

How Helen Keller used syntactic semantics to escape from a Chinese Room

William J. Rapaport

Received: 8 August 2006 / Accepted: 16 January 2007
© Springer Science+Business Media B.V. 2007

Abstract A computer can come to understand natural language the same way Helen Keller did: by using “syntactic semantics”—a theory of how syntax can suffice for semantics, i.e., how semantics for natural language can be provided by means of computational symbol manipulation. This essay considers real-life approximations of Chinese Rooms, focusing on Helen Keller’s experiences growing up deaf and blind, locked in a sort of Chinese Room yet learning how to communicate with the outside world. Using the SNePS computational knowledge-representation system, the essay analyzes Keller’s belief that learning that “everything has a name” was the key to her success, enabling her to “partition” her mental concepts into mental representations of: words, objects, and the naming relations between them. It next looks at Herbert Terrace’s theory of naming, which is akin to Keller’s, and which only humans are supposed to be capable of. The essay suggests that computers at least, and perhaps non-human primates, are also capable of this kind of naming.

Keywords Animal communication · Chinese Room Argument · Helen Keller · Herbert Terrace · Knowledge representation · Names · Natural-language understanding · SNePS · Syntactic semantics

Chinese Rooms, syntactic semantics, and SNePS

Philosophy is the history of a deaf-blind person writ large.

—Helen Keller (1903, p. 25), cited in Leiber (1996, p. 426).

The shortest answer to the question, “How can a computer come to understand natural language?” is: “The same way Helen Keller did.”

—Albert Goldfain (personal communication, 2006).

W. J. Rapaport (✉)

Department of Computer Science and Engineering, Department of Philosophy, Department of Linguistics, and Center for Cognitive Science, State University of New York at Buffalo, Buffalo, NY 14260-2000, USA
e-mail: rapaport@cse.buffalo.edu

This essay investigates the nature of Chinese Rooms and how to escape from them. We consider a few real-life approximations of them before focusing on one in particular: Helen Keller's experiences growing up deaf and blind. We will examine what her experiences were like locked in a sort of Chinese Room yet learning how to communicate successfully—indeed, brilliantly—with the outside world.

In Keller's own view, learning that “everything has a name” was the key to her success. Using the SNePS knowledge-representation and reasoning system, we will examine the sense in which this was true and its relationship to “syntactic semantics”—a theory of how the kind of syntactic symbol manipulation that computers do so well suffices for providing the kind of semantics needed for natural-language understanding (NLU) (*pace* Searle, 1980, 1993, p. 68).

I also critique Herbert Terrace's Keller-like theory of naming. Terrace believes that only humans are capable of naming in his sense, hence that non-human primates cannot learn language. I argue that at least computers, and possibly non-human primates, are also capable of naming in Terrace's sense, hence that, if it *is* necessary for learning language, then computers (and, perhaps, non-human primates) can understand language.

A quiz

Who said the following, and when?

Consider a box B inside of which we have a man L with a desk, pencils and paper. On one side B has two slots, marked *input* and *output*. If we write a number on paper and pass it through the input slot, L takes it and begins performing certain computations. If and when he finishes, he writes down a number obtained from the computation and passes it back to us through the output slot. Assume further that L has with him explicit deterministic instructions of finite length as to how the computation is to be done. We refer to these instructions as P. Finally, assume that the supply of paper is inexhaustible, and that B can be enlarged in size so that an arbitrarily large amount of paper work can be stored in it in the course of any single computation. ...I think we had better assume, too, that L himself is inexhaustible, since we do not care how long it takes for an output to appear, provided that it does eventually appear after a finite amount of computation. We refer to the system B–L–P as M. ... In the approach of Turing, the symbolism and specifications are such that the entire B–L–P system can be viewed as a digital computer Roughly, to use modern computing terms, L becomes the logical component of the computer, and P becomes its program. In Turing's approach, the entire system M is hence called a *Turing machine*.

If you answered “John Searle, in 1980”, you've got the wrong person and you're off by some 20 years: it was written by Hartley Rogers (1959, pp. 115, 117)! Searle (1980), of course, later made such a “box” famous as the “Chinese Room” (just replace ‘a number’ in Rogers's quote with ‘a sentence (in Chinese)’ or ‘squiggles that are in fact Chinese writing’), used in a thought experiment to create a counterexample to the Turing test (Rapaport, 2005c, d; Shieber, 2004; Turing, 1950).

Chinese Rooms

Let us define a “Chinese Room” (CR) as a kind of Turing machine as Rogers specified, i.e., a system consisting of two subsystems: a program and a processor.

The program

The (computer) program is an implementation, in some language, of an algorithm. In a CR, the program accepts as input what are in fact (written or spoken) expressions in a natural language (NL) (archetypally Chinese), and manipulates the input to yield as output “appropriate” expressions in that language.¹

The “appropriateness” of the output depends on the kind of input. If the input is a story together with “reading comprehension” questions about it, then “appropriate” output would be whatever would count as “appropriate” answers to the questions. They need not be correct or non-trivial² answers, but at least they should be relevant ones. If the input is a fragment of a conversation in Chinese, then “appropriate” output would be whatever would count as an “appropriate” continuation of the conversation. This could include changing the subject in the way that sometimes happens in ordinary conversations, but should probably not include the kinds of abrupt changes that one sees in Eliza-like programs that try to avoid topics that they are not designed to handle (Weizenbaum, 1966).

These are *de re* descriptions of the I/O, i.e., descriptions from a third-person point of view. Viewed *de dicto*—i.e., from a first-person point of view (from the *program’s* point of view, so to speak)—the program accepts as input uninterpreted marks (“meaningless squiggles”; Searle, 1980, p. 417),³ and manipulates them to produce other such marks as output.

This is an NL-processing CR. There could also be CRs for visual processing or other kinds of cognition. Given the “AI-completeness” of these tasks,⁴ probably any CR needs to handle all such tasks. Thus, a more general description of a CR would not restrict the I/O to NL expressions but—as with Rogers’ box—would allow any expressions from some syntactic system.

¹ Often, programs are characterized merely in input–output (I/O) terms. I use the phrase “manipulate the input to yield the output” in order to emphasize the *algorithmic* nature of the program. (The “manipulation”, of course, could be a null operation, in which case the algorithm (or program) would indeed be a mere I/O specification.) However, to be even more accurate, we should say that the program *describes how to* accept and manipulate the input, since a program is a static (usually textual) object, as opposed to a dynamic “process”. See Sect. “The processor” below.

² Albert Goldfain pointed out to me that some of Eliza’s responses are only trivially appropriate (see Weizenbaum, 1966).

³ On the relation of “symbols” to “marks” (roughly, uninterpreted symbols, for readers who will excuse the apparent oxymoronic nature of that phrase), see Rapaport (1995, 2000).

⁴ A problem is AI-complete if a (computational) solution for it requires or produces (computational) solutions to *all* problems in AI (Shapiro, 1992).

The processor

Since a program is a static object, a second subsystem is required that causes the program to be executed, i.e., that creates a dynamic “process” consisting of actually *taking* the input, *manipulating* it, and *producing* the output.⁵

In a computer, this second subsystem is the central processing unit (CPU), fetching and executing the instructions (by implementing the switch-settings specified by the program). In the Chinese Room Argument (CRA), the second subsystem is a human, named ‘Searle’, who (consciously) follows the instructions of the program. In Rogers’s box, it is L (“the logical component of the computer”).

In the CRA, Searle-in-the-room (i.e., the processor) by hypothesis does not understand Chinese. Thus, the I/O “squiggles” are meaningless *to Searle-in-the-room*; he does not interpret them. All that he has access to is the *syntax* of the input. On Searle-the-philosopher’s interpretation of the CRA, even though Searle-in-the-room does not understand Chinese, he is having a fluent conversation in Chinese (i.e., is passing a Turing test for Chinese NLU). Therefore (by analogy? by universal generalization?), no computer—more precisely, no CR processor—equipped with a Chinese NLU program could understand Chinese. Thus, NLU, or cognition more generally, is not computable.

The system

Searle-the-philosopher assumes that it is Searle-in-the-room (the processor) who does not understand Chinese. But Searle-the-philosopher concludes that no *computer* can understand Chinese solely computationally. Yet the computer is a third thing over and above the two subsystems; it is the system itself, consisting of the processor *interacting with* (i.e., executing) the program. If the combined system is not a single (perhaps “emergent”) thing but is completely analyzable (or reducible) to its two subsystems, then so much the worse for Searle’s analogy. For humans are surely single entities (at least from the perspective of cognitive psychology), and it is a human who understands NL, not a *part* of a human (i.e., a brain). Since no one thinks it necessary for a CPU to “know” or “understand” what it is doing, there is no reason for Searle-in-the-room’s lack of knowledge or understanding of Chinese to prevent the *system* from understanding Chinese. It is the entire computer that processes data, not any single part of it (e.g., only its CPU or its memory; cf. Rogers, above; Sect. ‘Syntactic semantics’, below).

But the system is *not* a mere *set* of its two subsystems. The program by itself can do nothing; the processor without a program to execute has nothing to do. On this view, too, it is not Searle-in-the-room *simpliciter* (the processor) who is conversing in Chinese. Rather, the entire CR *system*—the processor executing the program—is conversing in Chinese. This “systems reply” to the CRA seems clearly related to notions of “wide” or “extended” cognition (Clark & Chalmers, 1998; Hutchins, 1995a, b; cf. Giere, 2002). But to say that the *system* understands Chinese might be to describe the situation from the system’s first-person point of view. Alternatively, one

⁵ Cf. Richmond Thomason’s (2003) characterization of a computer as a device that “change[s] variable assignments”, i.e., that accepts certain assignments of values to variables as input, changes (i.e., manipulates) them, and then outputs the changed values. (The I/O processes do not actually have to be part of the computer, so-defined.)

could say that the system can be *said* to understand Chinese. But this must be to describe the situation from a third-person point of view, say, the point of view of the native Chinese speaker outside the room. A third alternative might be this: Chinese NLU is being produced. This more neutral description seems to be a “view from nowhere” (Nagel, 1986), but can there be understanding without an understander? (Cf. my discussions of Taylor, 2002 and of Proudfoot, 2002 in Rapaport, 2006).

Whether it is Searle-in-the-room or the entire CR system that is alleged to understand Chinese, or whether “it is understanding” in the same way that we say “it is raining” (i.e., understanding is occurring without anything *doing* the understanding, just as it rains without a “rainer”), Searle-the-philosopher argues that the mere syntactic symbol-manipulation undertaken in the Room does not suffice to yield semantic understanding (Searle, 1980, 1993). What appears to be missing are meanings to be “attached” to the squiggles (Searle, 2002, p. 53; cf. Rapaport, 2006). In the “robot reply” to the CRA, these enter the Room via sensors.⁶ According to Searle-the-philosopher, these are just more symbols, not meanings, so nothing is gained. But more symbols is all that a cognitive agent gets; hence, they *must* suffice for understanding, so nothing is lost.

Syntactic semantics

How could Searle-in-the-room come to have knowledge of the semantics of the (Chinese) squiggles? How can *we* have knowledge of the semantics of *our* language? This is the challenge posed by Searle’s CRA. Allegedly, such knowledge is required in order for Searle-in-the-room to “understand” Chinese (more generally, to “understand” NL; more generally still, to “understand” *simpliciter*). As we have seen, however, the processor in a CR does *not* need to have knowledge of the semantics of the input, for it is not the processor who needs such knowledge or who understands. Rather, it is the entire system that needs such knowledge or that understands. In Rogers’s system, it is not L who computes, but “the entire system M” consisting of L, the program, and the box itself—i.e., the Turing machine.

For Searle-the-philosopher, this is the task of explaining how Searle-in-the-room (or the system) could know what the symbols *are about*, not (merely) what their grammatical syntax is, i.e., how they are related *to other things* (their meanings), not merely how they are related *to each other* (their grammatical syntax). Of course, Searle-in-the-room (or the system) also needs to know the grammatical syntax. Given such knowledge, how much semantics can be learned? Quite a bit, or so I have argued in a series of earlier essays on the theory of “syntactic semantics”.⁷ This theory has three basic theses (cf. Rapaport, 2002).

Thesis 1

A computer (or a cognitive agent) can take two sets of symbols with relations between them and treat their union as a single syntactic system in which the

⁶ Effectors are also provided, to enable Searle-in-the-room or the system to manipulate the environment, though this may be less essential, since almost no one would want to claim that a quadriplegic or a brain-in-a-vat with few or no effectors was not capable of cognition. Cf. Maloney (1987, 1989); Rapaport (1993, 1998, 2000); Anderson (2003, Sect. 5); Chrisley (2003, fn. 25).

⁷ Rapaport (1985, 1986b–1990, 1993–2003a, 2005b, 2006).

previously “external” relations are now “internalized”. Initially, there are *three* things: two sets (of things)—which may have a non-empty intersection—and a third set (of relations between them) that is external to both sets. One set of things can be thought of as a cognitive agent’s mental entities (thoughts, concepts, etc.). The other can be thought of as “the world” (in general, the meanings of the thoughts, concepts, etc., in the first set). The relations are intended to be the semantic relations of the mental entities to the objects in “the world”. These semantic relations are neither among the agent’s mental entities nor in “the world” (except in the sense that *everything* is in the world). At best, they are accessible only from a third-person point of view (though I will argue later that even that is not the case). In a CR, one set might be the squiggles, the other set might be their meanings, and the “external” relations might be semantic interpretations of the former in terms of the latter.

When the two sets are unioned, the meanings become “internalized”. An agent understands the world by “push[ing] the world into the mind” (Jackendoff, 2002, Sect. 10.4; cf. Rapaport, 2003a). Now the agent’s mental entities include *both* the original thoughts, concepts, etc., *and* representatives of the (formerly) external objects, *and* representatives of the (formerly) external relations. Before the union, the *semantic* relations obtained between *two* sets; now, their mental analogues obtain within a *single* set. *Hence, they are syntactic*: Semantics is the study of relations between two sets, whereas syntax is the study of relations among the members of a single set (Morris, 1938).

As an example, Stanislas Dehaene (1992, pp. 30–32) argues that “numbers may be represented mentally in three different codes”: an “auditory verbal code” (i.e., a mental “word sequence”), a “visual arabic [sic] code” (i.e., mental Arabic numerals), and an “analogue magnitude code” (i.e., a mental number line). The first two are “purely” syntactic representational systems. The third is a “semantic” syntactic system: it is syntactic, because it is mental—implemented in the very same kind of neurons as them. Yet it provides an internal semantic interpretation for the other two codes.

In this way, the syntax of the new, unioned set *can* suffice for semantics; i.e., it can syntactically handle the semantic system. In this way, computers and cognitive agents *can* understand NL via a syntactic, holistic, conceptual-role semantics, with the help of negotiation (Rapaport, 2002, 2003a).⁸ Note that, for a network, any semantics that does not go outside the network must be a conceptual-role semantics (Rapaport, 2002).⁹

Thesis 2

Some might prefer to say that the purely syntactic system (merely) “models” the semantic relationship between the symbols (e.g., the Chinese squiggles) and their (real-world/external-world) meanings. But I take the internalization of these real-world/external-world meanings seriously; this is, in fact the second thesis of the

⁸ And/or some kind of “dynamic” or “incremental” semantics along the lines of, e.g., Discourse Representation Theory (Kamp & Reyle, 1993) or Dresner’s (2002) algebraic-logic approach.

⁹ Such a theory has received a partial computational implementation in our research project on “contextual vocabulary acquisition”: Ehrlich (1995, 2004); Ehrlich and Rapaport (1997, 2004); Rapaport (2003b, 2005a); Rapaport and Ehrlich (2000); Rapaport and Kibby (2002); Kibby, Rapaport, Wieland, & Dechert (forthcoming).

theory: we only have direct access to internal representatives of objects in the external world. I discuss this further in Sects. ‘The Japanese Room’ and ‘The Library Room’.

Thesis 3

The third thesis is that understanding is recursive. Briefly, we understand one system (a syntactic domain) in terms of another (a semantic domain) that is antecedently understood. The base case must be a system that is understood in terms of itself, i.e., syntactically. We will look at this at the end of the essay (Sect. ‘Concluding remarks’).

Discussion

Syntactic semantics shows how syntax suffices for semantics (contra Searle, 1980), allowing Searle-in-the-room to “escape” from the CR (Rapaport, 2000). The present essay offers further evidence in support of this.

One objection is that the unioned set is “closed” or “circular” (Harnad, 1990; cf. Rapaport, 1995).¹⁰ Just as all words in a dictionary are ultimately defined in terms of other words in the dictionary (cf. Spärck Jones, 1967), in our unioned set all elements are highly interconnected and, almost by definition, not connected to anything “outside” the set. To fully understand any term defined in a dictionary, one apparently has to have some independent knowledge that links the dictionary information to the external world. First, this may not strictly be true even for dictionaries, especially if they contain pictures corresponding to internalizations of the external objects. More significantly, our brain is just such a closed system: all information that we get from the external world, along with all of our thoughts, concepts, etc., is represented in a single system of neuron firings. Any description of that real neural network, by itself, is a purely syntactic one. Yet the syntax of that real neural network is what provides our semantic understanding of language and the world. If this implies some kind of Quinean holism or indeterminacy, so be it (see note 22).

The SNePS knowledge-representation, reasoning, and acting system

Before examining some real-life CRs, we need some background that will be useful later on.

SNePS is a computational system for knowledge-representation, reasoning, and acting. It can be viewed as a propositional semantic network whose nodes represent propositions and (their constituent) concepts (including embedded propositions) and whose labeled, directed arcs encode the relationships of constituent concepts to the concepts and propositions of which they are constituents.¹¹

¹⁰ Objections to holistic theories in general are replied to in Rapaport (2002, 2003a).

¹¹ For other ways of viewing SNePS, see Shapiro and the SNePS Research Group (2006). For details, see, e.g., Shapiro (1979); Shapiro and Rapaport (1987, 1992, 1995); Shapiro (2000); and online at: [http://www.cse.buffalo.edu/sneps] and [http://www.cse.buffalo.edu/~rapaport/snepskr.html].

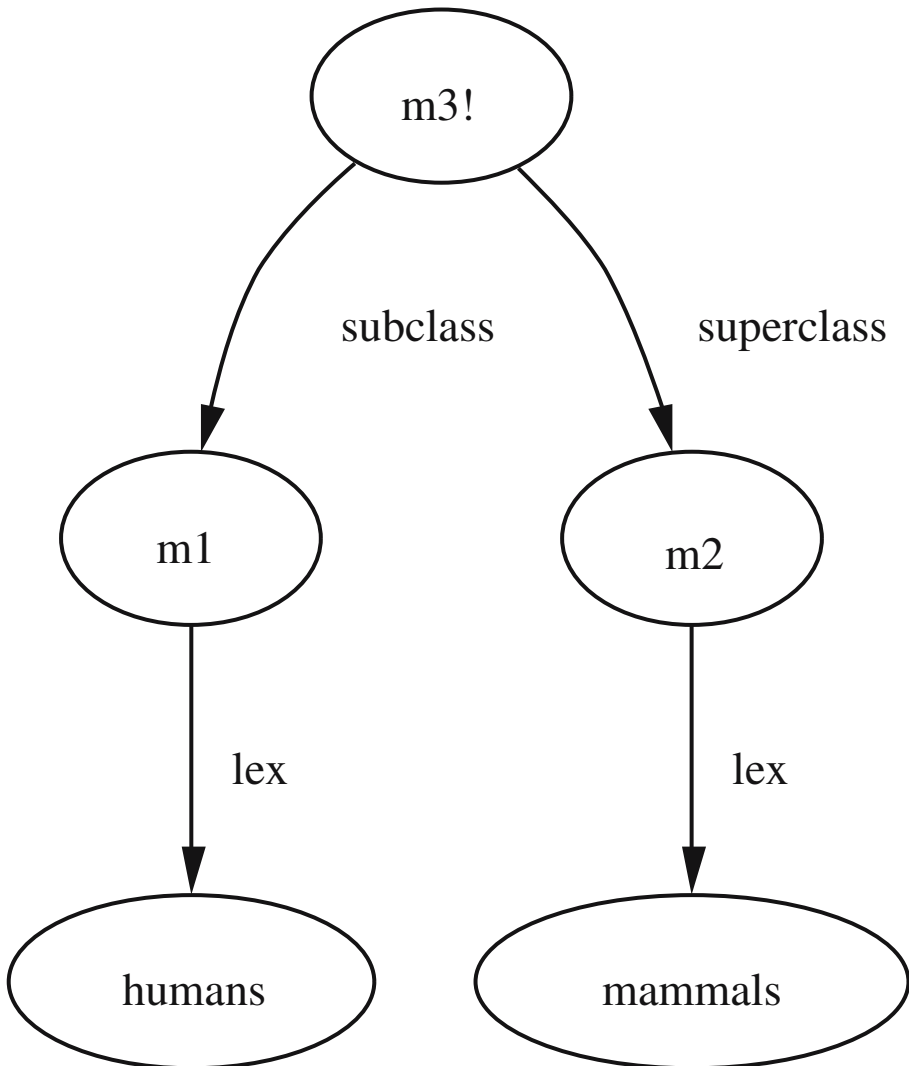


Fig. 1 A simple SNePS network. (See, Sect. ‘The SNePS knowledge-representation, reasoning, and acting system’ for an explanation. The graphical package that drew the networks for this essay places all arc labels to the *right* of their arc)

Figure 1 shows a simple SNePS network. Node *m3* represents the proposition that humans are mammals.¹² The exclamation point on the node label indicates that the proposition is “believed” (i.e., accepted as true; technically, “asserted”) by the

¹² Here, it is represented via a subclass–superclass relationship. There are other ways to represent this in SNePS, e.g., as one of the universally quantified propositions “For all x , if object x has the property of being human, then x is a member of the class mammals” or else “For all x , if x is a member of the class humans, then x is a member of the class mammals”, or in other ways depending on the choice of ontology, which determines the choice of arc labels (or “case frames”). SNePS leaves this up to each user. E.g., an alternative to the *lex* arc is discussed in Sect. ‘What really happened at the well house?’

system. The proposition is analyzed as claiming that a class represented by node *m1*, and expressed in English by the lexeme ‘humans’, is a subclass of a superclass represented by node *m2*, lexicalized as ‘mammals’.¹³ Although the network “encodes” the notions of “subclass”, “class”, and “lex”ical entry, it does not explicitly *represent* these. The only things that are explicitly represented are (a) two lexical entries (represented by: humans, mammals), (b) two entities (represented by: *m1*, *m2*) that are characterized only structurally as concepts associated with those lexemes (the association is shown by the arcs; see Shapiro & Rapaport, 1987 on the semantics of the lex case frame), and (c) another entity (represented by: *m3*) that is an asserted proposition structurally characterized only in terms of *m1* and *m2*.¹⁴ SNePS can only “talk” about or “have beliefs” about *nodes*, not arcs.¹⁵ This network shows that Cassie (the computational cognitive agent implemented in SNePS) believes that humans are mammals.¹⁶

Information can be entered into SNePS in one of two ways: The SNePS User Language (SNePSUL) uses a Lisp-like syntax to describe networks directly in terms of their graph structure. E.g., the network of Fig. 1 can be built by first defining the arc labels *lex*, *subclass*, and *superclass*, and then issuing the command

```
(assert subclass (build lex humans) superclass (build lex mammals))
```

This can be read “inside out”: first, nodes *m1* and *m2* are built (or retrieved, if they had previously been built) with *lex* arcs emanating from them, one to a node labeled humans and one to a node labeled mammals. Second, node *m3* is built (or retrieved) with a *subclass* arc to *m1* and a *superclass* arc to *m2*. Finally, *m3* is “asserted” (i.e., treated as a proposition that Cassie “believes”, denoted by the exclamation mark; unasserted propositional nodes represent propositions that Cassie is merely thinking about without being committed to).

The second language for entering information into a network is SNePSLOG, whose syntax is like that of a language for predicate logic. In such languages, a predicate or function-symbol is characterized in terms of a *sequence* of arguments. In SNePSLOG, a function symbol¹⁷ is defined in terms of a *set* of arguments (a “case frame”) distinguished by (arc) labels. To create the network of Fig. 1 using SNePSLOG, we first define two case frames (i.e., sets of arc labels). In SNePS, the user (or

¹³ The labels on the nodes at the heads of the *lex* arcs are arbitrary. For expository convenience, I use English plural nouns. But they could just as well have been singular nouns or even arbitrary symbols (e.g., “b1”, “b2”). The important points are that the nodes (1) *represent* lexical expressions in some language and (2) are “aligned” (in a technical sense; see Shapiro & Ismail, 2003) with entries in a lexicon. E.g., had we used “b1” instead of “humans”, the lexical entry for b1 could indicate that its morphological “root” is ‘human’, and an English morphological synthesizer could contain information about how to modify that root in various contexts. See Sect. ‘A SNePS analysis of what Keller learned: preliminaries’, n. 42.

¹⁴ On the notion of “structural”, as opposed to “assertional”, characterizations, see Woods (1975); Shapiro and Rapaport (1987, 1991).

¹⁵ For further discussion, see Shapiro and Rapaport (1987) and Sect. ‘What really happened at the well house?’, below.

¹⁶ For more information on Cassie, see Shapiro and Rapaport (1987, 1991, 1992, 1995); Shapiro (1989, 1998); Rapaport, Shapiro, and Wiebe (1997); Rapaport (1991a, 1998, 2000, 2002, 2003a); Ismail and Shapiro (2000); Shapiro, Ismail, and Santore (2000); Shapiro and Ismail (2003); Santore and Shapiro (2004).

¹⁷ SNePS *only* has function symbols, no predicates. All well-formed formulas are terms; none are sentences in the usual sense, although some terms can be “asserted”, meaning that Cassie (or the system) treats them as (true) sentences.

“knowledge engineer”) is free to choose arc labels and case frames, except for some pre-defined ones used by the reasoning and acting systems (see Shapiro & Rapaport, 1987). In SNePSLOG, the user is also free to choose function symbols that, essentially, name these case frames. A function symbol (i.e., a case-frame name) can be identical to one of its arc labels; this often has mnemonic value for the human user. For clarity, however, I will usually use different names. Here are two function-symbol definitions:

1. `define-frame thing-called (nil lex)`
2. `define-frame AKO (nil subclass superclass)`

Case-frame 1 defines a function symbol that can be used to build nodes `m1` and `m2` of Fig. 1, and case-frame 2 defines a predicate (actually, another function symbol) that can be used to build node `m3`, as follows:

3. `AKO(thing-called(humans),thing-called(mammals))`.

(The period causes `m3` to be asserted.) The semantics of these case frames is as follows:¹⁸

1' `[[thing-called(w)]]` = the concept associated with (or expressed by) the lexical entry (i.e., the word) *w*.

2' `[[AKO(x,y)]]` = *x* is a kind of *y* (or: *x* is a subclass of the superclass *y*).

(The “nil” in the `define-frame` syntax serves to ensure that the function symbol is not also an arc label; hence, it plays no role in the semantics.) Thus, (3) can be understood as saying that a thing called ‘humans’ is a kind of (or: is a subclass of a) thing called ‘mammals’.

In what follows, I use SNePSLOG syntax for readability (because of its similarity to a language for predicate logic)¹⁹ and network diagrams for their graphical value.

What is it like to be in a Chinese Room?

Before seeing how to “escape” from a CR, let us consider a few *actual* situations that suggest what it might be like to be *in* one.

The Japanese Room

We can get a feel for what it might be like to be Cassie by considering the network in Fig. 2 for a Japanese-speaking computational cognitive agent implemented in SNePS (Arahi & Momouchi, 1990, p. 2). My first reaction on seeing this SNePS network was that indeed I could not understand it. But why should I? It only matters for *Cassie* (or her Japanese counterpart) to understand it. But even that isn’t quite accurate. Cassie only has to be able to understand *with* it; i.e., she only has to *use* that network in order to understand *other* things (e.g., language). After all, we humans do

¹⁸ I use double-brackets (“[[]]”) to denote a function that takes as input the symbol inside the brackets and returns a meaning for it.

¹⁹ The *logic* underlying SNePS is *not* that of ordinary predicate logic! It is a paraconsistent relevance logic with an automatic belief-revision system that can allow for any believed (i.e., asserted) proposition to be withdrawn (i.e., unasserted), along with any of its inferred consequents, in the light of later beliefs. See Shapiro and Rapaport (1987), Martins and Shapiro (1988), Shapiro (2000).

1. はじめに

自然言語理解システムの構築はマン・マシン・インタフェースおよび電子化された文書データの高利用という視点から非常に重要な課題である。しかし、自然言語を理解する仕組みの構築には、汎用性や実用性の観点から多くの問題が依然として残されている¹⁾。この問題に対して、我々は自然言語理解システム構築のための一つの基礎として、日本語文に典型的に出現する名詞述語文の理解について研究を進めている。ここで、名詞述語文とは、主語と述語の対立の中で、述語が名詞で作られる文のことである²⁾。名詞述語文については、日本語学の観点から種々の考察が行われている^{3, 4)}。しかし、計算言語学的観点から十分研究がなされているとはいえない。このような立場から我々は、従来より名詞述語文「<名詞句1>は<名詞句2>である。」を対象として、名詞句間の意味関係の学習と解析を行う手法の基礎的な考察を行っている⁵⁾。

本稿では、名詞句間の意味関係を学習・解析する手法の概要とその有効性を評価するため本手法に基づく実験システムを作成して行った実験結果について述べる。

2. 名詞句間の意味関係

名詞述語文における<名詞句1>と<名詞句2>の間の基本的な意味関係を以下に示す。

- (1) 下位・上位関係

<名詞句2>が<名詞句1>の上位概念を表わす。
例：鯨はは乳類です。
- (2) 同一関係

<名詞句1>と<名詞句2>が同一概念を表わす。
対象レベルと概念レベルでの同一関係がある。
例1：富士山は日本一の山です。
(対象レベルでの同一関係)
例2：正方形は直角正四辺形です。
(概念レベルでの同一関係)
- (3) 対象・属性関係

<名詞句1>が対象で<名詞句2>が<名詞句1>の属性値を表わす。
例：太郎は医者です。
- (4) 対象・事象関係

<名詞句1>が対象で<名詞句2>が<名詞句1>の関与する事象を表わす。
例：父が帰宅です。
- (5) 要素・集合関係

<名詞句1>が<名詞句2>の表わす集合の要素となっている。
例：雄は性別です。
- (6) うなご文関係⁶⁾

「太郎は鯨です。」に代表される文で文脈に依して、例えば「太郎は鯨を食べます。」という意味に解釈される。これは、太郎という名前の鯨がいて文字どおりそれが鯨であるという解釈も可能である。
- (7) 比喩関係

<名詞句1>の属性に<名詞句2>の顕著な属性を重ねる隠喩を表わす。
例：太郎は鯨です。
(太郎は鯨のようにぬらりくらりと言いつをする。)
- (8) 同語反復同一関係

<名詞句1>と<名詞句2>に同じ名詞句を置き、文脈(状況)の中で修飾的解釈が行なわれる。

例：太郎は太郎です。自分で考え、自分で行動すべきです。

ここでは、文字どおりの意味である(1)~(5)を対象とし、意味関係の学習と解析を行う。文字どおり以外の意味である(6)~(8)についての学習と解析については、今後研究を進め別の機会に報告したい。

3. 知識表現

我々は、知識表現として、S. C. Shapiroらによって提案された意味ネットワークSNePS^{7, 8)}に基づく表現を用いている。図1に「人間はは乳類です。」のSNePS表現の例を示す。SNePSでは節(nodes)が、命題、対象、属性、関係などを表わし、弧(arc)はそれらの間の構造的、意味的つながりを表わす。ここで、ある概念からある概念までの弧の連なりを二つの概念間のパス(経路)と呼ぶ。図1で<lex>につながる<人間>と<は乳類>は語彙であり、<m1>と<m2>は<人間>と<は乳類>の概念であり、<m3>は<m1>と<m2>が下位・上位関係であるという概念である。




図1 SNePSの例

ここで用いる概念生成モデルは、事物の概念である具象概念を事物の性質に対応する属性概念によって制限することにより新たな概念である派生概念を生じ、それが具象概念あるいは属性概念として知識の中に組み込まれてゆくというものである。したがって、概念は具象概念、属性概念、派生概念のいずれかに分類される。これは、学習の段階でヒューリスティックスを用いて行われる。

ここで、2章で述べた(1)~(5)までの意味関係をネットワーク上での弧で表現する際の表記について述べる。ここで、(1)、(3)、(4)、(5)については、それぞれ<名詞句1>を始点とするか<名詞句2>を始点とするかで2通り考えられるので、実際には以下のような9つの関係が考えられる。

- (1) 下位・上位関係 : <-subclass--:superclass->
- 上位・下位関係 : <-superclass--:subclass->
- (2) 同一関係 : <-equiv--:equiv->
- (3) 対象・属性関係 : <-object--:property->
- 属性・対象関係 : <-property--:object->
- (4) 対象・事象関係 : <-object--:event->
- 事象・対象関係 : <-event--:object->
- (5) 要素・集合関係 : <-member--:class->
- 集合・要素関係 : <-class--:member->

4. 概要

Fig. 2 A Japanese-speaking computational cognitive agent (Arahi & Momouchi, 1990, p. 2).

not need to understand our neuron firings in order for us to be able to understand with, or by means of, our neuron firings.

I, of course, can only understand the semantic network in Fig. 2 by mapping it to my concepts, and there's insufficient information in that figure alone for me to do that in any but a non-arbitrary way. Each SNePS network in Arahi and Momouchi (1990) (as in Fig. 2) has English arc labels but Japanese node labels for the nodes at the heads of lex arcs (call these "lex nodes", for convenience). The Japanese

networks' use of English arc labels makes it *appear* that the arc labels convey some information to Cassie. They don't. They only convey information to *us*; but that's irrelevant (McDermott, 1981). They serve as "punctuation" or structuring devices only. True, the "reader" of a network uses the arc labels to help understand the network, but that's akin to the use of self-documenting variable names in a high-level programming language: useful to an "external" reader but not essential to the compiler or interpreter. True, too, SNePS's NL module uses the arc labels to generate appropriate English expressions (see Sect. 'A SNePS analysis of what Keller learned: preliminaries', below), but—again—the arc labels could equally well have been more-or-less arbitrary strings of characters. What counts is *that* the arcs are labeled in certain ways, not *what* the labels are. The "*content*" of the networks is in the nodes (in particular, in the *lex* nodes; cf., e.g., Rapaport, 1988b).

Nonetheless, the networks seem incomprehensible! When I first saw them, I felt like Searle-in-the-room. I can manipulate the networks, but I don't know what they mean. What's missing is a link—a correspondence—between the Japanese *lex* nodes and *my* neural analogues of *lex*-nodes *or else* between the Japanese *lex* nodes and things in the world. The latter, though, is impossible; at best that would be *representations* in *my* mind of things in the world. So this reduces to the first alternative (cf. Rapaport, 2000, 2002, 2003a).

At this point, recall Thesis 2, the Kantian thesis that we cannot have direct access to "things-in-themselves", i.e., to real-world objects external to our minds. We can only have access to internal entities that are representatives of them, filtered through our cognitive mechanisms. (Following Shapiro & Ismail, 2003, I will usually call the external items "objects" and the internal ones "entities".) (The legion of readers who may disagree with Thesis 2 can consult Rapaport, 2000, 2005b for my reasoning, which is a version of the Argument from Illusion. See also Sect. 'The Library Room', below.)

On this view, semantic relations obtain between two systems of mental entities. One might be internal representatives of Chinese squiggles; the other might be internal representatives of external objects. The latter would have to be augmented with other mental entities that might not have direct external counterparts (e.g., such "abstract" concepts as love, justice, $\sqrt{-1}$, etc., as well as "non-existents"—or Meinongian objects—such as unicorns, golden mountains, or round squares).²⁰ Each of these two systems will have its own grammatical syntax. (The syntax of a set of squiggles will be a grammatical syntax in the linguistic sense. The syntax of a set of internal representatives of external objects will be an ontology, in the AI sense.) The fully internal relations between the two systems (which can be considered as internal representatives of the classical word-world semantic relations) will be purely syntactic, too. More precisely, their union will have two kinds of (internal) relations: grammatical ones and semantic ones. E.g., in NL, there are grammatical properties and relations among the words in the sentence "Ricky loves Lucy" (subject, verb, object), and there are semantic relations between the words 'lawyer' and 'attorney' (synonymy), or 'tall' and 'short' (antonymy). But, in the classical sense of syntax (Morris, 1938), both of these sorts of relations are syntactic.

If all items are internal (e.g., all are represented by neuron firings), then they are purely syntactic. The alternative is that we are restricted to relations between internal entities and external objects. In that case, we are in no different a situation

²⁰ Cf. Meinong (1904); Rapaport (1976, 1978, 1981, 1985/1986, 1991b).

than a computer: any relations that we have with the external world could be had by a computer. But, if we can understand the semantics of NL in this way, so can a computer. Note, though, that such relations are neither in our minds nor among the objects in the world that we are trying to understand (to repeat: except insofar as everything is in the world), so it is difficult to account for how we would use or understand them.

Here is another example: I recently attended a lecture on phonology that was almost incomprehensible to most of the non-phonologists in the audience because of the speaker's extensive use of unexplained technical jargon. Most of the unfamiliar terms were names or descriptions of certain (kinds of) sounds and of various parts of our anatomy. I felt as if I were in a CR listening to aural squiggles. My first reaction was Searlean: I wanted to know the meanings of these terms—I needed *semantics*. Upon calmer reflection, I realized that what I really needed was to *hear the sounds* or to *see pictures* of the anatomical structures. A computer would need, and could have, the same *additional input*—i.e., annotations (as in the Semantic Web). But such input, together with the original terms and links among them, would just be more internal syntax. No external semantic links would be needed.

In the Semantic Web, Web pages containing text and images are associated (or “annotated”, typically on their source pages) with other symbolic (i.e., *syntactic*) information that provides “semantic” information about the Web-page content. (For this reason, Ceusters, 2005 cynically refers to the Semantic Web as the “Syntactic” Web!) Tim Berners-Lee appears to be of two minds concerning this. In 1999, he said

a piece of information is really defined only by what it's related to, and how it's related. There really is little else to meaning. The structure is everything The brain has no knowledge until connections are made between neurons. All that we know ... comes from the way our neurons are connected. (Berners-Lee & Fischetti, 1999, p. 12; cf. pp. 184ff.)

This sounds very much like a purely syntactic, conceptual-role semantics. Yet, a few years later, he wrote this (perhaps under the influence of his co-authors)

The computer *doesn't truly “understand”* any of this information [provided by the Semantic Web], but it can now *manipulate the terms much more effectively*²¹ in ways that are useful and meaningful to the human user. (Berners-Lee et al., 2001; my emphasis.)

We are not told what “true understanding” is. Perhaps it is first-person understanding—understanding *for the computer* (or for the Web itself?)—as opposed to third-person understanding *for a “human user”*. The theory of syntactic semantics claims that first-person understanding is just this sort of syntactic manipulation, and the more ways there are to “manipulate the terms”, the more understanding there will be. A cognitive agent (including a cognitive computer) *can have* such first-person—or “intrinsic”—understanding: e.g., a computer that can calculate greatest common divisors (GCDs) can be said to understand them if it has a sufficiently large network of information about GCDs and mathematics in general (Goldfain, 2004, 2006; Rapaport, 1988b).

²¹ As Goldfain (personal communication, 2006) put it: It can *understand with* the Semantic Web information.

Searle's frustration stems from his desire to understand such networks semantically. But the only option open to him is to understand them syntactically. Granted, there aren't enough samples of the Japanese networks for *me* to be able to make much sense out of them. But given enough networks (produced by encounters with new contexts), I should be able to.²²

The Library Room

But where is the *experience*, the qualium, the feeling of "Aha! Now I understand" (Galbraith & Rapaport, 1995)? The short answer is that the understanding is "in" the complex network. On the face of it, that's not very satisfactory: If I can detect no understanding from a few networks, why should it arise from a few more? The long answer is that the more "interlocking" networks there are, and the more experience the person in the room has in manipulating them, the more understanding there will be. How might such understanding and experience of understanding arise?

Consider a second CR: my understanding of Library of Congress catalog numbers. I have neither studied nor know the rules of syntax or semantics for them, but I do not need to in order to use them "fluently" to find books. However, through frequent use of them, I have come to learn (inductively) certain rules of syntax, such as

LOC-catalog-number ::= letter₁ + numeral + '.' + letter₂ + numeral + year-numeral

and of semantics, such as

- [[letter₁]] = a category of certain books;
- usually, [[letter₂]] = initial letter of author's last name;
- usually, [[year-numeral]] = year of publication.

The more links I make with my knowledge of books and libraries, the more I know of the syntax and semantics; and the more I know of the syntax and semantics, the more I understand of what I'm doing. Searle-in-the-Library-of-Congress-room would also come to have such understanding. Why shouldn't Searle-in-the-CR?

Matching LOC codes with actual books requires some experience with the external world. So it might be objected that I am not dealing only with a syntactic symbol system, but—inevitably—with such a system *and* with the real world. Objections to theories are best viewed as problems that the theory must deal with (Castañeda, 1980, 1984; Rapaport, 1982). My preferred way of dealing with this particular problem is, once more, Thesis 2, namely, internalizing my experiences

²² Would I be able to do so uniquely or "correctly"? Or would the theoretical existence of an infinite number of models for any formal theory mean that the best I might be able to hope for is an understanding that would be unique (or "correct") "up to isomorphism"? Or (following Quine, 1960, 1969) might the best I could hope for be something like equivalent but inconsistent translation manuals? John McCarthy (personal communication, April 2006) thinks that I would eventually come to a unique or correct understanding. These important issues are beyond the scope of this essay. For present purposes, it suffices that the understander be able to make *some* sense out of the networks, even if it is not the intended one, as long as the understanding is consistent with all the given data and modifiable (or "correctable") in the light of further evidence (Rapaport, 2005a; Rapaport & Ehrlich, 2000).

with the real world. In the present case, I go to the library with an LOC number to find the book whose number it is, I find the book, see that the letter categorizes it along with other books on roughly the same topic, see that the letter is the initial letter of the author's last name, and so on. *But my visual perceptions of these things are internal*—they are images in my mind, *caused* by external objects, but as much in my mind as are my mental images of the LOC numbers. The semantic relations between LOC numbers and their meanings are all (internal) syntactic relations among two sets of (mental) symbols.

If you don't agree with this argument, then, for me to explain what the semantics of LOC numbers is, I need to describe the numbers and their meanings to you. But those descriptions of their meanings—of the real-world objects—are syntactic (linguistic) items that are themselves representatives of the real-world objects. Once more, we are stuck with syntactic relations between two symbol systems (Rapaport, 1995; Smith, 1982).

The Helen Keller Room

Has anyone ever *really* been in a CR for a sufficiently long time? Yes—Helen Keller:²³

The morning after my teacher [Anne Sullivan] came she ... gave me a doll. ... When I had played with it a little while, Miss Sullivan slowly spelled into my hand the word “d-o-l-l.” I was at once interested in this finger play and tried to imitate it. When I finally succeeded in making the letters correctly I was flushed with childish pleasure and pride. Running downstairs to my mother I held up my hand and made the letters for doll. *I did not know that I was spelling a word or even that words existed; I was simply making my fingers go in monkey-like imitation.* (Keller, 1905, p. 35; my emphasis.)

At the beginning of this passage, one expects that the antecedently-played-with doll would be associated with the finger-spelled word ‘d-o-l-l’. But, as can be seen from Keller's later claim of ignorance (“I did not know ...”), her statement that she “made the letters *for doll*” (my italics) must be taken *de re*,²⁴ since, clearly, Keller did not know that she was “making ... *letters*” (my italics) or that they were “for doll”.

²³ A similar observation has been made by Justin Leiber (1996, esp. p. 435):

[T]he suspicion that Keller did not live in the real world, could not mean what she said, and was a sort of symbol-crunching language machine ... suggests a prejudice against the Turing test so extreme that it carries the day even when the Turing test passer has a human brain and body, and the passer does not pass as simply human but as a bright, witty, multilingual product of a most prestigious university, and professional writer about a variety of topics.

For readers unfamiliar with Helen Keller (1880–1968), she was blind and deaf because of a childhood illness (see below), yet graduated from Radcliffe College of Harvard University, wrote three autobiographies, and delivered many public lectures on women's rights, pacifism, and helping the blind and deaf.

²⁴ When I was very young, I heard what I took to be a single word that my parents always used when paraphrasing something: ‘inotherwords’. (At the time, I had no idea how to spell it.) It took me several hearings before I understood that this was really the three-word phrase ‘in + other + words’. Similarly, from Keller's point of view, her finger spellings were not letters + for + dolls, but an unanalyzed “lettersfordolls”.

The italicized last sentence of Keller's quote is significant: it is a wonderful description of pure syntax; Searle would be pleased. Sullivan, playing native-Chinese speaker to Keller's Searle-in-the-room, no doubt would have had reason to believe that Keller *did* know what she was doing.

The passage continues (Keller, 1905, p. 35)

In the days that followed I learned to spell in this uncomprehending way a great many words, among them *pin, hat, cup* and a few verbs like *sit, stand* and *walk*. But my teacher had been with me several weeks before I understood that everything has a name.

Here, 'name' is synonymous with 'word'; for Keller and Sullivan, even the adjective 'very' is the "name" of a "thing" (Keller, 1905, p. 261). Although sounding a bit odd to modern philosophical ears, this is not entirely idiosyncratic: "When saying something of an object, one always uses a name of this object and not the object itself, even when dealing with linguistic objects" (Tarski, 1969).

Again, these descriptions of Keller's own experiences, given long after the fact, are best understood *de re*. She experienced external things, and she experienced meaningless finger manipulations, but she did not link them in an appropriate way. Such linking between experiences of a word (a finger spelling) and experiences of an external object should have yielded semantic understanding. According to syntactic semantics, semantic understanding would actually have come via Keller's linking of *internal* representations of *both* of those external experiences. They would have played different roles: one the role of syntax, one that of semantics. One (the finger spellings) would not have been comprehended; the other (the physical objects) would have been antecedently familiar (they were, after all, part of the world she lived in every day). One would have been a name for a thing; the other, a thing named.

Keller succeeded remarkably well in her understanding of NL. Here is how she described her experience immediately after the well-house²⁵ event during which one of her hands was immersed in water while Sullivan finger-spelled 'w-a-t-e-r' in her other hand and Keller allegedly first realized that the finger-spellings had meaning (we'll examine this in detail later)

As we returned to the house every object which I touched seemed to quiver with life. That was because I saw everything with the strange, new sight that had come to me. (Keller, 1905, p. 36.)

This would eventually be the experience of Searle-in-the-room, who would then have both *semantic* methods for doing things and purely *syntactic* ones. The semantic methods, however, are strictly internal: they are not correspondences among words and things, but *syntactic* correspondences among *internal nodes* for words and things.

Swan (1994) points out how important Keller's *hand* was to her ability to communicate.²⁶ At the well house, both the name and the object were communicated via the *same* sense modality: touching her hand. He also points out how she had to learn about the visual aspects of the world *as expressed via language* (rather than

²⁵ Keller referred to a "well"-house, whereas Sullivan referred to a "pump"-house. Keller's house is now a museum; their website (helenkellerbirthplace.org) also calls it a "pump". I use Keller's term in this essay.

²⁶ On hands and cognition, cf. Papineau's (1998) review of Wilson (1998).

perceived). This lends credence to the view that Keller's understanding of language and the world was an internal understanding, with all input encoded into a single syntactic system. Mental concepts for both words and objects were built on the basis of tactile (and olfactory) sensation. One of the reasons the well-house episode was significant was that it was the event that enabled Keller to distinguish some of her concepts from others, categorizing some as representing the world and others as names of the former. For Keller, initially, language and the non-linguistic part of the world were indistinguishable.

Swan discusses the difficulty for Keller of distinguishing between self and other, between her words and those of others.²⁷ Before the well-house episode, she could use signs, but had difficulties with certain ones, in particular, with those for container versus contained ('mug' versus either 'milk' or 'water'; see Sect. 'Names for things: from "monkey-like imitation" to natural-language understanding', below). Perhaps, *before* the well-house episode, she could not distinguish words from (other) objects: words *were* objects, part of the holistic fabric of her world. *After* the well-house episode, she could distinguish between two kinds of objects in the world: (non-linguistic) objects and words for them. That placed a syntactic-semantic structure on her mental network. And it resulted, as we know, in the blossoming of her understanding. Searle-in-the-room could do no worse.

Gilbert Harman (1987, p. 57) has said that

a language, properly so called, is a symbol system that is used both for communication and thought. If one cannot think in a language, one has not yet mastered it. A symbol system used only for communication, like Morse code, is not a language.

Before the well house, Keller used symbols to communicate, but not to think. Her pre-well-house signs were like non-understood Morse code. (Do you understand what "... . -.. -.. - -" means if you know no Morse code? But if I finger-spelled that into your hand each time I saw you, and forced you to do the same to me, would you get the hang of doing it? Would you *then* understand it?)

Unless the symbols are part of a larger network, they have very little meaning. To that extent, perhaps Searle has a point. But the more they *are* used for thinking, the more language-like they are. And they *have* to be part of a larger network—separated into syntactic and semantic regions—else they could not be used to communicate. They have meaning if and only if (and to the extent that) they are part of such a larger, "partitioned" network.²⁸ Searle denies the "if" part of this, i.e., that being part of such a network confers meaning. Keller, I suggest, was a living counterexample.

Names for things: from "monkey-like imitation" to natural-language understanding

A puzzle

The case of Helen Keller offers a real-life CR situation, and I have given some reasons why the epiphenal well-house episode—paradigmatic of the relationship

²⁷ Keller had been accused of plagiarism, when, in fact, it is possible that she had merely made a grievous use-mention confusion, viz., not having learned how to use quotation marks; cf. Keller (1905).

²⁸ As Goldfain pointed out to me, 'partitioned' needs scare quotes, because I don't want to imply an empty intersection between the syntactic domain and the semantic domain. Cf. Sect. 'Thesis 1', above.

between a syntactic domain and a semantic domain, with Keller simultaneously having one hand immersed in syntax and the other in semantics—was so significant for her.

But why should it have been? By Keller's and Sullivan's own testimony, Keller seemed able to use (finger-spelled) words for things, as well as (self-invented) signs and gestures for things, *before* the well house. What made the well house so significant?

What did Helen Keller understand, and when did she understand it?

It is not easy to determine the chronology of Keller's language learning. There are at least two distinct, if not independent, first-person accounts: (1) the student's—Keller's (1905) autobiography, written, of course, long after the events, and (2) the teacher's—Sullivan's contemporaneous letters, which are probably to be trusted more.²⁹ As Keller herself says, “When I try to classify my earliest impressions, I find that fact and fancy look alike across the years that link the past with the present. The woman paints the child's experiences in her own fantasy”³⁰ (Keller, 1905, p. 23). Even though Sullivan's letters are “incomplete” as scientific “records” (Macy, in Keller, 1905, pp. 224, 239; cf. p. 241), they are the closest thing available. Together, the two sources provide a reasonable—if tantalizing—picture. For Keller's earliest years, however, we must rely on her own report, since Sullivan was not present.

Before the illness

Keller was born, healthy, on 27 June 1880. At 6 months, she could speak a few words or phrases, e.g., ‘How d’ye’, ‘tea’, and—significantly—‘wah-wah’ (“water”) (p. 25).³¹ But did she understand them?

My son at the same age could only say “Da-da-da”, and said it for the first time when *not* in my presence. My granddaughter at the same age did not even say that much. The difficulty of interpreting what infants “say”, hence of deciding what they understand, is illustrated nicely in a “For Better or for Worse” cartoon in which a girl (Elizabeth) holds her baby sister (April), who is looking down (presumably at their dog) and saying “Goggg-gog Gogg! Go-Go Gogg Gogog”. Elizabeth says, “Mom! Listen! April's trying to say *doggie!*” Their mother explains, “She's just making sounds, honey. She doesn't know what she's saying.” Then April, in the presence of her bottle, says, “Babab. Ba Ba. Ba-ba. Bababa. Ba-ba. Ba!”, and Elizabeth exclaims, “She's saying Ba-Ba for bottle!” Again, their mother rationally explains, “She doesn't know the word ‘bottle’, Elizabeth. It's all just noise to her.” April then looks at her mother and utters, “Mum-mum mamamamama ma-ma”; the

²⁹ Trusting a third-person over a first-person viewpoint is consistent with trusting the native Chinese speaker's viewpoint over Searle-in-the-room's (Rapaport, 2000). There are also Keller's somewhat more contemporaneous letters, but these—while intrinsically interesting, and exhibiting (especially in the early ones) her gradual mastery of language—do not contain much information on *how* she learned language. There are also Sullivan's speeches and reports. Although they contain some useful information and some valuable insights—especially into the nature of teaching—they, like Keller's autobiography, were written *ex post facto*; cf. Macy, in Keller (1905, p. 278.)

³⁰ This has an overtone of holistic reinterpretation (Rapaport, 1995, Sect. 2.6.2); we understand the present in terms of all that has gone before, and the past in terms of all that has come after.

³¹ Such page citations are to Keller (1905) unless otherwise indicated.

astonished mother calls her friend Connie on the phone to announce, “It’s true, Connie! I heard her! She said Ma-Ma!”. Similarly, my granddaughter at age 1;2,21³²—still without language but babbling with English intonation—pointed to a swing and said something that sounded to us like “ist?”; *we* interpreted her as saying ‘this’. Linguistic communication is at least 50% the *hearer’s* interpretation.

The illness

In February 1882 (age 1;7), Keller contracted the illness that left her deaf and blind, and, like many deaf children, she did not learn to speak. Nevertheless, she could make sounds—again, significantly, the sound ‘wah-wah’ for “water”, which “I ceased making ... only when I learned to spell the word” (p. 25).

For another anecdotal comparison, at about this age (1;6,23–1;7,5), my son could say (impolitely!) to our neighbor’s 2-year-old daughter Marielle, “Shut up, Ma’iel”; he used and seemed to understand “OK” and “good”; he could say “See ya” to his older brother leaving the house; he used “zipper” and “wait” correctly; etc. I noted in my diary at the time that each day he had many new words and new abilities to comprehend. It occurred to me that he was not merely learning to speak English; rather, he already spoke a “foreign” language that his mother and I could interpret to a limited extent. Words meant different things for him than for us. For example, ‘Bambi’ sometimes meant “videotape”; ‘water’ sometimes meant “liquid”. But these seemed neither to be overgeneralizations nor metaphors; such characterizations, in any case, are *de re*, third-person descriptions. From my son’s (*de dicto*) point of view, those words *didn’t* mean what *we* think they do.

After the illness

After her recovery, Keller could communicate via touch, via “crude signs”—

A shake of the head meant “no” and a nod, “yes”, a pull meant “come” and a push meant “go” (p. 27)

—and via (other) imitative motions, including some rather complex ones, e.g.:

when I wanted to go egg-hunting, ... I would double my hands and put them on the ground, which meant something round in the grass, and Martha [a childhood friend] always understood. (p. 28.)

She familiarized herself with the outdoors, guided by her sense of smell. This continued well after Sullivan’s arrival: they often studied outside (p. 43; cf. Sullivan’s comments on the significance for Keller of the sense of smell, pp. 293ff). She also, of course, had a sense of taste, learning thereby that the ocean was salty—a bit of commonsense knowledge that she lacked because no one thought to tell her such an obvious thing (p. 230)! She also had a “sense” of vibration, being able to sense when a door closed (p. 27). I am not sure whether to count this as part of her sense of touch, or as a remnant of a sense of hearing, which is, after all, a sensitivity to vibrations (cf. p. 208 on her ability to sense music).

³² I.e., at age 1 year, 2 months, 21 days.

Age 5

By age 5, she could perform rather complex tasks, such as folding and separating clean clothes (p. 27), and she knew that her father did something mysterious by holding a newspaper in front of his eyes. Imitating this did not illuminate the mystery (p. 30).³³ Similarly, she knew that others communicated, not with signs, but by moving their lips; imitation of this, too, was not successful (pp. 27–28).

Age 6

Fortunately, others understood her signs (p. 28). When she was 6, she tried to teach her dog some of these (with no success, of course; p. 29), though Sullivan tells a similar story (p. 253) about Keller trying to teach *finger spelling* to her dog at about the same age or a bit later (20 March 1887, to be exact—*after* Keller had begun to learn words but *before* the well house) (p. 253). She certainly, at about this time, had a desire to express herself (p. 32).

Sullivan's arrival: finger-spelling in a Chinese Room

On 3 March 1887, Sullivan arrived at Keller's home to become her teacher; Keller was now about 6;8. Almost immediately upon her arrival, Sullivan and Keller began to communicate with each other using signs and gestures (p. 245). The next day, Sullivan began teaching Keller finger spelling, presenting her with an object or action and finger-spelling its name: 'doll', 'pin', 'hat', 'cup', 'sit', 'stand', and 'walk' are the words Keller remembered. Sullivan cites 'doll', 'cake', and (sewing) 'card'. 'Cup' is of some interest, since 'mug' was to give Keller a notorious difficulty a few weeks later.

To what extent did Keller understand these words? As we saw, she herself considered this to have been "monkey-like imitation" (p. 35): finger spelling was an activity to be performed upon presentation of certain objects. It was a ritual, with a syntactic structure: there were right and wrong ways to perform it. But Keller did this "in ... [an] uncomprehending way" (p. 35) and did not yet understand "that everything has a name" (p. 35). It certainly *seems* that she was in a CR.

Was Keller really so uncomprehending at this stage? Recall that she had already developed her own system of signs and gestures for communicating her needs and wants.³⁴ Surely, this is evidence of a semantic correspondence.

I suppose that it is remotely possible that even Keller's early, self-invented signs were ritual movements performed uncomprehendingly in certain circumstances, yet rituals that just happened to convey appropriate information to others. In Robert

³³ Cf. what I have called the "miracle of reading": "When we read, we seemingly just stare at a bunch of arcane marks on paper, yet we thereby magically come to know of events elsewhere in (or out of!) space and time" (Rapaport, 2003a, Sect. 6; a typo in the original is here corrected). Clifton Fadiman once observed that

[W]hen I opened and read the first page of a book for the first time, I felt that this was remarkable: that I could learn something very quickly that I could not have learned any other way ... [I] grew bug-eyed over the miracle of language ... [viz..] decoding the black squiggles on white paper. (Quoted in Severo, 1999.)

Hofstadter (2001, p. 525) makes similar observations.

³⁴ Leiber (1996) also notes that Keller had linguistic knowledge and abilities both before and immediately after her illness.

Sheckley's story "Ritual" (1954), creatures living on a remote planet perform a series of ritual "welcoming the gods" dances as a religious ceremony. The dance in fact consists of the preparations for the arrival of a spaceship. When a spaceship finally does arrive after centuries without a landing, the villagers perform their "dance", which just happens to facilitate the spaceship landing.

But I doubt that Keller's signs were such rituals. Had *all* of Keller's gestures been such conveniently coincidental ("extensional") rituals, she would not have been able to do the complex tasks she did, or to satisfy her needs, or to have the background knowledge that became the basis for her language learning.

All that Sullivan was doing can be seen as offering Keller a new system for accomplishing her communicational goals. Keller might not have realized this, so that, for her, her *own* gestures for an object *did* constitute a semantic correspondence while Sullivan's finger spellings did not. However, that Keller *was* able to associate finger spellings with objects and actions surely indicates that she had the means to promote these to *semantic* correspondences.

There is, in fact, evidence that she did so: the day that Sullivan arrived, she taught Keller 'cake', and the next day she taught her 'card'. Keller ...

... made the "c-a," then stopped and thought, and making the sign for eating and pointing downward she pushed me [Sullivan] toward the door, meaning that I must go downstairs for some cake. The two letters "c-a," you see, had reminded her of Friday's "lesson"—not that she had any idea that *cake* was the name of the thing, but it was simply a matter of association, I suppose. (p. 246.)

I would argue that Keller *did* have the idea that 'cake' "was the name of the thing"—but that she had that idea *de re*, not *de dicto*: she did not yet have the *concept* of names for things. She could certainly associate words (i.e., finger spellings) with objects

Then I [Sullivan] spelled "d-o-l-l" and began to hunt for it. She [Keller] follows with her hands every motion you make, and she knew that I was looking for the doll. She pointed down, meaning that the doll was downstairs. ... [S]he ran downstairs and brought the doll ... (pp. 246–247),

although not without a reward of some cake.

As Keller built up a vocabulary of finger-spelled words and made mental links between (her internal representations of) these and (her internal representations of) objects and actions, she was building a "semantic network" of associated representations that she could and did use in a language-like way. Searle would argue that she did not *understand* language. I prefer to say that she did not understand language *de dicto*—she *did* understand it *de re*, in the sense that she was using it, but did not realize that she was using it or how it worked. I.e., perhaps she understood language, but did not understand that (or how) she understood it.

She was, in fact, at the same stage of language development as a normal child would have been at a much earlier age. Are we prepared to say that normal children at this stage do not understand language? Perhaps. But eventually they *do*, and eventually Keller did. Why not Searle-in-the-room or a computer? What is the crucial step (or steps) that must be taken to move from *this* level of understanding (or, if you prefer, from this level of *not* understanding) to the level that we adult speakers of language are at? We'll return to this in Sect. 'The significance of names'.

Ages 6;8,12–6;8,14

By 11 March 1887, Sullivan says that “Keller knows several words now, but has no idea how to use them, or that everything has a name” (p. 251). Yet two days later, Keller can associate words with objects: “when I give her the objects, the names of which she has learned, she spells them unhesitatingly” (p. 251).

Age 6;8,21

Around 20 March 1887, Keller reports that she was confused by ‘mug’ and ‘water’: Keller and Sullivan

... had a tussle over the words “m-u-g” and “w-a-t-e-r.” Miss Sullivan had tried to impress upon me that “m-u-g” is *mug* and that “w-a-t-e-r” is *water*, but I persisted in confounding the two. (p. 36.)

Apparently, she confused the container with the contained; perhaps this was because they always appeared together. Although she may have had a *mug* by itself, perhaps the *water* was always *in* the mug.

Here is Sullivan’s account (concerning a different liquid)

Helen has learned several nouns this week. “M-u-g” and “m-i-l-k” have given her more trouble than other words. When she spells “milk,” she points to the mug, and when she spells “mug,” she makes the sign for pouring or drinking, which shows that she has confused the words. She has no idea yet that everything has a name. (Sullivan, 20 March 1887, pp. 252–253.)

By ‘learning’, Sullivan must mean the ability to spell, to make the finger movements. ‘Mug’ and ‘milk’ (‘water’?) give Keller trouble, *not* in terms of the movements, but in terms of how to use them (what they refer to, or name). But that assumes that Keller knows that they have a *use*, which is plausible, as we’ve seen, though not altogether clear. Moreover, there are other interpretations of Keller’s actions: pointing to the mug *could* also be pointing to the milk (or water) *in* the mug, and making signs for pouring or drinking *could* refer to what one does with a *container* as well as with its *contents*. Keller’s own version suggests that she was not making *any* distinctions at all, rather than merely confusing the mug and the liquid.

Age 6;9,9

Sullivan’s second version of the confusion supports my interpretation that Keller was aware only of *events* considered as unanalyzed wholes

... “mug” and “milk” had given Helen more trouble than all the rest. She confused the nouns with the verb “drink.” She didn’t know the word for “drink,” but went through the pantomime for drinking whenever she spelled “mug” or “milk”. (Sullivan, 5 April 1887, p. 256.)

I think it is significant that Sullivan reports the confusion as between ‘mug’ and ‘milk’, where Keller reports it as between ‘mug’ and ‘water’.³⁵ First, and most importantly (if only for Freudian reasons), Keller’s one remaining *spoken* word was,

³⁵ In an earlier autobiography, Keller also called this a ‘mug’/‘milk’ confusion (p. 364). And in Sullivan’s description of the well-house episode (see Sect. ‘Epiphany’, below), she describes “w-a-t-e-r” as a “new word” for Keller (p. 257).

you will recall, ‘water’ (‘wah-wah’). Second, if Sullivan’s report is the one to be trusted, besides the semantic-domain confusion between container and contained, there might also have been a syntactic-domain confusion between two words beginning with ‘m’: recall the earlier “confusion” between ‘ca[ke]’ and ‘ca[r]d’.

Just before the well-house

There were a few days to go before the visit to the well house. What did Keller learn in those days between (a) her confusing the word for a mug with the word for its liquid contents and (b) her later epiphany? By 20 March, according to Sullivan, Keller knew 12 word–object combinations (p. 255) yet instinctively used her own signs—not finger-spelled words—to *communicate*. By 1 April, Sullivan reports, Keller’s vocabulary had increased to 25 nouns and 4 verbs³⁶—including, significantly, ‘mug’, ‘milk’, and ‘water’. Yet, two days later, Sullivan says that Keller “has no idea what the spelling means” (p. 256). I take it that, from Sullivan’s point of view, Keller’s “knowledge” of these words was at least associative and probably even communicative, yet not “conscious”. But not “conscious” in what sense? Keller apparently could ask for the finger spellings that corresponded to certain objects (the ones marked ‘x’ in note 36). What more could Sullivan want at this stage?

Searle would say that for real NLU, a lot more is wanted. I’d have to agree: Keller could not yet have passed a Turing test. So although imagining what Keller was like at this stage may give us an insight as to what Searle-in-the-room is like, there is a large gap between the two. Searle-in-the-room, after all, passes the Turing test.

Perhaps what Keller “knew” at this stage was an association of these words with certain complex, unanalyzed events, and what she learned at the well house was that the events have parts, each of which is associated with a word. If so, then what she learned was as much about the semantic domain as it was about the association between the two domains. Of course, she also presumably learned then that the *words* did not refer to complex events but only to *parts* of them. So she learned something about the syntactic domain, too.

Helen Keller and the “miracle worker” at the well house

Epiphany

The magical day was 5 April 1887. Sullivan, having failed to clarify the difference between ‘mug’ and ‘milk’, took Keller for a walk to the well house.

This morning, while she was washing, she wanted to know the name for “water.” ... I spelled “w-a-t-e-r” ... [I]t occurred to me that with the help of this new word I might succeed in straightening out the “mug–milk” difficulty. We went out to the pump-house, and I made Helen hold her mug under the spout while I pumped. As the cold water gushed forth, filling the mug, I spelled “w-a-t-e-r” in Helen’s free hand. The word coming so close upon the sensation of cold water rushing over her [other] hand seemed to startle her. She dropped

³⁶ “Doll, mug, pin, key, dog, hat, cup, box, water, milk, candy, eye (x), finger (x), toe (x), head (x), cake, baby, mother, sit, stand, walk. ... knife, fork, spoon, saucer, tea, paper, bed, and ... run” (p. 256). “Those with a cross after them are words she asked for herself” (p. 256).

the mug and stood as one transfixed. A new light came into her face. She spelled “water” several times. Then she dropped on the ground and asked for its name and pointed to the pump and the trellis, and suddenly turning round she asked for my name. I spelled “Teacher.” Just then the nurse brought Helen’s little sister into the pump-house, and Helen spelled “baby” and pointed to the nurse. All the way back to the house she was highly excited, and learned the name of every object she touched, so that in a few hours she had added 30 new words to her vocabulary. ...

... Helen got up this morning like a radiant fairy. She has flitted from object to object, asking the name of everything and kissing me for very gladness. Last night when I got in bed, she stole into my arms of her own accord and kissed me for the first time, and I thought my heart would burst, so full was it of joy. (Sullivan, 5 April 1887, pp. 256–257.)

A few observations on this passage and on the well-house episode are in order.

1. On “wanting to know the name for ‘water’ ”: clearly, Keller wanted to know the name for water (the *stuff*), not for ‘water’ (the *word*); she did not want to know the name for a name. However, Sullivan is not to be blamed for this particular use-mention confusion! On the other hand, hasn’t Sullivan repeatedly told us that Keller did not know that things have names? Then why does she report Keller as asking for the *name* of water? Perhaps this needs to be taken *de re*: note that it’s quite possible that what Keller wanted to know was the appropriate finger spelling for *washing*!
2. Sullivan’s comment about “straightening out the ‘mug-milk’ difficulty” can be interpreted as supporting my suggestion that the mug-milk confusion was one of container vs. contained or of unanalyzed events.
3. Note that, at the well house, there was little chance to “confound” two objects—there was a direct and simultaneous association of word with object. Although the mug in Keller’s hand *might* have caused some interference, Keller’s own account indicates that it did not

We walked down the path to the well-house Some one was drawing water and my teacher placed my hand under the spout. As the cool stream gushed over one hand she spelled into the other the word *water*, first slowly, then rapidly. I stood still, my whole attention fixed upon the motions of her fingers. Suddenly I felt a misty consciousness as of something forgotten—a thrill of returning thought; and somehow the mystery of language was revealed to me. I knew then that “w-a-t-e-r” meant the wonderful cool something that was flowing over my hand. ...

... As we returned to the house every object which I touched seemed to quiver with life. That was because I saw everything with the strange, new sight that had come to me. (p. 36.)

Moreover, if, indeed, it was ‘*milk*’—not ‘*water*’—that Keller had been confusing with ‘mug’, then the well-house experience was a controlled experiment, filling the mug with *water* instead of milk.

4. The finger-spelled word ‘w-a-t-e-r’ meant “the wonderful cool something that was flowing over my hand”: ‘W-a-t-e-r’ was antecedently meaningless; “the

- wonderful cool something ...” was antecedently understood. The semantic relation is asymmetric (Rapaport, 1995); here we have the intensional, asymmetric equivalence of a definition.
- Note that Keller did *not* say that ‘water’ meant H₂O: *Twin Keller* (living on Putnam’s (1975) *Twin Earth*, where ‘water’ refers to XYZ, not H₂O) would have had the *same* experience, and ‘water’ would have meant exactly the same thing for her (modulo the essential indexical ‘my’), viz., “the wonderful cool something that was flowing over my hand”.
 - Keller’s post-well-house experiences of seeing “everything with the strange, new sight” should be the eventual experience of Searle-in-the-room, who would then have *semantic* methods for doing things in addition to purely syntactic ones. Crucial to promoting semantics-as-correspondence to semantics-as-meaning—semantics-as-understanding—is that the semantic domain must be antecedently understood. This, as we shall see shortly, was crucial for Keller’s progress.

Aftereffects

Five days later, Sullivan reports Keller replacing her own signs by the corresponding finger-spelled words as soon as she learns them (p. 257). Clearly, Keller had realized the advantages of this more efficient and expressive code for communication. Equally crucially, as Sullivan observes (p. 258), Keller *understood* what the finger-spelled words referred to *before* she was able to “utter” them: “the idea always precedes the word” (Sullivan, 8 May 1887, p. 260). As Sullivan noted later (pp. 291ff), Keller had her *own* signs for things before she had *words* for them, still using her signs when she had not yet learned the words (p. 260), so she was using two codes. She had several ways to communicate her ideas, preferring one (words), but using whatever was at hand (so to speak).

Two other observations that Sullivan made are worth mentioning at this point. First, it was important for Keller to *generate* language, not merely to *understand* it, in order to help build her vocabulary (Sullivan, 16 May 1887, pp. 262ff); interactive conversation is crucial (cf. Rapaport, 2000, Sect. 8; Rapaport, 2003a). Second,

Language grows out of life, out of its needs and experiences. ... *Language* and *knowledge* are indissolubly connected; they are interdependent. Good work in language presupposes and depends on a real knowledge of things. As soon as Helen grasped the idea that everything had a *name*, and that by means of the manual alphabet these names could be transmitted from one to another, I proceeded to awaken her further interest in the *objects* whose names she learned to spell with such evident joy. *I never taught language for the purpose of teaching it*; but invariably used language as a medium for the communication of *thought*; thus the learning of language was *coincident* with the acquisition of knowledge. In order to use language intelligently, one must have something to talk *about*, and having something to talk about is the result of having had experiences; no amount of language training will enable our little children to use language with ease and fluency unless they have something clearly in their minds which they wish to communicate, or unless we succeed in awakening in them a desire to know what is in the minds of others. (Sullivan, p. 317.)

Jerome Bruner has observed much the same thing (1983, pp. 103–104)

So at the end of this first round of examining the simplest form of request—asking for objects—we are forced to a tentative conclusion. Language acquisition appears to be a by-product (and a vehicle) of culture transmission. Children learn to use a language initially (or its pre-linguistic precursors) to get what they want, to play games, to stay connected with those on whom they are dependent. In so doing, they find the constraints that prevail in the culture around them embodied in their parents’ restrictions and conventions. The engine that drives the enterprise is not language acquisition per se, but the need to get on with the demands of the culture. ... Children begin to use language ... not because they have a language-using capacity, but because they need to get things done by its use. Parents assist them in a like spirit: they want to help them become “civilized” human beings, not just speakers of the language.

This is an insight that—beyond its evident importance for education *in general*—is of importance for *computational* NLU systems, too. Knowledge (especially “prior” or “background” knowledge) is also important for figuring out a meaning for an unfamiliar word “from context”. Indeed, the “context” from which this can be done must include not only the unfamiliar word’s *textual* context (i.e., the surrounding words, or “co-text”) but also the reader’s prior knowledge (Rapaport, 2003b).

Keller’s language learning proceeded apace after the well house. Two months later, she wrote her first letter to a friend (p. 123). Her vocabulary learning was cyclic and recursive—each new encounter with a word serving to clarify and enhance what she already knew (p. 40).

Words for abstract concepts (e.g., ‘love’, ‘think’)—concepts that could not be “shown”, hence for which there was nothing apparent to associate them with—were harder for her to learn, but not impossible (for the details, see pp. 40f, 300). In April 1887, she learned prepositions by direct experience—standing *on* a chair or *in* her wardrobe (Sullivan’s account, p. 279). Keller’s own account of learning sentence structure is reminiscent of Russellian propositions: she would paste pieces of paper with words written on them “in raised letters” (p. 279) onto the things they named: She would put her doll on the bed, the doll labeled ‘doll’, the bed labeled ‘bed’, with labels for ‘is’ and ‘on’ placed near the doll, on the bed; or she would put the label ‘girl’ on herself, the labels for ‘is’, ‘in’, and ‘wardrobe’ on the wardrobe, and then she would stand in the wardrobe, thus labeled.

Over a year later, by which time her language was of Turing-test quality, she would, nonetheless, use some not-yet-understood words in “parrotlike” fashion (Macy, p. 134) until she learned how to use them properly (until she learned their meaning?). These included “words of sound and vision which express ideas outside of her experience” (Macy, pp. 134–135). We do the same, with words like ‘pregnant’ used by a male (Rapaport, 1988b, 2003a). Evidently, though, much more of Keller’s knowledge is knowledge by description than is ours. (Of the paintings in an art gallery, she said, “I have at least the satisfaction of seeing them through the eyes of my friends”—p. 200.)

What really happened at the well house?

The well-house association of ‘water’ with water was not different in kind from previous word–object associations that Keller had made and had used for

communication. Sullivan was not trying to teach Keller something new but merely trying to reinforce something she had more or less successfully taught her before. Various incidental experiences—Keller’s mug/water-or-milk confusion, her memory of the spoken word ‘wah-wah’, and the perhaps unique “co-activation” of word and object (cf. Mayes, 1991, p. 111)—no doubt contributed to making the well-house experience the significant event it was. But what exactly *did* she learn, and why was it so significant?

What did Keller learn?

Keller learned something she had not been taught. In her own and Sullivan’s words, she learned that things have “names”. But merely knowing that ‘w-a-t-e-r’ or ‘d-o-l-l’ were the appropriate finger spellings to perform when in the presence of water or a doll—or, significantly, when *not* in the presence of water or a doll, but desiring one or (more importantly; see Sect. ‘The significance of names’, below) merely wishing to convey one’s thoughts about water or a doll—could be described as knowing that those things had names.³⁷

Keller learned something more: *that some things in the world (viz., finger spellings) were names of other things in the world*. More precisely, she learned the concept of a name, thereby learning a metalinguistic fact:³⁸ Her mental world became more than an unstructured associative network of concepts; it developed a syntactic–semantic structure, by means of which some of the entities in it (her internal representations of words) “named” others (her internal representations of objects, events, ideas, etc.).

A SNePS analysis of what Keller learned: preliminaries

It may prove helpful to use SNePS to understand Keller’s accomplishment. To do this, I will first replace the *lex arcs* of Figs. 1 and 2 by a different mechanism for relating words to things: a “name-object” case frame.³⁹ The proposition that humans are mammals will now be represented as in Fig. 3, instead of as in Fig. 1.

In SNePSLOG, we can define a predicate ISA-Thing-Named:

define-frame ISA-Thing-Named (nil object name)

where

$[[\text{ISA-Thing-Named}(o, n)]]$ = o is the concept associated with (or expressed by) the lexical entry (i.e., the word, or “name”) n .⁴⁰

³⁷ Of course, ‘name’ (or ‘word’) might be overly simplistic. A simple (to us) finger-spelled “name” might be interpreted as a full sentence: possibly, ‘d-o-l-l’ means “Please give me my doll.” Cf. “Please machine give cash” as the meaning of pushing a button on a cash machine; see Sect. ‘T-naming’.

³⁸ As David Wilkins pointed out to me.

³⁹ This has been called an *expression–expressed* case frame (Neal & Shapiro, 1987; Shapiro et al., 1996) or a *word–object* case frame. I will use ‘name–object’ instead of ‘word–object’ for consistency with Keller’s and Sullivan’s terminology. The replacement of *lex arcs* with a name–object case frame is not essential to my point, but makes the exposition clearer.

⁴⁰ More precisely, o is the concept associated with (or expressed by) the lexical entry *aligned with* n . See n. 13.




Fig. 3 $[[m6]]$ = Humans are mammals, using the name-object case frame instead of a lex arc. Nodes b1 and b2 correspond to nodes m1 and m2 of Fig. 1. The nodes labeled “humans” and “mammals” represent lexical entries exactly as the corresponding nodes of Fig. 1 do. Nodes m4 and m5 represent propositions relating entities to names for them (see text)

(Cf. the semantics of the function symbol *thing-called*, in Sect. ‘The SNePS knowledge-representation, reasoning, and acting system’, above.) The network of Fig. 3 can then be constructed by asserting

ISA-thing-named(b1,humans).
 ISA-thing-named(b2,mammals).
 AKO(b1,b2).

(I.e., roughly—but only very roughly! (see below)— b1 is a thing named ‘humans’; b2 is a thing named ‘mammals’; and b1 is a kind of b2—i.e., humans is a kind of mammals, or—more idiomatically— humans are mammals.)

In Fig. 1, nodes m1 and m2 are functional terms representing what we have called “structured individuals” in previous SNePS writings. They have what philosophers might call “essential” features (in this case, their lexical expression). In Fig. 3, nodes b1 and b2 are base nodes (i.e., nodes with no arcs emanating from them), which are like Skolem constants, representing what philosophers might call “bare particulars”,⁴¹ with *no* essential features, only “accidental” features asserted about them (in this case, their names—i.e., their lexical expression).⁴²

In Fig. 1, Cassie can *use* the words ‘humans’ and ‘mammals’ but has no beliefs *about* them. In Fig. 3, she can *both* use the words *and* has beliefs about them. For example, node m4 represents a belief about something, viz., $[[b1]]$ (which is, in fact, the class of humans) and a name (or word) for it, viz., ‘humans’. However, just as the lex-arc representation is not expressible by Cassie, $[[m4]]$ is a belief that is also not expressible easily, if at all. The lesson of the discussion that follows is that a cognitive agent can have names for things without realizing or being able to say that they are names. The latter is a greater cognitive achievement than the former. This is what “understanding that everything has a name” really means.

Only m6 is easily expressible; to express it, Cassie uses the words at the heads of m4’s and m5’s name arcs. SNePS’s algorithm (Shapiro, 1982) for expressing a node (such as m6) that has a subclass–superclass case frame is to generate an expression for the node at the head of the subclass arc (b1), followed by the word ‘are’, followed by an expression for the node at the head of the superclass arc (b2). This (recursive) algorithm for expressing a node (like b1, or b2) that is at the head of an object arc

⁴¹ Sellars (1963, p. 282); Allaire (1963, 1965); Chappell (1964); Baker (1967).

⁴² Instead of using “base-node” labels b1, b2, etc., we could have used mnemonic (for *us*) English words, like humans1 and mammals1, as is often done in AI. For present purposes, as well as for McDermott-1981-like reasons, these would be needlessly confusing.

emanating from an asserted node that has a *name* arc emanating from *it* is to generate the lexeme at the head of the *name* arc that emanates from the asserted node (*m4!*, or *m5!*) that has the *object* arc to *b1* (or *b2*). Thus, *m6* would be expressed as “Humans are mammals”.

But to express *m4* or *m5*, Cassie would also have to use those words to express *b1* or *b2*—so how would she express the nodes at the heads of *m4*’s or *m5*’s *name* arcs? One (bad) way to express *m4* is to generate an expression for *b1*, as before, followed by the word ‘is’, followed by the word at the head of the *name* arc, viz., ‘humans’. But this leads to the awkward and ambiguous “Humans is humans”.

Another (bad) way is to express the base node *b1* with a demonstrative like ‘this’, followed by the word ‘is’, followed by the word ‘humans’. However, if we apply this to Fig. 4’s node *m8*, we wind up saying “this is wet”, which is ambiguous between the intended “this thing is called ‘wet’ ” and the more natural “this thing has the property of being wet”.⁴³ Actually, it would be better to express *m8* with something like “this is wetness”, but it is difficult to generate such “philosophical” nouns without further network information to signal to the NL generator when (and how) to use them.

Yet another possibility is to express Fig. 3’s *m4* by saying “something that is a subclass of mammals is humans”, i.e., to describe *b1* in terms of the *assertional* information about it in the rest of the network (“something that is a subclass of mammals”), followed by the word ‘is’, followed by the word ‘humans’: “Something that is a subclass of mammals is humans.” This is better, but still suffers from an ambiguity in the word ‘is’: the first occurrence (in “something that is a subclass of mammals”) is the “is” of predication; the second (in “is humans”) is elliptical for something like “is called” or “is named”. And *that* has the further problem that it suggests that Cassie has a belief about “names” or about “calling”, which, at this stage, she does not, for the following reason.

In SNePS, Cassie only has beliefs about entities represented by *nodes*; arcs and their labels are merely structuring devices. From a network point of view, arc labels serve to identify argument positions for an implicit (or unnamed) function symbol. Thus, for example, the *object–property* case frame of Fig. 4 is an implicit, 2-place function-symbol, the order of whose arguments is irrelevant since each argument place has a label, viz., an arc label, in a Davidsonian fashion (Davidson, 1967). SNePSLOG makes that function symbol explicit (here, we “name” it ‘Is’) and fixes the order of the arguments (so that the arc labels do not need to be visible). Just as in predicate logic, we cannot speak about *either* argument positions (arcs or their labels) *or* function symbols (whether implicit or explicit). We can only speak about terms (i.e., nodes). This is a linguistic analogue of Quine’s (1953) dictum that to be is to be the value of a bound variable; or, as Wittgenstein might have put it, whereof we cannot quantify over, thereof we must be silent. (See Sect. ‘The SNePS knowledge-representation, reasoning, and acting system’, above, and Shapiro & Rapaport, 1987 for further discussion.)

⁴³ Fig. 4 can be created, using SNePSLOG, by defining: `define-frame Is (nil object property)`, where `[[Is(o,p)]] = object o has property p`, and then asserting: `ISA-thing-named(b3,water)`. `ISA-thing-named(b4,wet)`. `Is(b3,b4)`.




Fig. 4 [[m9!] = Water is wet. Here, the proposition is analyzed as representing that an object [[b3]—i.e., something expressed by the “name” ‘water’—has the property [[b4]—i.e., something expressed by the “name” ‘wet’

The SNePS analysis: “promotion” of a structural arc to a conceptual node

We are now in a position to see Keller’s accomplishment through the eyes of SNePS. In SNePS terms, *before* the well house, Keller’s belief that water is wet could be represented as in Fig. 4. Figure 4 is very similar to Fig. 3, the major difference being that, in Fig. 3, we used the subclass–superclass case frame to represent the “is” of class subsumption (“Humans *are*—in the sense of *is a subclass of*—mammals”), whereas, in Fig. 4, we use an object–property case frame to represent the “is” of predication (“Water *is*—in the sense of *has the property of being*—wet”). In both cases, we represent the principal concepts (*human, mammal, water, wet*) by base nodes—nodes with no structural information. As noted above, logically and linguistically, these are very much like Skolem constants; ontologically, they are very much like bare particulars. We are able to talk about them by giving them names, using the name–object case frame. However, although we (and Cassie) can express b3, for instance, by *using* the “name” ‘water’, there is no explicit belief that ‘water’ *is a name for* b3.

After the well house, I suggest, Keller “promoted” the arc-label name to a node, about which she can have an explicit belief, as shown in Fig. 5. (From the SNePSLOG point of view, one *might* say that she “demoted” a predicate to a term, but really she “promoted” a (hidden) argument position to a term.) Here, we use the object–rel–possessor case frame to represent general possessive expressions of the form “*x is y’s z*”, as follows: object *x* stands in the *z* relation to possessor *y* (see Chun, 1987 for details).⁴⁴ In SNePSLOG, we could define a predicate:

define-frame Possession (nil object possessor rel)
 where

[[Possession(*o,p,r*)]] =object *o* is possessor *p*’s *r*, or: *p*’s *r* is *o*, or: object *o* stands in the *r* relation to possessor *p*. Then Fig. 5 is the result of asserting

⁴⁴ Briefly, this case frame is the SNePS analogue of several English possessive constructions (*’s, of,* etc.), neutral as to whether the possession is that of whole-to-part, ownership, kinship, etc. E.g., “Frank is Bill’s father” might be represented as: Possession(Frank, Bill, father), understood as expressing the relationship that Frank *is the father of* Bill. If, in general, Possession(*o,p,r*), it might follow that *o* is an *r*—e.g., Frank is a father—but only because, in this case, $\exists p$ [Frank is the father of *p*]. As Stuart C. Shapiro (personal communication, August 2006) observes, if one man’s meat is another’s poison, it doesn’t follow that the first person’s meat is meat *simpliciter*, but only meat *for him*, since it is also poison (and thus *not* meat) *for* someone else.




Fig. 5 A network containing the network of Fig. 4 as a subnetwork. Here, $[[m11]] = \text{Water's name is 'water'}$; i.e., $[[b3]]$'s $[[b5]]$ is 'water'. $[[m12]] = \text{Wet's name is 'wet'}$; i.e., $[[b4]]$'s $[[b5]]$ is 'wet'. $[[m10]] = \text{the inexpressible proposition asserting the name of } [[b5]]$ —i.e., the concept of a name—to be the lexical entry 'name'

ISA-thing-named(b5,name).
 Possession(water,b3,b5).
 Possession(wet,b4,b5).

Thus, e.g., [[m11]] = ‘Water’ is [[b3]]’s name (i.e., the object ‘water’ stands in the “name” relation to possessor b3). Cassie (or Keller) can now *say* that ‘water’ is a name for water, because she has an explicit concept of a name, represented by a node, rather than an inexpressible arc label.

Note that the now-explicit concept of a name ([[b5]]) is itself expressed by the name ‘name’. At the risk of making the network of Fig. 5 illegible, we should also show a node m13 that represents the proposition that [[b5]]’s name is ‘name’. This is shown separately, for the sake of clarity, in Fig. 6.⁴⁵

Roger Brown (1973, p. 3) observed that “linguistic processes, in general tend to be invisible. The mind’s eye seeks the meaning and notices the medium as little as the physical eye notices its own aqueous humor”. The well-house experience made one crucial linguistic process visible to Keller. Keller learned more than how to *use* words or signs; she learned that certain things *were* words or signs. She learned more than how to “see through” a sign to its meaning; she learned how to see the sign *as a sign*.

The significance of names

Terrace’s theory of names

But why is it so significant to learn what a name is? A possible answer consistent with Keller’s post-well-house behavior, can be found in Herbert S. Terrace’s essay, “In the Beginning Was the ‘Name’ ” (1985). In this section, we’ll look at Terrace’s theory of why “naming” is important, whether his notion of “naming” is akin to Keller’s, the extent to which his theory is supported by more recent observations, and the relevance of all this to computational NL competence. (To distinguish Terrace’s terms from, say, Keller’s, I will refer to Terrace’s theory of names and naming as the theory of ‘T-names’ and ‘T-naming’.)

Overview of T-naming

Terrace summarized his theory in a letter to the editor of the *New York Review of Books*

Before speculating about the origins of grammar, it is prudent to ponder the origins of the *referential* use of individual words. Unlike apes, children use individual words to comment about objects **for the sheer joy of communicating**. Adults do not reward a child with a tree when she points to one and then says *tree* By contrast, there is no evidence that apes communicate *about* things. As Lord Zuckermann observes, apes use language not as “... a way of

⁴⁵ In SNePSLOG, node m13 could be constructed by asserting Possession(name,b5,b5). i.e., [[b5]]’s name is ‘name’. This is the base case of a recursion when rule (1) of Sect. ‘SNePS analysis of learning to name’, below, is applied to m10. (The first consequent of that rule would not build a new node (because of the Uniqueness Principle; Shapiro, 1986); rather, it would return the already-existing node m10. The second consequent builds m13.




Fig. 6 Node m10 is the hard-to-express node that represents that Cassie can *use* the name ‘name’ to talk about [[b5]]. Node m13 represents the proposition that allows Cassie to talk about the naming relationship between [[b5]] and ‘name’, viz.: [[m13]] = ‘name’ is [[b5]]’s name. I.e., the node labeled name represents [[b5]]’s [[b5]]; i.e., it represents [[b5]]’s name; i.e., it represents the name ‘name’

conversing, but a game associated with pleasurable reward.” Although the origins of human language are unclear, one contributing factor must be the adaptive value of communicating meanings that cannot be expressed in a single word (e.g., *the large tree* or *the single-tusked elephant at the large tree*). It appears, therefore, that the cognitive leap to language occurred in two stages: first, developing the lexical competence to use arbitrary symbols to refer to particular objects and events, and then the syntactic competence to combine and inflect those symbols systematically so as **to create new meanings**. (Terrace, 1991, p. 53; italics in original, boldface mine.)

The game-like nature of language use by apes is reminiscent of Keller’s *pre*-well-house use of language, and the child’s use of words for objects “for the sheer joy of communicating” clearly describes Keller’s *post*-well-house behavior. So, *prima facie*, T-naming might well be what Keller learned to do at the well house. It was there that, by her metalinguistic discovery, she “develop[ed] the lexical competence to use arbitrary symbols to refer to particular objects and events”.

Antonio Damasio (1989, p. 25) has observed that one of the “stages” of concept formation “is that of generating names that are pertinent to the stimulus and are usable to narrate the primary display when inserted in appropriate grammatical structures.” Perceiving an object causes neuronal activity representing its features and structure. These, in turn, are linked to other neuronal structures that “generate names”, which allow me to communicate to you that I am thinking of an object and what it is.

Bruner (1983, p. 114) makes a similar observation

In object request the principal task is to incorporate reference into request. When the child finally masters nominals, he [sic] need no longer depend upon the interpretive prowess of his mother or the deictic power of his indexical signaling. The demands of dealing with displaced reference in requesting objects provide an incentive.

Having (T-)names gives one power. Keller, apparently, lacked this ability before the well house.

One doesn’t, however, *create* new meanings (*pace* Terrace’s quote, above). Rather, one creates new (combinations of) symbols that are able to be associated with meanings that one had no words for before (cf. Elgin, 1984 for a literary exploration of this theme). Zuckermann’s way of putting this is better (1991, p. 53):

“the additional adaptive value of joining lexical items in ways that multiplied the meanings that they can convey”.

So, for Terrace, syntax is built on top of lexical semantics, as it seems to have been for Keller, too. Bruner, after observing two children’s acquisition of language, concurs (1983, pp. 97–98)

[R]equesting absent objects ... requires a degree of specification not needed when an object is within reach or sight. An object out of sight requires the use of nominals for easy specification. ...

Remote or displaced requests began at the landmark age of fourteen months in both children.

T-naming enables conversation—the exchange of *information*, distant or displaced in space and/or time, with no immediate, practical goals (other than, perhaps, to satisfy one’s curiosity or to be sociable).⁴⁶

T-naming

Terrace (1985, p. 1011) considers “The ability to refer with names” to be “perhaps” the most “fundamental” and “uniquely human skill”. This referential ability appears to be akin to symbol grounding. It is the link between word and world (reminiscent of “alignment” in SNePS; Shapiro & Ismail, 2003). But reference, as Frege taught us, is not all there is to meaning: is Terrace’s notion of “referring” *Bedeutung*? Is it *Sinn*? What would he say about the “referring” use of a name like ‘unicorn’ or ‘Santa Claus’? (Note that it is only under certain circumstances that one would use such “non-referring” names to request the named entity.)

SNePS analysis of learning to name

How did Helen Keller learn to name? She learned the nature of the relationship between a name (i.e., a finger spelling) and an object. She learned the name of naming. So, is it possible for Cassie to learn to name? Given the network of Fig. 4, we’d like her to be able to say, when asked, something like “‘water’ is the *name* of b3”, or (more likely) “‘water’ is the name of water”, or (even more likely) something along the lines of “‘water’ is the name of this stuff” (since [[b3]] is not otherwise characterized; cf. the earlier discussion of this in Sect. ‘A SNePS analysis of what Keller learned: preliminaries’). In any case, the point is for her to do something more than merely express b3 *using* ‘water’. But—as we have seen—that would require a network like that of Fig. 5.

Similarly, consider the network in Fig. 7. This is one way to represent that something is a red, round ball (represented by nodes m19, m17, and m15, respectively).⁴⁷ In SNePSLOG, Fig. 7 is produced by first defining a classic AI predicate for class membership:

⁴⁶ A relevant (humorous) take on the nature of conversation appeared in a *New Yorker* cartoon showing a man, two women, and a chimp, all dressed in suits and ties, sitting in a bar at a table with drinks; the chimp thinks, “Conversation—what a concept!”.

⁴⁷ Ann Deakin pointed out to me that color is not a good example for Helen Keller! Perhaps taste or smell would be better? On the other hand, for *Cassie*, color might be more accessible than taste or smell (cf. Lammens, 1994)!

define-frame ISA (nil member class)

and then asserting:⁴⁸

```

;;; Something is a thing named "ball"
ISA-thing-named(b6,ball).
;;; Something is a ball
ISA(b7,b6).
;;; Something is a thing named "round"
ISA-thing-named(b8,round).
;;; The ball is round
Is(b7,b8).
;;; Something is a thing named "red"
ISA-thing-named(b9,red).
;;; The round ball is red
Is(b7,b9).

```

In Fig. 4, Cassie can *use* ‘water’ as the name of [[b3]] (i.e., she can call [[b3]] ‘water’). But without a node explicitly asserting (i.e., naming) the relationship between the node labeled “water” and b3, she does not understand—*de dicto*—that ‘water’ is [[b3]]’s *name*. She has no theory of names or naming. Of course, we would like Cassie to be able to say that red’s name is ‘red’, etc. This would require additions to the network of Fig. 7 as shown in Fig. 8, produced, using SNePSLOG, by asserting

```

Possession(ball,b6,b5).
Possession(round,b8,b5).
Possession(red,b9,b5).

```

Similarly, we might want Cassie to be able to say, when asked, not merely that the round ball is red, but explicitly that red is a *property* of the round ball. So, we would need to be able to have Cassie answer this sort of question: given a propositional node asserting that some object *a* has some property *F*, what is the (name of the) relationship between *a* and *F*? As with the case of names, without a node explicitly asserting (i.e., naming) the relationship between b9 and b7 in Fig. 7, she does not understand—*de dicto*—that b9 (i.e., red) is a *property* of b7. She has no theory of properties, either. (Most people other than cognitive scientists probably don’t!)

Although these two situations are analogous, there is, perhaps, a slight advantage to the naming case over the property case (in addition to the fact that most people who are not cognitive scientists *do* have a theory of naming!). For in order for Cassie to *utter* something about [[b3]] or [[b4]] (in Fig. 4), she must use a word for it, whereas she would not need to use the word ‘property’ in uttering m17 or m19 (in Fig. 7). (This may explain why theories of naming are probably more common than theories of properties. ‘Property’ (in this sense) is a technical term of philosophy, not a term of ordinary language.)

So Cassie could, perhaps, recognize that there is a relationship between her concept of an object and the word she uses when she says something about it (e.g., in Fig. 4, between b3 and the node labeled *water*), though this is a task for empirical investigation. From this recognition, she could come to believe a proposition that would link these nodes with a node for the relationship, which, if she asked us, we

⁴⁸ Lines beginning with semicolons are comments.




Fig. 7 $[[m15]] = \text{Something (viz., } [[b7]] \text{) is a member of the class } [[b6]]$, which is expressed by the lexical entry 'ball'. $[[m17]] = \text{It is round (i.e., it has the property expressed by 'round')}$. $[[m19]] = \text{It is red (i.e., it has the property expressed by 'red')}$. N.B.: these nodes are *added* to the networks of Figs.5 and 6




Fig. 8 Nodes added to the networks of Figs. 5–7 after Cassie learns the concept of “names for things”. (As in Figs. 5 and 6, $[[b5]]$ = the concept of “a name”) $[[m20]]$ = $[[b6]]$'s name is ‘ball’, $[[m21]]$ = $[[b8]]$'s name is ‘round’, $[[m22]]$ = $[[b9]]$'s name is ‘red’

could tell her was called ‘naming’. So she could assert node m11 in Fig. 5. In any event, let us suppose that this can be done. It is, it would seem, what Keller did.

In SNePS, this could be done via rules such as the following:

1. If x is y (in the sense that x is named, or called, y), then y is x 's name;⁴⁹

more precisely:

For any (two nodes) x and y , if object x has name y , then object y is possessor x 's name,

$$\text{all}(x,y)(\text{ISA-thing-named}(x,y) \Rightarrow \{\text{ISA-thing-named}(b5,\text{name}),\text{Possession}(y,x,b5)\})$$

or, in SNePSLOG

2. If x is y (in the sense that x has property y), then y is a property;⁵⁰

more precisely:

For any (two nodes) x and y , if object x has property y , then y is a member of the class of things called ‘property’;

or, in SNePSLOG

$$\text{all}(x,y)(\text{Is}(x,y) \Rightarrow \{\text{ISA-thing-named}(b10,\text{property}), \text{ISA}(y,b10)\}).$$

To get Cassie to *learn* such rules is a worthy topic for future research.

Terrace's and Keller's theories of naming

Some of Terrace's claims about the linguistic abilities of apes are reminiscent of Keller's pre-well-house linguistic abilities:

... even though apes can learn substantial vocabularies of arbitrary symbols, there is no a priori reason to regard such accomplishments as evidence of human linguistic competence. After all, dogs, rats, horses, and other animals can learn to produce arbitrary “words” to obtain specific rewards. (Terrace, 1985, p. 1012.)

Keller “learn[ed] substantial vocabularies of arbitrary symbols”, too. (Cf. Sullivan's use of expressions like “Keller knew n words”.) But what kind of learning is this? Given the context of Terrace's paper, it does not seem to mean that apes (or Keller) could link the arbitrary symbols to objects. And given Terrace's belief in the logical and chronological priority of naming over syntax, it does not seem to mean that the apes (or Keller) knew the syntactic roles of the symbols.

But Keller could “produce arbitrary ‘words’ to obtain specific rewards”. So, by ‘learning a symbol’, Terrace must mean producing the symbol in return for a (not necessarily associated) reward, without any (semantic) linking of the symbol with the world.

⁴⁹ Note that, arguably, it does not necessarily also follow that y is a name *simpliciter*; see n. 44. A similar rule appears in Shapiro (2003); cf. Shapiro and Ismail (2003).

⁵⁰ Here, it arguably *does* make sense to say that y is a property *simpliciter*, not merely that it is x 's property; see n. 49.

It would be just as erroneous to interpret the rote sequence of pecks [by a pigeon], red → green → blue → yellow, as a sentence meaning, *Please machine give grain*, as it would be to interpret the arbitrary sequence of digits that a person produces while operating a cash machine as a sentence meaning *Please machine give cash*. In sum, a rote sequence, however that sequence might be trained, is not necessarily a sentence. (Terrace, 1985, p. 1014.)

This sounds like Keller's pre-well-house use of language. But *why* aren't those rote sequences sentences with those meanings? Granted, perhaps, they lack the exact grammatical *syntax* of those *sentences*, but why say that they lack the same meaning? Propositions (meanings) can be *implemented* (i.e., expressed) in numerous different ways, even within the same language; after all, paraphrases exist. When I push a sequence of buttons on a cash machine, why am I *not* telling (asking) it to give me a certain amount of money? Isn't that what the I/O encoding scheme amounts to? Isn't there at least an aspect (or a kind) of meaning that is determined by how an expression is used or by the context that it appears in (cf. Kalderon, 2001; Rapaport & Ehrlich, 2000)? Granted, perhaps what the symbols mean *to me* is not what the symbols mean (if anything) *to the machine*, but as long as we can communicate (so as to overcome misunderstandings), what's the difference? And, in any case, it *does* mean *something* to the machine—it has syntactic meaning (internal meaning).

A brief example of a symbol having two such different meanings (a meaning for the “speaker” or user, and another for the “hearer” or “receiver”) might be instructive. In my university library, when I push the button for the elevator on the ground floor, the button lights up. I have come to learn empirically that if the light stays on when I release the button, it means that the elevator is not on the ground floor. If the light immediately goes off, it means that the elevator *is* on the ground floor and that in a few seconds the door will open. The light's going off is a symbol that *I* interpret to mean “the elevator is on the ground floor; the door will open shortly”. This is its meaning *for me*: *my* interpretation of it. It is important to note that I have determined this meaning empirically (in particular, contextually; cf. Rapaport & Ehrlich, 2000), and I could be wrong. If the light goes off and no elevator door opens within a few seconds (and it is in working order), I would have to revise my beliefs.

There is also its meaning *for the elevator system*: the role that the light going off plays in the electrical network that controls the elevator. A study of the wiring diagram might reveal, let us suppose, that when the button is pushed, a circuit is closed that lights the button and a test is conducted to determine the location of the elevator. If the elevator is on the ground floor, the circuit is opened, turning off the light, and, a short time later, another circuit is closed, opening the door; else, the circuit remains closed, keeping the light on, and another circuit is closed, sending the elevator to the ground floor, where the light-circuit is opened, turning off the light, followed by the door opening as before. The meaning of the light's going off, then—its role in that network—is correlated with the elevator's being on the ground floor. From the elevator's point of view, so to speak, the only way the light going off would *not* mean that the elevator is on the ground floor would be if the elevator were broken.

One thing missing from such behavioral uses of language is the *intention* to communicate an idea by using a certain word (Terrace, 1985, p. 1017). Non-human animals who have been trained, behavioristically, to “use language” (which I place

in scare quotes so as not to beg any questions about what it is they are actually doing) *seem* to communicate intentionally with each other. But, Terrace points out,

That would presuppose not only that Jill [one of the pigeons mentioned earlier] could discriminate each color from the others (when she clearly could) but that Jill also understood that (a) some arbitrary communicable symbol described color_{*i*}, (b) she sought to communicate to Jack [the other pigeon] that the color she saw was color_{*i*}, and (c) Jack would interpret Jill's message as Jill intended it. There is no evidence to support any of these suppositions. (Terrace, 1985, p. 16.)

This, of course, does not affect experiments in computational linguistics that provide mechanisms (based on speech-act theory) for implementing intentions to speak.⁵¹ One of the advantages a computer has over non-human animals is that, even if the latter lack intentions to communicate, the former can be given them!

But, according to Terrace, even if it could somehow be shown that a non-human animal *intended* to use a certain word to communicate the idea that it wanted a specific object, that would not suffice to show that it was using the word as a *name* for the object. This is because the animal might believe that using that word is the best way to cause the listener to *give it* the desired object. Roughly, the animal might be ascribed the belief, “If I make such and such a noise [or use such and such a sign], my trainer will bring me one of those sweet, red, round things to eat, so I'll make that noise [or use that sign]”, without the animal having a belief that the sign “names” the object.

So, *T-naming* is a very special activity

... the main function of such words [viz., the “use of a symbol as a name”] in the use of human language—[is] the transmission of information from one individual to another *for its own sake*.

... a child will utter a name of an object, person, color, and so on, simply to indicate that she or he knows that the object she or he is attending to has a name and also to communicate the fact that she or he has noticed that object

... In many instances, the child refers to the object in question spontaneously and shows no interest in obtaining it. The child not only appears to enjoy sharing information with his or her parent but also appears to derive intrinsic pleasure from the sheer act of naming. (Terrace, 1985, pp. 1016–1017.)

A mere *associative* link between an arbitrary symbol and a specific object⁵² is not sufficient for a *semantic* link, according to Terrace. What is also needed is *intending* to use the symbol for the object *only to indicate that you are thinking of the object*.

So, one difference between Keller's pre- and post-well-house language might be this: *Before*, she didn't have such intentions; *after*, she did. I.e., *after*, she *intended* to refer to water by ‘water’. To do this, she needed to be able to think and talk about the naming relationship. But is it really the case that she lacked that intention

⁵¹ Allen and Perrault (1980); Bruce (1975); Cohen and Levesque (1985, 1990); Cohen and Perrault (1979); Grosz and Sidner (1986); Haller (1994, 1995).

⁵² Perhaps like that between the node labeled humans and node m1 in Fig. 1 or between the node labeled humans and node b1 in Fig. 3. (The latter associative link is represented by node m4.)

before? Although the evidence is at best unclear, I think she did have the intention, but not the name of the naming relationship, so that her intentions were often frustrated.

Searle might say that Searle-in-the-room (or the CR itself, on the systems reply) *says* things but doesn't *mean* them. What, then, does it mean to mean something by a word? In Cassie's terms, it would be this: Cassie has a concept that she wants to communicate to Oscar.⁵³ She has a name for the concept. So, she utters that name, *assuming* that Oscar uses that word for the "same" (i.e. the corresponding) concept (in his mind), and *intending* that Oscar will think of that concept—that that concept will be "activated" in Oscar's mind. As Terrace puts it,

In most human discourse, a speaker who utters a name expects the listener to interpret the speaker's utterance as a reference to a jointly perceived (or imagined) object (Terrace, 1985, p. 1017.)

"Jointly imagined" is where such entities as unicorns and the Hob-Nob witch⁵⁴ come in. So, T-naming is more appropriately a relationship between a name and a mental concept. That is, in answer to the question raised at the beginning of this section, it is more like *Sinn* than *Bedeutung* (actually, it is more Meinongian—cf. Rapaport, 1981). So, in the CRA, what's missing from Searle's description is the *intention* to communicate.⁵⁵ So one *could* argue that if Searle's CR is to be taken literally, then it *doesn't* understand, but that would only be because Searle hasn't fleshed out the full theory of computational NL understanding and generation. And if it's to be taken as schematic for whatever would be the full theory, then he's wrong.

As both Terrace and Bruner (1983; cf. Rapaport, 2003a) point out—and Anne Sullivan long before them—

... language draws upon certain kinds of nonlinguistic knowledge. For example, before learning to speak, an infant acquires a repertoire of instrumental behavior that allows her or him to manipulate and/or approach various objects. An infant also learns how to engage in various kinds of social interaction with her or his parents—for example, being able to look where the parent is looking or pointing. Eventually, the child learns to point to things that he or she would like the parent to notice. In short, the infant first masters a social and conceptual world onto which she or he can later map various kinds of linguistic expression. (Terrace, 1985, p. 1018.)

⁵³ Oscar is the Other SNePS Cognitive Agent Representation, first introduced in Rapaport, Shapiro, and Wiebe (1986).

⁵⁴ "Hob thinks a witch has blighted Bob's mare, and Nob wonders whether she (the same witch) killed Cob's sow" (Geach, 1967, p. 628).

⁵⁵ I.e., the intention to communicate should be one of the features of computational NL understanding and generation in addition to those I cited in Rapaport (2000, Sect. 8). There, I said that a computational cognitive agent must be able to "take discourse (not just individual sentences) as input; understand all input, grammatical or not; perform inference and revise beliefs; make plans (including planning speech acts for NL generation, planning for asking and answering questions, and planning to initiate conversations); understand plans (including the speech-act plans of interlocutors); construct a "user model" of its interlocutor; learn (about the world and about language); have lots of knowledge (background knowledge; world knowledge; commonsense knowledge; and practical, "how-to", knowledge ... and remember what it heard before, what it learns, what it infers, and what beliefs it revised And it must have effector organs to be able to generate language. In short, it must have a mind."

So, *internal concepts* are learned first, via links between visual input and mental concepts.⁵⁶ Names are attached later.

This suggests that the first rule in Sect. ‘SNePS Analysis of Learning to Name’ need not incorporate a name right away; i.e., Cassie might have an unnamed concept. (Again, cf. Elgin, 1984.) Thus, the rule might be:

If x is (named) y , then y is x ’s *something-or-other*;
more precisely:

For all (nodes) x and y , if object x ’s name is y , then object y is possessor x ’s *something-or-other*,

or, in SNePSLOG:

$\text{all}(x,y)(\text{ISA-thing-named}(x,y)\Rightarrow \text{Possession}(y,x,\text{something-or-other}))$

The “base” node (roughly, Skolem constant) labeled *something-or-other* might later be “given” a name by asserting a proposition to the effect that whatever object is represented by that node is expressed by the word ‘name’; in SNePSLOG:

$\text{ISA-Thing-Named}(\text{something-or-other},\text{name}).$

(Here, ‘name’ appears embedded in the predicate, about which predicate we cannot speak; it is hidden as the label for an argument position in the definition of the predicate, about which position we also cannot speak; and it appears explicitly as a term, which *can* be spoken about.) Alternatively, we could have a “defeasible” rule (with three conjoined antecedents) to the following effect:

$\text{all}(x,y,z,r)(\{\text{Is}(z,\text{unknown}),\text{Possession}(x,y,z),\text{Possession}(x,y,r)\}\&\Rightarrow \text{Equivalent}(\{z,r\})).$

where⁵⁷

$\text{define-frame Equivalent}(\text{nil equiv})$

and

$[[\text{Is}(o, \text{unknown})]] = \text{object } o \text{ has the property of being } [[\text{unknown}]]$
 $[[\text{Equivalent}(\{x_1, \dots, x_n\})]] = \text{mental entities } x_1, \dots, x_n \text{ correspond to the same actual object.}$

I.e., if z is an unknown (or unfamiliar) concept, and if x is y ’s z , and if x is (also) y ’s r (where r is some antecedently *familiar* concept), then (defeasibly, of course!) z is (or: might be) equivalent to r . (The “defeasibility” is automatic in SNePS: virtually all beliefs are revisable in the light of later beliefs, so there is very little “cost” to this somewhat overgenerous rule; in any case, the antecedent about being “unknown” limits its application somewhat.) So, if

⁵⁶ SNePS pic arcs, like lex arcs, point to SNePS nodes representing pictorial images (Rapaport, 1988b; Srihari, 1991a, b, 1993, 1994; Srihari & Rapaport, 1989, 1990). Also cf. anchoring or “alignment”; Shapiro and Ismail (2003).

⁵⁷ Note that the predicate “Equivalent” is defined in terms of a single arc (“equiv”) that can point to a set of nodes; this has the same effect as having a set of arcs with the same label, each of which points to a node. See Maida and Shapiro (1982); Shapiro and Rapaport (1987); Rapaport, Shapiro, and Wiebe (1997) for further discussion of the SNePS notion of equivalence.

something-or-other is an unknown concept, and if x is y 's something-or-other, and if x is also y 's name, then (perhaps) a something-or-other is a name. This is, at least, a plausible hypothesis, subject to further modification (Rapaport, 2005a; cf. Rapaport & Ehrlich, 2000).

Keller was able to learn language once she grasped the concept of naming—"the conventions of using symbols and words that do the work of referring" (Terrace, 1985, p. 1021). Apes, according to Terrace, lack this ability (pp. 1021, 1023–1024): "an ape does not naturally refer to an object to which it attends *solely for the purpose of noting that object to a conspecific*. ... [W]hatever referential skills an ape displays naturally seem to be in the service of some concrete end" (p. 1024, my italics). Keller's post-well-house interest in the names of objects seems to have been for its own sake. Her own pre-well-house signs were always "in the service of some concrete end".

Can apes speak for themselves?

Is Terrace's emphasis on intentional but non-purposive naming a reasonable one? Why does Terrace think that using a sign in order to get a reward is *not* using it linguistically (Terrace, 1985, pp. 1016–1017)? What is so important about T-naming? And how do we know that apes *don't* have intentions when they "use language"? Finally, is there any evidence that apes *do* T-name?

Purposive naming

Terrace has two reasons for believing that purposive-only use of names is *not* language. First, he claims that it is simply a matter of fact that apes *don't* use signs except when they want something, whereas human children at 18 months *do* (Terrace, 1985, p. 1023). So, at least, what apes do is *preliminary* to full human-like language use.

Second—and this is what's important about T-naming, according to Terrace (1985, p. 1011)—without it, grammatical syntax would not develop

... when there is a desire simply to communicate information about a relationship between one object or action and another, about some attribute of an object, or about past or future events ... ungrammatical strings of words would not suffice—hence the functional value of syntax. (Terrace, 1985, p. 1026.)

His argument seems to be this: if I want something that is present in our common environment, I can use a name or make some (perhaps non-semantic) sign to get you to give it to me. If a single name is insufficient (say, because there are two things, and I want a specific one, e.g., the *large* banana), I can combine two (or more) names. But it doesn't really matter in what order I combine them (since we can both see the items), so grammar (beyond mere juxtaposition) is not necessary. If I don't know a name for the object, I can point to it. But for *absent* (displaced) objects, pointing won't work. And if I wanted to communicate about some *feature* of an (absent) object, then grammar facilitates my communication: if I wanted to talk about the color of the banana, it becomes crucial *which* of the previously juxtaposed signs is the color-term and which the noun; grammar enters upon the scene. Note that it is highly unlikely that I *want* the color; rather, I just want you to know that I'm *thinking* about the color—that's T-naming.

Intentions

We *don't* know that apes don't have intentions. I'd bet on evolutionary grounds that they do. We *can*, however, insure that a *computer's* use of language *would be* intentional, by having the NL understanding and generating program include speech-act or intention-and-action modules. Searle might object that that's just more syntax, not "real" intentions or desires. But what would be the difference? What is a "real" intention? Moreover, desires (and intentions) *can* be adequately incorporated. In our contextual–vocabulary–acquisition project, there are times when the system, in order to settle on a definition, needs more information. That need—and it is a real need—could prompt it to seek that information, to ask questions. Surely, these would be real desires for information or intentions to ask questions.⁵⁸

One might reply that such computational desires or intentions have no qualitative "feel" to them. Perhaps. Qualia may best be seen as a feature of the implementing medium (Rapaport, 2005b). So, of course, the computational desires and intentions *might* have a "feel", depending on their implementation. Or they might not. But why would a lack of "feeling" disqualify them as "real" desires or intentions?

The thoughts, desires, and intentions of a language user that is not an ordinary human—an ape or a computer (or even, perhaps, a Helen Keller)—might be very different in kind from those of a normal human (as the cartoon described in note 46 suggests). They might very well depend on the language user's body and purposes. But they would be thoughts (or desires or intentions) nonetheless.

From representation to language

Is Terrace right about the inability of non-human primates to T-name? Several papers published *after* Terrace's deal with the issue specifically.

Terrace claims that apes lack something that humans have that enables humans to have language. Vauclair (1990) is sympathetic to this, though it's not clear that that something is T-naming.

According to Vauclair (p. 312), both human and non-human primates have "basically similar ways of coding environmental stimuli in terms of cognitive organization"; i.e., they have "mental representations". But they do *not* share language. For Vauclair, this is partly by definition

Representation is an individual phenomenon by which an organism structures its knowledge with regard to its environment. This knowledge can take two basic forms: either reference to internal substitutes (e.g., indexes or images) or use of external substitutes (e.g., symbols, signals, or words).

Communication is a social phenomenon of exchanges between two or more **conspecifics** who use a code of specific signals usually serving to meet common adaptive challenges (reproduction, feeding, protection) and promote cohesiveness of the group. (Vauclair, 1990, p. 312; italics in original, my boldface.)

⁵⁸ In Shapiro (1981), SNePS asks questions as a result of back-chaining.

Since apes and humans are not conspecifics, they cannot, *by definition*, communicate with each other. Even if that restriction is lifted, it is not clear whether T-naming is Vauclairian communication unless, perhaps, as a by-product, it “promote[s] cohesiveness of the group”—perhaps that’s the function of conversation (as the cartoon described in note 46 suggests).

Language is conceived as a system that is both communicational and representational: it is grounded in social conversation that attributes to certain substitutes (called *signifiers*) the power to designate other substitutes (called *referents*). (Vauclair, 1990, p. 313.)

So, apes and humans could never have a common language, because language is communicational (and, by definition, apes can’t communicate with humans). But how does this definition of language make it *human*-specific? Perhaps it is the social-communication aspect or negotiation (Rapaport, 2003a). After all, apes and humans don’t share a common “society”.

The closest Vauclair gets to supporting Terrace’s theory is in claiming that two of the marks of language are its ability to deal with displacement *in* space and time (i.e., things not in the presence of the speaker) and its ability to deal with what might be called displacement *from* space and time (i.e., dealing with non-existents) (Von Glasersfeld, 1977, cited in Vauclair, 1990, p. 316). T-naming, however, is logically *independent* of this. For one could, in principle, be able to refer to something displaced in (or from) space or time either if one wanted it or if one merely wanted to talk about it “for its own sake”. And, clearly, one could be able to refer to something in one’s current environment for either of those reasons.

So several issues are still open: *Do* non-human primates T-name? (Terrace, of course, says ‘no’.) Do they use language to talk about displaced objects? One would expect Vauclair to delve into this. Instead, he locates the gap between ape and human language elsewhere:

I am convinced that apes display the most sophisticated form of representation in the animal kingdom ..., but this phenomenon is insufficient in itself to qualify for linguistic status. To go beyond the 1–1 correspondence between the sign and the actual perceptual situation, we need to introduce a third term. The relation between symbol and object is more than the simple correspondence between the two. Because the symbol is tied to a conception, we have a triangular connection among objects, symbols, and concepts: “it is the conceptions, not the things, that symbols directly mean” (Langer, quoted in von Glasersfeld, 1977). (Vauclair, 1990, p. 320.)

Now, I am happy to agree with Langer; her claim is consistent with syntactic semantics. But it’s not clear what that has to do with Vauclair’s point. He *seems* here to be saying that what’s missing is the concept: no concept, no language. Yet earlier he claimed that representation *required* concepts: although ‘concept’ is not part of his *definition* of ‘representation’, he talks about “internal processing”, “internal representation”, “cognitive maps”, “internal coding”, and “internal substitutes” (pp. 313ff). What are these if not concepts?

Later (p. 321), he locates the gap “in the emergence in humans of verbal language”, but he is silent on what these emergent features are; perhaps it is T-naming. Or perhaps it is being intentional:

The specificity of human language is above all of functional order. First, this system uses representative stimuli that allow the sender to know the status of the sent message, to control it, and to endow it with intentions. (Vauclair, 1990, p. 321.)

As we have seen, this won't distinguish between human and *computer* use of language. Perhaps, however, this *was* something Helen Keller lacked before the well house.

The other thing that non-human primates lack is the social convention, the Brunerian negotiation of meaning (p. 322). This, however, seems irrelevant to T-naming. In any event, Helen Keller, arguably, had this *before* the well house, and computers certainly can have it (witness, e.g., the vocabulary-acquisition process described in Rapaport & Ehrlich, 2000; also see Rapaport, 2003a).

Orangutan reference

Can non-human primates T-name? Miles's (1990) work with the orangutan Chantek is suggestive of T-naming in such a primate.

Chantek was clearly capable of "displaced reference" (pp. 520–523), and four out of about 97 cited uses of *names* do not *appear* to involve wanting the object: making the signs (1) 'car' "as he passed [his] caregiver's car on a walk", (2) 'time' "when [his] caregiver looked at her pocket watch", (3) 'Coke drink' "*after* finishing his Coke" (my italics), and (4) 'time drink' "when [his] caregiver looked at her watch" (pp. 520–523). Each of these, however, could be interpreted otherwise: (1) Perhaps Chantek was tired of walking and wanted to ride in the car; (2) perhaps he wanted to know the time (though it's hard to believe that he had the appropriate concepts for understanding time) or perhaps 'time' was also his sign for the watch itself (we are not told); (3) perhaps he wanted another Coke to drink; (4) perhaps he was thirsty. It is hard to know when a naming act is a *T-naming*. So T-naming may be an overly restrictive criterion.

On the other hand, those who are more sympathetic than Terrace to the view that apes can use language tend to have criteria that are overly permissive. Consider Miles's three "elements" of "linguistic representation" (p. 524)

1. A sign must designate an element of the real world.
2. A shared cultural understanding about its meaning must exist.
3. The sign must be used intentionally to convey meaning.

The first element is surely too strong, since we can talk about nonexistents. Moreover, it would seem better to say that a sign must *be used by someone* to designate something (where 'something' is construed along Meinongian lines).

The second element seems to rule out interspecies linguistic representation and, perhaps, computer language. What might happen if the knowledge-representation language (i.e., language of thought) of a computer system that can learn concepts differs significantly from that of humans? (See Rapaport, 2003a, Sect. 10.) According to Winston (1975/1985, p. 143) (and Wittgenstein (1958, p. 223)—"If a lion could talk, we could not understand him."), what would happen is that the two systems—computer and human (or lion and human)—would not be able to understand each other. However, in the various ape experiments, both subject and experimenter use an artificial language, so they do have a shared cultural understanding, where the "culture" is that of the laboratory. Granted, the sign for Coke may have all sorts of connotations for the human but not the chimp. But that's no different from the fact that 'Coke' has all

sorts of connotations for *me* but not *you* because of our different experiences with it. The case of the computer is a bit easier, since *we* get to *give it* its cultural knowledge. Hence, insofar as the computer has a “mind” (i.e., a knowledge base; cf. Rapaport, 1995, Sect. 1.1.3), both it and we can have “shared cultural understanding”.

Keller’s pre-well-house uses of finger-spelled words seem in some cases to have designated in the sense of element (1) (e.g., some of her uses of ‘cake’ and ‘doll’). Even her confused use of ‘mug’ and ‘milk’/‘water’ might be taken to have designated the mug + liquid complex. Clearly, before the well house, she could designate via her own signs. Arguably, her inability to clearly designate with finger spellings could be attributed to insufficiencies in her *shared* cultural understanding. She clearly shared in some cultural understanding—after all, she was a human, living with other humans. But, of course, she was blind and deaf, hence cut off from much that the rest of us share without even realizing it. Finally, though she used her own signs intentionally to convey meaning, most of her pre-well-house use of finger spellings was no doubt mere mimicry.

Again, Miles’s criteria for referential use of words or signs is weaker than Terrace’s:

first, that signs can be used to indicate an object in the environment; second, that signs are not totally context dependent; third, that signs have relevant semantic domains or realms of meaning; fourth, that signs can be used to refer to objects or events that are not present. (Miles, 1990, p. 524.)

(I take the third criterion to mean that there is a systematic correlation between sign and referent.) All of these are necessary—but not sufficient—for T-naming. One of the essential aspects of T-naming is that there be no desire to *have* the object named—no ulterior motive.

However, Chantek showed some behavior that seems to be part of T-naming when he would show his caregivers some object (pp. 524f). Since he already had the object, it would seem that he had no other purpose for showing it than to get his caregivers to understand that he was thinking about it. This behavior, when combined with displaced reference, surely lays the groundwork for eventual T-naming.

Is T-naming a significant mark either of human language development in particular or of language development *simpliciter*? Granted that Helen Keller exhibited it after (and apparently *only* after) the well house, it would seem that it is significant for humans (or, at least, for her). And if Chantek either could easily have exhibited T-naming, or in fact did exhibit it (on occasion), it might *not* be unique to human language. It certainly makes for more sophisticated use of language (the ability to tell stories, the ability to fabricate), and it does make language learning easier. Yet there’s an awful lot of linguistic behavior that apes such as Chantek are capable of that makes one wonder why Terrace requires that, in order to T-name, the language user must *not* want the object. Chantek, for instance, learned labels for things he wanted, displayed displacement reference for things he wanted, and used language to deceive in order to get something (pp. 526–529). And Chantek, apparently, was capable of a metalinguistic achievement that, again, could underlie eventual T-naming:

Chantek used the blades of a scissors instead of his hands to make the sign for biting By transferring the total shape of the sign, including configuration and movement, to another means of expression, he showed that he understood that

the sign was an abstract representation in which the composite elements stood for something else. (Miles, 1990, p. 530.)

Indeed, some of the beginnings of what looks like T-naming can be seen in the following passages:

The second stage of development, that of subjective representation ... ranged from 2 years to almost $4\frac{1}{2}$ years of age.... In this stage, Chantek used his signs as symbolic representations, but his perspective remained subjective. He gave the first evidence of displacement ... and developed proximal pointing, which indicated that he had mental representations. ... He elaborated his deception and pretend play He showed evidence of planning through mental representations and signed to himself about objects not present. ... For the first time he also used signs in his deceptions.

... The third stage, nascent perspective taking, ranged from about $4\frac{1}{2}$ years to over 8 years of age, during which his vocabulary increased to 140 signs Chantek's representations became more objective and moved toward perspective taking, the ability to utilize the point of view of the other. ... Most important, he was able to take the perspective of the other by getting the caregiver's attention and directing the caregiver's eye gaze before he began to sign.

It was at this point that he invented his own signs. ... He clearly understood that signs were representational labels, and he immediately offered his hands to be molded when he wanted to know the name of an object. (Miles, 1990, pp. 534–535.)

How reminiscent of Helen Keller's post-well-house behavior, whether or not it is T-naming!

Against T-naming

Two arguments can be mounted *against* the significance of T-naming. The first, due to Greenfield and Savage-Rumbaugh (1990), is based on possible biases on the part of researchers. Terrace's claim that apes don't T-name is apparently supported by evidence such as that "Kanzi [a pygmy chimpanzee] had a much smaller proportion of indicatives to statements (4%) in comparison with requests (96%), than would be normal for a human child" (Greenfield & Savage-Rumbaugh, 1990, p. 568). But, as Greenfield and Savage-Rumbaugh point out, an alternative explanation is that this is an artifact of their artificial, human-controlled environment, in which they *must* request things. By contrast, "In the wild, a given animal might *state* his planned activity, rather than *requesting* it" (p. 568). They suggest that if we studied human language development without the assumption that children will eventually succeed in learning language, we might *not* ascribe T-naming to them at the analogous developmental stage at which we deny it to apes (p. 571).

The second, perhaps weaker, argument against T-naming focuses on just what it is that a speaker intends to communicate. T-naming certainly involves a desire to communicate—but to communicate what? For Terrace, it is the desire to communicate that the speaker is thinking of a distal object. The speaker is playing a sort of "guess what I'm thinking about" game, using a word that means *what he or she is*

thinking about. But that notion—what the speaker is thinking about—is ambiguous between the actual object (a *de re* interpretation) and the concept in the speaker's mind (a *de dicto* interpretation). However, since the speaker can be thinking of an object that doesn't exist, or a proposition that may lack a truth value, the *de re* interpretation can fail. Only the *de dicto* interpretation can be consistently maintained in all cases (Rapaport, 1976, 1978, 1981, 1985/1986, 1986a; Rapaport et al., 1997). As a consequence, *all* uses of names appear to turn out to be T-naming, i.e., the use of a name for something that one is merely thinking of.

However, whether or not non-humans are capable of it, T-naming by itself is probably not sufficient for full language use in conversation. As Fitch (2006, p. 370) puts it, "Chimpanzees certainly have important ingredients necessary for human language (for example, the ability to pair arbitrary signals with meanings), but they are still not discussing philosophy or even what they had for dinner yesterday."

Return to the well house

What was the significance of the well-house episode?

Negotiation is crucial to understanding language (Arrighi & Ferrario, 2005; Rapaport, 2002, 2003a). When a speaker uses a word in a conversation, all participants in the conversation must try to align the concepts that each finds or constructs in their minds. Often, an interlocutor has to *merge* two concepts (e.g., "Oh, John Williams the former conductor of the Boston Pops whom you're talking about is John Williams the composer whom I'm thinking of!") or to *split* one into two (e.g., "Oh, John Williams the classical guitarist whom you're talking about isn't John Williams the conductor and composer whom I'm thinking of!"); cf. Maida & Shapiro, 1982). So, one thing that was significant about Keller's experience at the well house was that two of *her* concepts merged or were aligned: *her* concept of water (previously linked to 'wah-wah') and *her* concept of *Sullivan's* concept of water. Prior to the well house, *Sullivan thought* that these were aligned ideas, but, in fact, they weren't.

Moreover, the well house itself played a significant role

... a key feature of human referring acts ... [is that] [t]hey are highly context sensitive or deictic. Parties to a referring act infer its referent from an *utterance* in a *context*. ... John Lyons argues that deixis is the source of reference, that "locating in context" rather than simply "tagging" is the heart of reference ... (Bruner, 1983, pp. 69–70).⁵⁹

Keller's experience was significant because the context was extremely simple: water in one hand, 'water' in the other.⁶⁰

One might reasonably expect to find, then, that the acquisition of referring procedures is heavily dependent on the "arranging" and simplifying of

⁵⁹ I consider other aspects of Bruner's book in Rapaport (2003a, Sect. 8). On the role of deixis in natural-language understanding, cf. Bruder et al. (1986); Rapaport, Segal, Shapiro, Zubin, Bruder, Duchan, Almeida et al. (1989); Rapaport, Segal, Shapiro, Zubin, Bruder, Duchan and Mark (1989); Duchan, Bruder, and Hewitt (1995).

⁶⁰ Actually, as we saw, there was a mug in the water hand, but it seems to have been ignored. Cf. Sect. 'Epiphany', observation 3, above.

contexts by the adult to assure that deictic demands be manageable for the child. (Bruner, 1983, p. 70.)

Keller's story is fascinating. Every teacher ought to read Sullivan's letters and reports on her teaching methods. Equally, Keller was an amazing pupil; one wonders what would have become of her had she *not* been blind and deaf! Consider the large number of syntactic systems with which she was familiar: finger spelling (the manual alphabet), lip reading (also tactually understood), the typewriter, three varieties of Braille, the Roman alphabet (again, tactually understood), oral speech (her own)—Keller's knowledge of speech is also akin to the CR; she had a "syntactic" knowledge of speech, since she couldn't hear herself (cf. Keller, 1905, p. 327)—Morse code, English, French, German, Latin, and Greek (and probably the Greek alphabet in some form).

Now, Searle-in-the-room also knows a syntactic system—squiggles—known to others as Chinese writing. The task for Searle-in-the-room is to get "beyond" the syntax. But what lies beyond? Ideas (mental entities?)? Objects? In general, of course, his task is to get to what the squiggles *mean*. How? Well, clearly, the more squiggles, the better. Note that *much* of Keller's learning was *book*-learning, which is purely syntactic (cf. Keller, 1905, pp. 30, 318; but cf. p. 317). But also Searle-in-the-room needs more experiences, even if only self-bodily ones. Ultimately, *all* such experiences are internal(ly represented), just as are (the experiences of) the squiggles.

Ditto for Keller. What she learned was not merely that everything has a name, but also that there is a naming relation, and she learned the name of naming. She was thus able to take a large network of mental entities, some of which were representatives of external objects (most of which, in turn, were internalized by the sense of touch) and some of which were representatives of external words (also internalized by the sense of touch), and partition it into two parts with explicit relations between them. She imposed a semantic structure on a domain hitherto only understood syntactically. When she was able to *organize* all her internal symbols such that some were names for others (and some of the "others" were directly linked to her experiences), she began to get beyond the syntax to the meanings (cf. Keller, 1905, p. 169). It was still a syntactic system, but now had semantic organization. *The organizing principle was discovered at the well house*. And that's how Helen Keller escaped from her CR.

Concluding remarks

In a footnote to an important essay on software engineering, Parnas wrote

It should be clear that while we cannot afford to use natural language specifications [for "program segments", p. 330] we cannot manage to do without natural language explanations. *Any formal structure is a hollow shell to most of us without a description of its intended interpretation.* ... [F]ormal specifications ... would be meaningless without a natural language description of the intended usage of the various functions and parameters. On the other hand, we insist that once the reader is familiar with the intended interpretation the specifications should answer all of his [sic] questions about the behavior of the programs without reference to the natural language text. (Parnas, 1972, p. 331, fn. 1; my italics.)

The italicized sentence in this passage is reminiscent of the CR. But note that, although we may understand “hollow” code in terms of natural language, Parnas assumes that we understand natural language directly. Or perhaps we only understand natural language via a further semantic interpretation, but then we must understand this semantics directly. The sequence of understanding one domain or language in terms of another must stop somewhere with a domain that we understand directly. And this “direct understanding” is what I have called “syntactic understanding” (Rapaport, 1986b, 1995, 2000; cf. Harel’s (1998) notion of “symbolic reasoning”).

Thesis 3 of syntactic semantics is that understanding is recursive—we understand one domain in terms of an antecedently understood one. The base case is a domain that is understood “directly”, in terms of itself. In such domains, some elements of the base domain are understood in terms of other elements of the base domain—i.e., they are understood syntactically. (For a related argument, see Kalderon, 2001.) In the case of language, linguistic elements (which are internal) are understood in terms of other internal, but non-linguistic (or “conceptual”), elements. This is how semantics can arise from syntax, and how natural-language understanding, by human or computer, is possible.

Acknowledgments An ancestor of this essay was originally written around 1992, was discussed in my seminar in Spring 1993 on “Semantics, Computation, and Cognition” at SUNY Buffalo, and first appeared in an unpublished technical report (Rapaport, 1996). I am grateful to Albert Goldfain, Frances L. Johnson, David Pierce, Stuart C. Shapiro, and the other members of the SNePS Research Group for comments.

References

- Allaire, E. B. (1963). Bare particulars. *Philosophical Studies*, 14; reprinted in Loux 1970, 235–244.
- Allaire, E. B. (1965). Another look at bare particulars. *Philosophical Studies*, 16; reprinted in Loux 1970, 250–257.
- Allen, J. F., & Perrault, C. R. (1980). Analyzing intentions in utterance. *Artificial Intelligence*, 15, 143–178.
- Anderson, M. L. (2003). Embodied cognition: A field guide. *Artificial Intelligence*, 149, 91–130
- Arahi, K., & Momouchi, Y. (1990). Learning of semantic concept in copular sentence (in Japanese), *IPSI SIG Reports*, 90(77).
- Arrighi, C., & Ferrario, R. (2005). The dynamic nature of meaning. In L. Magnani & R. Dossena (Eds.), *Computing, philosophy, and cognition* (pp. 295–312). London: [King’s] College Publications.
- Baker, R. (1967). Particulars: bare, naked, and nude. *Noûs*, 1, 211–212.
- Berners-Lee, T., & Fischetti, M. (1999). *Weaving the web*. New York: HarperCollins.
- Berners-Lee, T., Hendlar, J., & Lassila, O. (2001). The Semantic Web. *Scientific American* (17 May).
- Brown, R. (1973). *A first language*. Cambridge, MA: Harvard Univ. Press.
- Bruce, B. C. (1975). Generation as a social action. *Theoretical issues in natural language processing-I* (pp. 64–67). Morristown, NJ: Assoc. Comp. Ling.
- Bruder, G. A., Duchan, J. F., Rapaport, W. J., Segal, E. M., Shapiro, S. C., & Zubin, D. A. (1986). Deictic centers in narrative. *Tech. Rep. 86–20*. Buffalo: SUNY Buffalo Dept. Comp. Sci.
- Bruner, J. (1983). *Child’s talk*. New York: Norton.
- Castañeda, H.-N. (1980). *On philosophical method*. Bloomington, IN: Noûs Publications.
- Castañeda, H.-N. (1984). Philosophical refutations. In J. H. Fetzer (Ed.), *Principles of philosophical reasoning* (pp. 227–258). Totowa NJ: Rowman & Allenheld.
- Ceusters, W. (2005, October 27). *Ontology: The need for international coordination* [http://www.ncor.buffalo.edu/inaugural/ppt/ceusters.ppt

- Chappell, V. C. (1964). Particulars re-clothed. *Philosophical Studies*, 15, reprinted in Loux 1970, 245–249.
- Chrisley, R. (2003). Embodied artificial intelligence. *Artificial Intelligence*, 149, 131–150.
- Chun, S. A. (1987). SNePS implementation of possessive phrases. *SNeRG Tech. Note 19* (Buffalo: SUNY Buffalo Dept. Comp. Sci.) [<http://www.cse.buffalo.edu/sneps/Bibliography/chun87.pdf>]
- Clark, A., & Chalmers, D. J. (1998). The extended mind. *Analysis*, 58, 10–23.
- Cohen, P. R., & Levesque, H. J. (1985). Speech acts and rationality. *Proc. 23rd Annual Meeting, Assoc. Comp. Ling.* (pp. 49–60). Morristown, NJ: Assoc. Comp. Ling.
- Cohen, P. R., & Levesque, H. J. (1990). Rational interaction as the basis for communication. In P. R. Cohen, J. Morgan, & M. E. Pollack (Eds.), *Intentions in communication* (pp. 221–256). Cambridge, MA: MIT Press.
- Cohen, P. R., & Perrault, C. R. (1979). Elements of a plan-based theory of speech acts. *Cognitive Science*, 3, 177–212
- Damasio, A. R. (1989). Concepts in the brain. In Forum: What is a concept? *Mind and Language*, 4, 24–27.
- Davidson, D. (1967). The logical form of action sentences. In N. Rescher (Ed.), *The logic of decision and action*. Pittsburgh: Univ. Pittsburgh Press.
- Dehaene, S. (1992). Varieties of numerical abilities. *Cognition*, 44, 1–42.
- Dresner, E. (2002). Holism, language acquisition, and algebraic logic. *Linguistics and Philosophy*, 25, 419–452.
- Duchan, J. F., Bruder, G. A., & Hewitt, L. E. (Eds.). (1995), *Deixis in narrative*. Hillsdale, NJ: Erlbaum.
- Ehrlich, K. (1995). Automatic vocabulary expansion through narrative context. *Tech. Rep. 95-09*. Buffalo: SUNY Buffalo Dept. Comp. Sci.
- Ehrlich, K. (2004). Default reasoning using monotonic logic. *Proc., 15th Midwest Artif. Intel. & Cog. Sci. Conf.* (pp. 4–54) [<http://www.cse.buffalo.edu/~rapaport/CVA/ehrllich-maics-sub.pdf>]
- Ehrlich, K., & Rapaport, W. J. (1997). A computational theory of vocabulary expansion. *Proc. 19th Annual Conf., Cog. Sci. Soc.* (pp. 205–210). Mahwah, NJ: Erlbaum.
- Ehrlich, K., & Rapaport, W. J. (2004). A cycle of learning. *Proc., 26th Annual Conf., Cog. Sci. Soc.* Mahwah, NJ: Erlbaum, 2005, pp. 1555.
- Elgin, S. H. (1984). *Native tongue*. New York: DAW.
- Fitch, W. T. (2006). Hypothetically speaking. *American Scientist* (July–August), 369–370.
- Galbraith, M., & Rapaport, W. J. (Eds.). (1995). Where does I come from? Special Issue on subjectivity and the debate over computational cognitive science. *Minds and Machines*, 5, 513–620.
- Geach, P. T. (1967). Intentional identity. *The Journal of Philosophy*, 64, 627–632.
- Giere, R. N. (2002) Distributed cognition in epistemic cultures. *Phil. Sci.*, 69, 637–644.
- Goldfain, A. (2004). Using SNePS for mathematical cognition. [<http://www.cse.buffalo.edu/~ag33>]
- Goldfain, A. (2006). A computational theory of early mathematical cognition. [<http://www.cse.buffalo.edu/~ag33>]
- Greenfield, P. M., & Savage-Rumbaugh, E. S. (1990). Grammatical Combination in *Pan Paniscus*, in Parker & Gibson 1990 (pp. 540–577).
- Grosz, B. J., & Sidner, C. L. (1986) Attention, intentions, and the structure of discourse. *Computational Linguistics*, 12, 175–204.
- Haller, S. M. (1994). Interactive generation of plan descriptions and justifications. *Tech. Rep. 94-40*. Buffalo, NY: SUNY Buffalo Dept. Comp. Sci.
- Haller, S. M. (1995). Planning text for interactive plan explanations. In E. A. Yfantis (Ed.), *Intelligent systems* (pp. 61–67). Dordrecht: Kluwer.
- Harel, G. (1998). Two dual assertions. *The American Mathematical Monthly*, 105, 497–507.
- Harman, G. (1987). (Nonsolipsistic) Conceptual role semantics. In E. Lepore (Ed.), *New directions in semantics* (pp. 55–81). London: Academic.
- Harnad, S. (1990) The symbol grounding problem. *Physica D*, 42, 335–346.
- Hofstadter, D. R. (2001). Analogy as the core of cognition. In D. Gentner et al. (Eds.), *The analogical mind* (pp 499–538). Cambridge, MA: MIT Press.
- Hutchins, E. (1995a) *Cognition in the wild*. Cambridge, MA: MIT Press.
- Hutchins, E. (1995b). How a cockpit remembers its speeds. *Cog. Sci.*, 19, 265–288.
- Ismail, H. O., & Shapiro, S. C. (2000). Two problems with reasoning and acting in time. In A. G. Cohn et al. (Eds.), *Principles of knowledge representation and reasoning: Proc., 7th Int'l. Conf.* (pp. 355–365). San Francisco: Morgan Kaufmann.
- Iwańska, Ł. M., & Shapiro, S. C. (Eds.). (2000). *Natural language processing and knowledge representation*. Menlo Park CA/Cambridge MA: AAAI Press/MIT Press.

- Jackendoff, R. (2002). *Foundations of language*. Oxford: Oxford Univ. Press.
- Kalderon, M. E. (2001). Reasoning and representing. *Philosophical Studies*, 105, 129–160.
- Kamp, H., & Reyle, U. (1993). *From discourse to logic*. Dordrecht: Kluwer.
- Keller, H. (1903). *Optimism*. New York: Crowell.
- Keller, H. (1905). *The story of my life*. Garden City, NY: Doubleday (1954).
- Kibby, M. W., Rapaport, W. J., Wieland, K. M., & Dechert, D. A. (in press). CSI: Contextual semantic investigation for word meaning. In L. A. Baines (Ed.), *Multisensory learning*. Alexandria, VA: Association for Supervision and Curriculum Development [http://www.cse.buffalo.edu/~rapaport/CVA/CSI.pdf]
- Lammens, J. (1994). A computational model of color perception and color naming. *Tech. Rep. 94-26*. Buffalo: SUNY Buffalo Dept. Comp. Sci.). [http://www.cse.buffalo.edu/sneps/Bibliography/lammens94.pdf]
- Leiber, J. (1996). Helen Keller as cognitive scientist. *Philosophical Psychology*, 9, 419–440.
- Loux, M. J. (Ed.). (1970). *Universals and particulars*. Garden City, NY: Anchor.
- Maida, A. S., & Shapiro, S. C. (1982). Intensional concepts in propositional semantic networks. *Cognitive Science*, 6, 291–330.
- Maloney, J. C. (1987). The right stuff. *Synthese*, 70, 349–372.
- Maloney, J. C. (1989). *The mundane matter of the mental language*. Cambridge, UK: Cambridge Univ. Press.
- Martins, J., & Shapiro, S. C. (1988). A model for belief revision. *Artificial Intelligence*, 35, 25–79.
- Mayes, A. R. (1991). Review of H. Damasio & A. R. Damasio, *Lesion analysis in neuropsychology (inter alia)*. *British Journal of Psychology*, 2, 109–112.
- McDermott, D. (1981). Artificial intelligence meets natural stupidity. In J. Haugeland (Ed.), *Mind design* (pp. 143–160). Cambridge, MA: MIT Press.
- Meinong, A. (1904). Über Gegenstandstheorie. In R. Haller (Ed.), *Alexius Meinong Gesamtausgabe (Vol II. pp. 481–535)*. Graz, Austria: Akademische Druck- u. Verlagsanstalt (1971).
- Miles, H. L. W. (1990). The cognitive foundations for reference in a Signing Orangutan. (pp. 511–539).
- Morris, C. (1938). *Foundations of the theory of signs*. Chicago: Univ. Chicago Press.
- Nagel, T. (1986). *The view from nowhere*. New York: Oxford Univ. Press.
- Neal, J. G., & Shapiro, S. C. (1987). Knowledge-based Parsing. In L. Bolc (Ed.), *Natural language parsing systems* (pp. 49–92). Berlin: Springer-Verlag.
- Papineau, D. (1998). “Get a Grip”, review of Wilson 1998. *New York Times Book Review* (19 July), p. 9.
- Parker, S. T., & Gibson, K. R. (Eds.). (1990). “Language” and intellect in monkeys and apes. Cambridge, UK: Cambridge Univ. Press.
- Parnas, D. L. (1972). A technique for software module specification with examples. *Communication of the Association for Computing Machinery*, 15, 330–336.
- Preston, J., & Bishop, M. (Eds.). (2002). *Views into the Chinese Room*. Oxford: Oxford Univ. Press.
- Proudfoot, D. (2002). Wittgenstein’s anticipation of the Chinese Room, in Preston & Bishop 2002 (pp. 167–180).
- Putnam, H. (1975). The meaning of ‘meaning’. reprinted in *Mind, language and reality* (pp. 215–271). Cambridge, UK: Cambridge Univ. Press.
- Quine, W. V. O. (1953). “On what there is” and “Logic and the reification of universals”. In *From a logical point of view* (Chs. I, VI). New York: Harper & Row.
- Quine, W. V. O. (1960). *Word and object*. Cambridge, MA: MIT Press.
- Quine, W. V. O. (1969). Ontological relativity. In W. V. O. Quine (Ed.), *Ontological relativity and other essays* (pp. 26–68). New York: Columbia Univ. Press.
- Rapaport, W. J. (1976). *Intentionality and the structure of existence*. Ph.D. diss. Bloomington: Indiana Univ. Dept. Phil.
- Rapaport, W. J. (1978). Meinongian theories and a Russellian paradox. *Noûs*, 12, 153–110; errata, *Noûs*, 13 (1979) 125.
- Rapaport, W. J. (1981). How to make the world fit our language. *Grazer Philosophische Studien*, 14, 1–21.
- Rapaport, W. J. (1982). Unsolvable problems and philosophical progress. *American Philosophical Quarterly*, 19, 289–298.
- Rapaport, W. J. (1985). Machine understanding and data abstraction in Searle’s Chinese Room. *Proc., 7th Annual Meeting, Cog. Sci. Soc.* (pp. 341–345) Hillsdale, NJ: Erlbaum.
- Rapaport, W. J. (1985/1986). Non-existent objects and epistemological ontology. *Grazer Philosophische Studien* 25/26, 61–95.

- Rapaport, W. J. (1986a). Logical foundations for belief representation. *Cognitive Science*, 10, 371–422.
- Rapaport, W. J. (1986b). Searle's experiments with thought. *Philosophy of Science*, 53, 271–279.
- Rapaport, W. J. (1988a). To think or not to think. *Noûs*, 22, 585–609.
- Rapaport, W. J. (1988b). Syntactic semantics. In J. H. Fetzer (Ed.), *Aspects of artificial intelligence* (pp. 1–131). Dordrecht: Kluwer.
- Rapaport, W. J. (1990). Computer processes and virtual persons. *Tech. Rep. 90-13* (Buffalo: SUNY Buffalo Dept. Comp. Sci., May 1990) [<http://www.cse.buffalo.edu/~rapaport/Papers/cole.tr.17my90.pdf>]
- Rapaport, W. J. (1991a). Predication, fiction, and artificial intelligence. *Topoi*, 10, 79–111.
- Rapaport, W. J. (1991b). Meinong, Alexius I: Meinongian semantics. In H. Burkhardt, & B. Smith (Eds.), *Handbook of metaphysics and ontology* (pp. 516–519). Munich: Phil Verlag.
- Rapaport, W. J. (1993). Because mere calculating isn't thinking. *Minds & Machines*, 3, 11–20.
- Rapaport, W. J. (1995). Understanding understanding. In J. E. Tomberlin (Ed.), *Philosophical perspectives: AI, connectionism, and philosophical psychology* (Vol. 9, pp. 49–88). Atascadero, CA: Ridgeview.
- Rapaport, W. J. (1996). *Understanding understanding*, *Tech. Rep. 96-26*. Buffalo: SUNY Buffalo Dept. Comp. Sci. [<http://www.cse.buffalo.edu/tech-reports/96-26.ps>]
- Rapaport, W. J. (1998). How minds can be computational systems. *JETAI*, 10, 403–419.
- Rapaport, W. J. (1999). Implementation is semantic interpretation. *Monist*, 82, 109–130.
- Rapaport, W. J. (2000). How to pass a Turing test. *Journal of Logic Language and Information*, 9, 467–490.
- Rapaport, W. J. (2002). Holism, conceptual-role semantics, and syntactic semantics. *Minds & Machines*, 12, 3–59.
- Rapaport, W. J. (2003a). What did you mean by that? *Minds & Machines*, 13, 397–427.
- Rapaport, W. J. (2003b). What is the 'context' for contextual vocabulary acquisition? In P. P. Slezak (Ed.), *Proc., 4th Int'l. Conf. Cog. Sci./7th Australasian Soc. Cog. Sci. Conf* (Vol. 2, pp. 547–552). Sydney: Univ. New South Wales.
- Rapaport, W. J. (2005a). In defense of contextual vocabulary acquisition. In A. Dey et al. (Eds.), *Proc., 5th Int'l. & Interdisc. Conf., Modeling and Using Context* (pp. 396–409). Berlin: Springer-Verlag Lecture Notes in AI 3554.
- Rapaport, W. J. (2005b). Implementation is semantic interpretation: further thoughts. *JETAI*, 17, 385–417.
- Rapaport, W. J. (2005c). Review of Shieber 2004. *Computational Linguistics*, 31, 407–412.
- Rapaport, W. J. (2005d). The Turing test. In: *Ency. Lang. & Ling.*, (2nd ed., Vol. 13, pp. 151–159). Oxford: Elsevier.
- Rapaport, W. J. (2006). Review of Preston & Bishop 2002. *Australasian Journal of Philosophy*, 84, 129–133.
- Rapaport, W. J., & Ehrlich, K. (2000). A Computational theory of vocabulary acquisition. In Iwańska & Shapiro 2000 (pp. 347–375).
- Rapaport, W. J., & Kibby, M. W. (2002). Contextual vocabulary acquisition. In N. Callaos et al. (Eds.), *Proc., 6th World Multiconf., Systemics, Cybernetics & Informatics* (Vol. 2, pp. 261–266). Orlando: Int'l. Inst. Informatics & Systemics.
- Rapaport, W. J., Segal, E. M., Shapiro, S. C., Zubin, D. A., Bruder, G. A., Duchan, J. F., Almeida, M. J., Daniels, J. H., Galbraith, M. M., Wiebe, J. M., & Yuhan, A. H. (1989). Deictic centers and the cognitive structure of narrative comprehension. *Tech. Rep. 89-01*. Buffalo: SUNY Buffalo Dept. Comp. Sci. [<http://www.cse.buffalo.edu/~rapaport/Papers/DC.knuf.pdf>]
- Rapaport, W. J., Segal, E. M., Shapiro, S. C., Zubin, D. A., Bruder, G. A., Duchan, J. F., & Mark, D. M. (1989). Cognitive and computer systems for understanding narrative text. *Tech. Rep. 89-07*. Buffalo: SUNY Buffalo Dept. Comp. Sci.
- Rapaport, W. J., Shapiro, S. C., & Wiebe, J. M. (1986). Quasi-indicators, knowledge reports, and discourse. *Tech. Rep. 86-15*. Buffalo: SUNY Buffalo Dept. Comp. Sci; revised version published as Rapaport, Shapiro, and Wiebe (1997).
- Rapaport, W. J., Shapiro, S. C., & Wiebe, J. M. (1997). Quasi-indexicals and knowledge reports. *Cognitive Science*, 21, 63–107.
- Rogers, H. Jr. (1959). The present theory of Turing machine computability. *Journal of the Society for Industrial and Applied Mathematics*, 7, 114–130.
- Santore, J. F., & Shapiro, S. C. (2004). Identifying perceptually indistinguishable objects. In S. Coradeschi & A. Saffiotti (Eds.), *Anchoring symbols to sensor data* (pp. 1–9). Menlo Park, CA: AAAI Press.

- Searle, J. R. (1980). Minds, brains, and programs. *The Behavioral and Brain Sciences*, 3, 417–457.
- Searle, J. R. (1993). The failures of computationalism. *Think* (Vol 2 (June)), pp. 68–71. Tilburg: Tilburg Univ. Inst. Lang. Tech. & AI.
- Searle, J. R. (2002). Twenty-one years in the Chinese Room. In Preston & Bishop 2002 (pp. 51–69).
- Sellars, W. (1963). *Science, perception and reality*. London: Routledge & Kegan Paul.
- Severo, R. (1999). Clifton Fadiman, a wordsmith known for his encyclopedic knowledge, is dead at 95. *New York Times* (21 June): B5.
- Shapiro, S. C. (1979). The SNePS semantic network processing system. In N. Findler (Ed.), *Associative networks* (pp. 179–203). New York: Academic.
- Shapiro, S. C. (1981). COCCI: A deductive semantic network program for solving microbiology unknowns. *Tech. Rep. 173* Buffalo: SUNY Buffalo Dept. Comp. Sci.
- Shapiro, S. C. (1982). Generalized augmented transition network grammars for generation from semantic networks. *American Journal of Computational Linguistics*, 8, 12–25.
- Shapiro, S. C. (1986). Symmetric relations, intensional individuals, and variable binding. *Proceedings of IEEE*, 74, 1354–1363.
- Shapiro, S. C. (1989). The CASSIE projects. In J. P. Martins & E. M. Morgado (Eds.), *EPIA89: 4th Portuguese Conf. AI, Proc.* (pp. 362–380). Berlin: Springer-Verlag Lecture Notes in AI 390.
- Shapiro, S. C. (1992). Artificial intelligence. In S. C. Shapiro (Ed.), *Ency. AI* (2nd ed., pp. 54–57) New York: Wiley.
- Shapiro, S. C. (1998). Embodied cassie. In *Cog. Robotics* (pp. 136–143). Menlo Park, CA: AAAI Press.
- Shapiro, S. C. (2000). SNePS: A logic for natural language understanding and commonsense reasoning. In Iwańska & Shapiro 2000 (pp. 175–195).
- Shapiro, S. C. (2003). FevahrCassie. *SNeRG Tech. Note 35* (Buffalo, NY: SUNY Buffalo Dept. Comp. Sci. & Eng'g.) [<http://www.cse.buffalo.edu/~shapiro/Papers/buildingFevahrAgents.pdf>]
- Shapiro, S. C., & Ismail, H. O. (2003). Anchoring in a grounded layered architecture with integrated reasoning. *Robotics and Autonomous Systems*, 43, 97–108.
- Shapiro, S. C., Ismail, H. O., & Santore, J. F. (2000). Our dinner with Cassie. *Working notes for the AAAI 2000 Spring symposium on natural dialogues with practical robotic devices* (pp. 57–61). Menlo Park, CA: AAAI Press.
- Shapiro, S. C., & Rapaport, W. J. (1987). SNePS considered as a fully intensional propositional semantic network. In N. Cercone & G. McCalla (Eds.), *The knowledge frontier* (pp. 262–315). New York: Springer-Verlag.
- Shapiro, S. C., & Rapaport, W. J. (1991). Models and minds. In R. Cummins & J. Pollock (Eds.), *Philosophy and AI* (pp. 215–259). Cambridge, MA: MIT Press.
- Shapiro, S. C., & Rapaport, W. J. (1992). The SNePS family. *Computers & Mathematics with Applications*, 23, 243–275.
- Shapiro, S. C., & Rapaport, W. J. (1995). An introduction to a computational reader of narratives. In Duchan et al. 1995 (pp. 79–105).
- Shapiro, S. C., Rapaport, W. J., Cho, S.-H., Choi, J., Feit, E., Haller, S., Kankiewicz, J., & Kumar, D. (1996). A dictionary of SNePS case frames. [<http://www.cse.buffalo.edu/sneps/Manuals/dictionary.pdf>]
- Shapiro, S. C., & the SNePS Research Group (2006). SNePS. *Wikipedia* [<http://www.en.wikipedia.org/wiki/SNePS>]
- Sheckley, R. (1954). Ritual. In R. Sheckley (Ed.), *Untouched by human hands* (pp. 155–165). New York: Ballantine.
- Shieber, S. M. (2004). *The Turing test*. Cambridge, MA: MIT Press.
- Smith, B. C. (1982). Linguistic and Computational Semantics. *Proc., 20th Annual Meeting, Assoc. Comp. Ling.* (pp. 9–15). Morristown, NJ: Assoc. Comp. Ling.
- Spärck Jones, K. (1967). Dictionary circles, Tech. *Memo. TM-3304*. Santa Monica, CA: System Development Corp.
- Srihari, R. K. (1991a). PICTION: A system that uses captions to label human faces in newspaper photographs. *Proc., 9th Nat'l. Conf. AI* (pp. 80–85). Menlo Park, CA: AAAI Press/MIT Press.
- Srihari, R. K. (1991b). Extracting visual information from text. *Tech. Rep. 91-17*. Buffalo: SUNY Buffalo Dept. Comp. Sci.
- Srihari, R. K. (1993). Intelligent document understanding. *Proc., Int'l. Conf. Document analysis and recognition* (pp. 664–667).
- Srihari, R. K. (1994). Use of collateral text in understanding photos in documents. *Proc., Conf. Applied imagery and pattern recognition* (pp. 186–199).

- Srihari, R. K., & Rapaport, W. J. (1989). Extracting visual information from text. *Proc., 11th Annual Conf. Cog. Sci. Soc.* (pp. 364–371). Hillsdale, NJ: Erlbaum.
- Srihari, R. K., & Rapaport, W. J. (1990). Combining linguistic and pictorial information. In D. Kumar (Ed.), *Current trends in SNePS* (pp. 5–96). Berlin: Springer-Verlag Lecture Notes in AI 437.
- Swan, J. (1994). Touching words. In M. Woodmansee & P. Jaszi (Eds.), *The construction of authorship* (pp. 57–100). Durham, NC: Duke Univ. Press.
- Tarski, A. (1969). Truth and proof. *Scientific American*, 220, 63–70, 75–77.
- Taylor, J. G. (2002). Do virtual actions avoid the Chinese Room? In Preston & Bishop 2002 (pp. 269–293).
- Terrace, H. S. (1985). In the beginning was the ‘name’. *American Psychologist*, 40, 1011–1028.
- Terrace, H. S. (1991). Letter to the editor. *New York Review of Books*, 3(15), 53.
- Thomason, R. H. (2003). Dynamic contextual intensional logic. In P. Blackburn et al. (Eds.), *CONTEXT 2003* (pp. 328–341). Berlin: Springer-Verlag Lecture Notes in AI 2680.
- Turing, A. M. (1950). Computing machinery and intelligence. *Mind*, 59, 433–460.
- Vauclair, J. (1990). Primate cognition. In Parker & Gibson 1990 (pp. 312–329).
- Von Glasersfeld, E. (1977). Linguistic communication. In D. M. Rumbaugh (Ed.), *Language learning by a Chimpanzee* (pp. 55–71). New York: Academic.
- Weizenbaum, J. (1966). ELIZA—A computer program for the study of natural language communication between man and machine. *Communications of the Association for Computing Machinery* 9, 36–45.
- Wilson, F. R. (1998). *The hand*. New York: Pantheon.
- Winston, P. H. (1975). Learning structural descriptions from examples. Reprinted In R. J. Brachman, & H. J. Levesque (Eds.), *Readings in knowledge representation* (pp. 141–168). Los Altos, CA: Morgan Kaufmann (1985).
- Wittgenstein, L. (1958). *Philosophical investigations* (3rd ed., trans. by G. E. M. Anscombe). New York: Macmillan.
- Woods, W. A. (1975). What’s in a link. In D. G. Bobrow & A. Collins (Eds.), *Representation and understanding* (pp. 35–82). New York: Academic.
- Zuckermann, L. S. (1991). Letter to the editor. *New York Review of Books*, 3(15), 53.