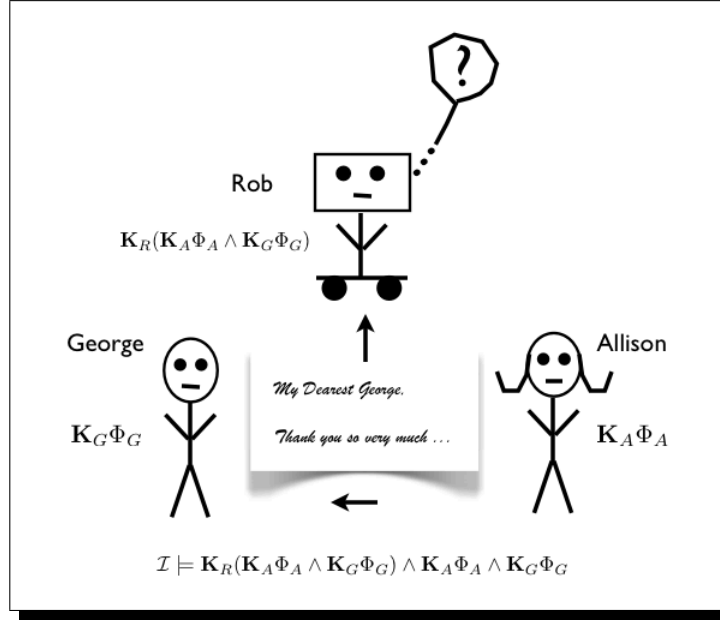
Suppose that Rob is a robot built and sold by Robots R Us, and that he comes upon Allison's letter. It would be very bad for the company, and for you as its owner, if Rob starting energetically setting out dead roses and cooking turned food.⁶ What does it take, overall and abstractly, for it to be the case that George understands what Allison has written, *and* that Rob understands the letter as well, and the background facts that it presupposes? For the reader's convenience, I've summed up the answer to this question in Figure 1.

The figure shows the three characters in my parable, and indicates the letter left for George by Allison. It also depicts the bare epistemic bones of the parable: Allison knows some collection of propositions expressed as a set of formulae Φ_A ; the same structure is true of George. In addition, were it to be the case that Rob genuinely understands, in "CRA-busting" and SGP-solving fashion, Allison's letter, he would know that Allison (George) knows Φ_A (Φ_B). I assume that all formulas are expressed in a formal logical intensional language that subsumes first-order logic.⁷ In addition, if you

⁶It would be easy enough to modify the example so that Allison's "recommended" actions were more violent, and hence the consequences of Rob's misunderstanding in turn more severe.

⁷We know that even problem-solving easier than that demanded by higher education requires understanding semantic content at the level of first-order logic; this is something e.g. Inhelder & Piaget (1958) showed. Of course, in addition, my parable presupposes at least a quantified epistemic logic, but for purposes of evaluating T&F's proposed solution to SGP, there is nothing to be gained by considering the details of such a logic. Interested readers may wish to see e.g. (Bringsjord & Govindarajulu 2012).

Figure 1: A Concretization of the SGP



examine the bottom line in the figure you will see that all the main top-level knowledge is true on some semantic model \mathcal{I} .⁸

T&F’s proposed solution, to which we now turn, must be measured against what is true in the Allison-George scenario, as summed up by Figure 1.

3 T&F’s Proposed Solution

The solution proposed by T&F is based on a new theory of meaning, **Action-based Semantics**, or just AbS, and on a new class of artificial agents, **two-machine artificial agents**, abbreviated AM². As T&F write in their abstract, their core innovation is that “AM²s implement AbS, and this allows them to ground their symbols semantically and to develop some fairly advanced semantic abilities ...”⁹ For T&F, a solution to SGP must satisfy

⁸There is by the way nothing to be gained by searching out the details of \mathcal{I} . All that is needed for my refutation is that \mathcal{I} subsumes a garden-variety truth-functional framework, such as the standard model-theoretic semantics of first-order logic.

⁹Alert, critical readers will find the “some fairly advanced” phrase ominously elastic, but I ignore this soon-to-be-the-crux issue for the moment.

the **Z Condition**, which means that the solution can't rely on any form of **innatism**, nor on any form of **externalism**. As T&F claim, a purported solution that allows “semantic resources” to be “pre-installed” in the AA is not a solution at all; and likewise, a proposed solution that allows semantic resources to be imported from the outside is likewise no solution.¹⁰ Needless to say, they hold that their solution to SGP satisfies the Z Condition.

The present their solution in a determinate sequence of moves; overall, there are two steps to their story. In the first step, T&F outline how new meanings are coupled with symbols; this is AbS. In the second step, they describe a machine in AM² that implements AbS without presupposing semantic resources or capacities; the second step is played out in a sequence of sub-steps: They first “refer to a specific learning rule and to an evolutionary scenario in order to show how a population of AM²s could develop its semantic abilities autonomously” (Taddeo & Floridi 2007, p. n). Then, after replying to three objections, the second step culminates with a final sub-step and what should be a crowning achievement: namely, a description of “how a population of AM²s can develop more complex semantic abilities, such as semantically grounded communication and a shared semantics” (Taddeo & Floridi 2007, p. n). Since, as we've seen, SGP is the problem of showing that a computational system can be capable of genuinely understanding the semantics of human-level natural language, the crowning achievement must provide what is needed to show that their proposal allows an artificial (computational) agent to grasp Allision's missive (or to grasp the Chinese symbols it receives as input from outside-the-Room native Chinese speakers).

Well, what is the crowing achievement? We are told to imagine a population of AM²s that interact with the environment in two ways: **feed**, or **hide**. The latter action is taken to avoid the attacks of three kinds of predators, α , β , and γ . We then read:

¹⁰There are clearly some immediate consequences of the affirmation of the Z Condition. First, subscribers to the doctrine of innate semantic capacities are rather left out in the cold. (For what it's worth, this includes me, since I e.g. believe that “background logic” — treated elegantly and rigorously in (Ebbinghaus, Flum & Thomas 1994) — and facility with it is innate. E.g., I hold that facility for using disjunctive syllogism is innate.) Second, subscribers to the doctrine that high-level semantic capacity at the level of neurobiologically normal human persons was indeed “injected” into *homo sapiens sapiens* by God are in the cold as well. (Hence Darwin is warm; but the co-discoverer of evolution by mutation and natural selection, Wallace, a theist, is by the Z Condition freezing.) Third, thinkers who hold that “Fodor's Paradox” (for a recent presentation, see e.g. Fodor 2008) (sometimes called “the paradox of learning”) isn't solvable by any move that fails to invoke either innatism or externalism are also automatically cast aside. (By the way, if T&F succeed in solving SGP under the constraint of the Z Condition, they have automatically also solved Fodor's Paradox!)

Suppose the AM²s involved in this scenario engage in a kind of adaptive language game, such as the *guess game*. A guess game is a technique used to study the development of common language in situated AAs. . . . Each AA involved in the game has a role: one is the *speaker*, and names the object that it perceives; the other is the *hearer* and has to find out the objects named by the speaker by trial and error. The speaker communicates only to convey the name of a perceived referent, and the hearer communicates only to inform the speaker about its guessing concerning the referent named by the speaker. During the game, the AAs interact and develop a common system of symbols. The game ends successfully if the two AAs develop a shared lexicon, grounded in the interaction among themselves and with the environment. In the case of the AM²s described above, the communicated symbols are related to the speaker's internal states and, indirectly, to the action that it performs, for example 'open your mouth' or 'hide yourself'. (Taddeo & Floridi 2007, p. n)

I assume for the sake of argument that roboticists following the guidance of T&F's proposal can thoroughly and completely succeed in the engineering of a population of AM²s wherein communication regarding predators is flawless. Actually, since I have no doubt that the engineering of such a population could be carried out in my own laboratory now, I assert that I *know* that roboticists could engineer, now, via an evolutionary approach consistent with the Z Condition, such a population. Please note, therefore, that I assume that the story T&F tell is as bright as it can possibly be.

3.1 Remembering Cog

I conclude this section by pointing out that if the history of AI and robotics is any guide, T&F's appeal to evolution isn't promising. In the annals of super-sanguine AI,¹¹ to which T&F have by any metric contributed, such appeals are at least tinged with ignominy, given the sad story of Cog, one made all the sadder by the arguably cooky optimism of Dennett, a member of the team. Here is Dennett in defense of the T&F-like idea that evolution will do the trick:

How plausible is the hope that Cog can retrace the steps of millions of years of evolution in a few months or years of laboratory exploration?

¹¹Contributions to which have been made by Turing (1950), who proclaimed that by the year 2000 the test that bears his name would be passed by computing machines, and who, in a brave but fideistic call for evolutionary approaches that would presumably resonate with T&F, saw little problem in setting down a machine in Cog-style and having it learn everything that otherwise engineers would painstakingly labor to figure out, and program in.

Notice first that what I have just described is a variety of Lamarckian inheritance that no organic lineage has been able to avail itself of. The acquired design innovations of Cog-I can be immediately transferred to Cog-II, a speed-up of evolution of tremendous, if incalculable, magnitude. Moreover, if you bear in mind that, unlike the natural case, there will be a team of overseers ready to make patches whenever obvious shortcomings reveal themselves, and to jog the systems out of ruts whenever they enter them, it is not so outrageous a hope, in our opinion. But then, we are all rather outrageous people. (Dennett 1994, p. 140)

And what's the trick that "guided" evolution was supposed to pull off? Essentially the trick of behaving at the conversational level of the box in CRA. Here's Dennett blithely sweeping aside the Chomskyan innatist position that SGP is unsolvable:

One talent that we have hopes of teaching to Cog is a rudimentary capacity for human language. And here we run into the fabled innate language organ or Language Acquisition Device made famous by Noam Chomsky. Is there going to be an attempt to build an innate LAD for our Cog? No. We are going to try to get Cog to build language the hard way, the way our ancestors must have done, over thousands of generations. Cog has ears (four, because it's easier to get good localization with four microphones than with carefully shaped ears like ours!) and some special-purpose signal-analyzing software is being developed to give Cog a fairly good chance of discriminating human speech sounds, and probably the capacity to distinguish different human voices. Cog will also have to have speech synthesis hardware and software, of course, but decisions have not yet been reached about the details. It is important to have Cog as well-equipped as possible for rich and natural interactions with human beings, for the team intends to take advantage of as much free labor as it can. Untrained people ought to be able to spend time—hours if they like, and we rather hope they do—trying to get Cog to learn this or that. Growing into an adult is a long, time-consuming business, and Cog—and the team that is building Cog—will need all the help it can get. (Dennett 1994, pp. 140–141)

Where is Cog now? Not in a place that augurs well for T&F: As Wikipedia reports, in cold understatement: "As of 2003, all development of the project had ceased. Today Cog is retired to the MIT museum."¹² Did any of the high hopes for Cog, in the direction of natural-language prowess, pan out? Alas, no, not a one. Cog isn't a single nanometer closer

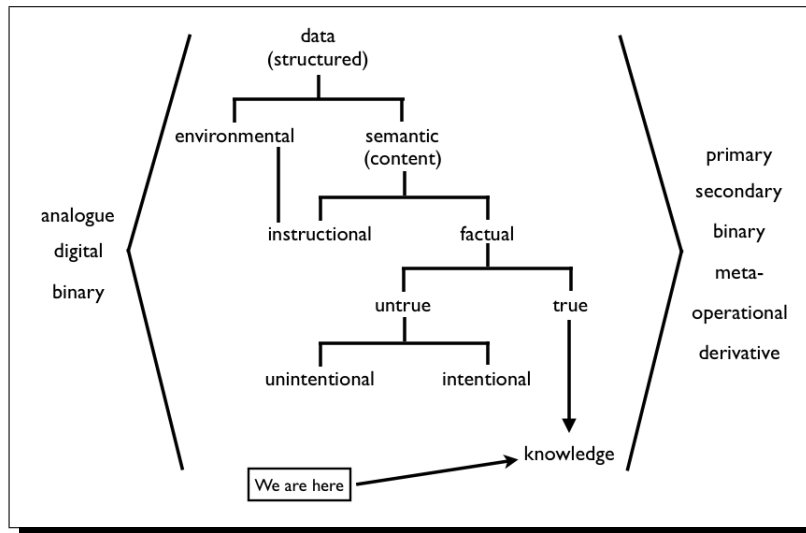
¹²[http://en.wikipedia.org/wiki/Cog_\(project\)](http://en.wikipedia.org/wiki/Cog_(project))

to being able to pass the Turing Test than the raw inanimate metal that went into it; *a fortiori*, Cog is no closer to vindicating an evolutionary approach to grounding symbols at the level required to address SGP than that same metal.

4 The Refutation

Certain ingredients are needed for my refutation, chief among which is Floridi's (2010) own account of information, summed up in Figure 2, which I reproduce directly from that elegant, succinct book. The reader will notice that in this figure I indicate that central to the present discussion is a specific form of factual and true semantic information: viz., *knowledge*. This is because, as indicated in Figure 1 and the prose exposition of it given earlier, we are specifically concerned with what Allison, George, and Rob know. The same thing is true of the Chinese Room, where we are concerned with whether Searle-in-the-Room knows what the Chinese characters mean. We can now proceed directly to the the refutation, in the form of a deductive argument that immediately follows.

Figure 2: Floridi's Ontology of Information



P1 SGP is solved by T&F iff we know that our hypothetical robot-cists can, given enough time and energy, construct, on the basis

of T&F’s recipe, a robot Rob whose performance in homes will be non-stupid.

P2s We know that our hypothetical (Robots R Us) roboticists can, given enough time and energy, construct, on the basis of T&F’s recipe, a robot Rob whose performance in homes will be non-stupid only if we know that their recipe is one the following of which results in a robot Rob capable of knowing (and a-ing generally) such things as George knows Φ_G and Allison knows Φ_A , in part as a result of having read T. (Here T is the brief missive regarding Valentine’s Day.)

P3 We don’t know that T&F’s recipe is one the following of which results in a robot Rob capable of knowing (and a-ing generally) such things as George knows Φ_G and Allison knows Φ_A , in part as a result of having read T.

Therefore (*modus tollens* \times 2):

C SGP isn’t solved by T&F.

This argument is formally valid, obviously. It has but three premises; if all true, T&F have failed. It’s exceedingly hard to see what could possibly be objectionable about P1 and P2. As to P3, it’s established by an additional deductive argument that couldn’t be simpler, to wit:

P4 If T&F don’t (even) discuss how their recipe engenders in a robot knowledge of semantic information at the level of human natural language, then P3.

P5 T&F don’t (even) discuss how their recipe engenders in a robot knowledge of semantic information at the level of human natural language.

Therefore (good ’ol *modus ponens*):

P3.

The textual evidence for **P5** is clear and indisputable, as T&F themselves admit that their account doesn’t extend to the case of knowledge in connection with human-level natural language (which is of course the level at which Allison’s missive is pitched). Put another way, the hide-and-feed game discussed above, which I claim to know to be easy enough to engineer with help from T&F’s paper, in no way provides any reason to rationally believe that a robot can have semantic knowledge of semantic information at the level of human natural language — and in fact this is something T&F *themselves* concede. For example, we read:

Still, the semantics elaborated by the AM²s has a certain lack of completeness and complexity. AbS provides a minimal and simple semantics. It is patently not truth-functional nor does it justify the elaboration of meanings through some explicit agreement among the agents. (Taddeo & Floridi 2007, p. n)

Since a solution to the concretized version of SGP must yield a semantics at the level of \mathcal{I} (recall the relevant equation in Figure 1), which subsumes simple truth-functional semantics, **P5** is verified by this part of T&F’s paper alone.

4.1 Anticipated Rebuttals

There is no way to know for sure what rebuttals will come from defenders of T&F’s proposed solution, but I anticipate a pair, and now give and promptly dispose of both.

Rebuttal 1

“Though you may not find our appeal to evolution convincing, and though that appeal may in fact not provide the engineering details you seek, certainly you cannot deny that, to put it starkly, evolution got the job done in the first place. After all, how did the semantic communicative capacity of human persons arrive on the scene, if not via evolution? At the end of the day, we are simply fleshing out and applying the existence proof that stands before our eyes, and yours; the existence proof that is us.”

Three problems derail this rebuttal:

1. As far as I’m concerned, it’s *not* clear that evolution is what “got the job done” when it comes to the mental powers of human persons. With Wallace, the co-discoverer with Darwin of the process of evolution by mutation and natural selection, I accept the sustained argument that is *Origin* (= Darwin 1859), but flatly reject the arguments of *Descent* (= Darwin 1997); that is, I reject the proposition (ME) that our mental powers are the product of evolution — on the strength of what is known as **Wallace’s Paradox**, which in a word flows from the fact that while our hunter-gatherer forbears had our brains, and hence our capacity for — say — working in the tensor calculus, this capacity was entirely superfluous with respect to picking berries and hunting antelope. I discuss Wallace’s Paradox at length in (Bringsjord 2001), and will not rehearse here the arguments I provide there in defense of Wallace’s argument for the falsity of (ME).

2. Despite the anti-innatism that the Z Condition, in harmony with for instance Dennett’s unforeseen rejection of Chomsky’s LDA, regiments, it is far from clear that innatism is false. In fact, the dominant response to Wallace from modern evolutionists (Gould) comes at least very close to affirming a form of innatism, because it views higher-level cognition at the level of that pointed to by Figure 1 as having suddenly arrived on the scene by way of a serendipitous mutation that remained dormant, and firmly innate, for tens of thousands of years.
3. Rebuttal 1 ignores what has been said above about the sad case of Cog. The moral of the carcass of Cog is that unless one knows how to engineer an AA capable of understanding semantic information, it’s mere wishfull thinking to assert that there is an evolutionary route to such an AA. Wishful thinking isn’t a promising engineering technique. It hardly follows from the fact that artificial agents able to thrive in the hide-and-feed world can arrive on scene, that that which enabled them to do so can allow an understanding Rob to arrive. Robots able to yell imperatives like “Hide!” or “Feed!”, which after all for purposes of the game in question could simply be “1” or “0,” don’t qualify as robots able to succeed in the Chinese Room.

Rebuttal 2

“You do concede, Bringsjord, that we solve *a* symbol grounding problem. We have shown that symbols carrying semantic information can be grounded in robots (or more generally, artificial agents (AAs)) in satisfaction of the Z Condition. This is all we intended to demonstrate.”

Yes, I do so concede. But there are two fatal problems with the rebuttal nonetheless. First, you have no scholarly entitlement to say that T&F have solved, or even pointed toward a solution of, SGP (vs. what we can call ‘sgp’). Second, absent a proof, or at least a strong argument, that sgp constitutes a rung on a ladder that will eventually reach a solution to SGP, the progress you have won isn’t necessarily progress at all. After all, to repeat, my own robots solve sgp routinely (and in microworlds larger than the feeding-predator game), but there is no such ladder to be seen in the case of this work. If the ladder exists in some hypothetical or mathematical space unexemplified, the onus is surely on T&F to communicate that unexemplified abstraction to the engineers craving a solution to SGP, and thereby provide at least a part of the antidote to Rob’s potential lack of understanding of Allison’s searing epistle.

5 Concluding Remarks

In light of the foregoing, Robots R Us is left with no assurance. For it remains the case that there is nothing on the scene, despite T&F's proposal, that can provide a guarantee that Rob will not innocently take it upon himself to furnish the house in which he toils with cheap, faded flowers, and stale food. This is so, in a nutshell, because the company has no reason to believe that Rob will understand the semantic information expressed by the sentences in Allison's letter, but *does* have reason, courtesy of CRA, to believe that Rob *won't* understand her letter. This epistemically uncomfortable situation is in stark contrast to what would be your own situation if you were hypothetically in the role of Rob. If you imagine yourself coming upon Allison's epistle while in the role of human housecleaner, you will see that it's *possible* that you too might provide dying flowers and old chocolate to those you serve — but this possibility would obtain for reasons other than a lack of understanding on your part. (You might get angry and decide to spite Allison and George, or you might decide to play a practical joke, and so on.) In other words, in your own case, which is after all the human, not the robot, one, there *is* a guarantee that rotten roses will not decorate the home of Allison and George as a result of an inability to ground the symbols in Allison's letter.¹³ SGP, which in distilled form simply is the vivid contrast between Robots R Us' confidence in you, versus its doubts about Rob, stands before us still, as untamed and undiminished as ever.

¹³And note that at worst you might be slightly unsure that Allison's was being sarcastic — in which case you would still genuinely understand the situation, and wouldn't energetically set out withered flowers.

References

- Bringsjord, S. (1992), *What Robots Can and Can't Be*, Kluwer, Dordrecht, The Netherlands.
- Bringsjord, S. (2001), 'Are We Evolved Computers? A Critical Review of Steven Pinker's *How the Mind Works*', *Philosophical Psychology* **14**(2), 227–243. A more detailed version of this paper is available from the author, and is currently available online at <http://www.rpi.edu/~faheyj2/SB/SELPAP/PINKER/pinker.rev2.pdf>.
- Bringsjord, S. & Bringsjord, E. (1993), 'Animal Communication of Private States Doesn't Illuminate the Human Case', *Behavioral and Brain Sciences* **16**(4), 645–646.
- Bringsjord, S. & Govindarajulu, N. S. (2012), 'Given the Web, What is Intelligence, Really?', *Metaphilosophy* **43**(4), 361–532. This URL is to a preprint of the paper.
URL: <http://kryten.mm.rpi.edu/SB NSG Real Intelligence 040912.pdf>
- Bringsjord, S. & Noel, R. (2002), Real Robots and the Missing Thought Experiment in the Chinese Room Dialectic, in J. Preston & M. Bishop, eds, 'Views into the Chinese Room: New Essays on Searle and Artificial Intelligence', Oxford University Press, Oxford, UK, pp. 144–166.
- Bringsjord, S., Noel, R. & Caporale, C. (2000), 'Animals, Zombanimals, and the Total Turing Test: The Essence of Artificial Intelligence', *Journal of Logic, Language, and Information* **9**, 397–418.
URL: <http://kryten.mm.rpi.edu/zombanimals.pdf>
- Bringsjord, S. & Schimanski, B. (2003), What is Artificial Intelligence? Psychometric AI as an Answer, in 'Proceedings of the 18th International Joint Conference on Artificial Intelligence (IJCAI-03)', Morgan Kaufmann, San Francisco, CA, pp. 887–893.
URL: <http://kryten.mm.rpi.edu/scb.bs.pai.ijcai03.pdf>
- Darwin, C. (1859), *On the Origin of the Species by Means of Natural Selection*, Murray, London, UK.
- Darwin, C. (1997), *The Descent of Man*, Prometheus, Amherst, NY. A reprint edition. The book was first published in 1871.
- Dennett, D. (1994), 'The Practical Requirements for Making a Conscious Robot', *Philosophical Transactions of the Royal Society of London* **349**, 133–146.
- Ebbinghaus, H. D., Flum, J. & Thomas, W. (1994), *Mathematical Logic (second edition)*, Springer-Verlag, New York, NY.
- Floridi, L. (2010), *Information: A Very Short Introduction*, Oxford University Press, Oxford University Press.

- Floridi, L. (2011), *The Philosophy of Information*, Oxford University Press, Oxford, UK.
- Fodor, J. (2008), *LOT2: The Language of Thought Revisited*, Oxford University Press, Oxford, UK.
- Harnad, S. (1990), ‘The Symbol Grounding Problem’, *Physica D* **42**, 335–346.
- Inhelder, B. & Piaget, J. (1958), *The Growth of Logical Thinking from Childhood to Adolescence*, Basic Books, New York, NY.
- Pollock, J. (1995), *Cognitive Carpentry: A Blueprint for How to Build a Person*, MIT Press, Cambridge, MA.
- Searle, J. (1980), ‘Minds, Brains and Programs’, *Behavioral and Brain Sciences* **3**, 417–424.
- Taddeo, M. & Floridi, L. (2005), ‘Solving the Symbol Grounding Problem: a Critical Review of Fifteen Years of Research’, *Journal of Experimental and Theoretical Artificial Intelligence* **17**(4), 419–445. This paper is reprinted in Floridi, L. (2011) *The Philosophy of Information* (Oxford, UK: OUP).
- Taddeo, M. & Floridi, L. (2007), ‘A Praxical Solution of the Symbol Grounding Problem’, *Minds and Machines* **17**(4), 369–389. This paper is reprinted in Floridi, L. (2011) *The Philosophy of Information* (Oxford, UK: OUP).
- Turing, A. (1950), ‘Computing Machinery and Intelligence’, *Mind* **LIX** (59)(236), 433–460.