

The Chinese Room Argument

The Chinese room argument - John Searle's (1980a) thought experiment and associated (1984) derivation - is one of the best known and widely credited counters to claims of artificial intelligence (AI), i.e., to claims that computers *do* or at least *can* (someday might) think. According to Searle's original presentation, the argument is based on two truths: *brains cause minds*, and *syntax doesn't suffice for semantics*. Its target, Searle dubs "strong AI": "according to strong AI," according to Searle, "the computer is not merely a tool in the study of the mind, rather the appropriately programmed computer really *is* a mind in the sense that computers given the right programs can be literally said to *understand* and have other cognitive states" (1980a, p. 417). Searle contrasts "strong AI" to "weak AI". According to weak AI, according to Searle, computers just *simulate* thought, their seeming understanding isn't real (just as-if) understanding, their seeming calculation as-if calculation, etc.; nevertheless, computer simulation is useful for *studying* the mind (as for studying the weather and other things).

[Table of Contents](#) (Clicking on the links below will take you to that part of this article)

- [The Chinese Room Thought Experiment](#)
- [Replies and Rejoinders](#)
 - [The Systems Reply](#)
 - [The Robot Reply](#)
 - [The Brain Simulator Reply](#)
 - [The Combination Reply](#)
 - [The Other Minds Reply](#)
 - [The Many Mansions Reply](#)
- [Searle's "Derivation from Axioms."](#)
- [Continuing Dispute](#)
 - [Initial Objections & Replies](#)
 - [The Connectionist Reply](#)
- [Summary Analysis](#)
- [Postscript](#)
- [Sources](#)

The Chinese Room Thought Experiment

Against "strong AI," Searle (1980a) asks you to imagine yourself a monolingual English speaker "locked in a room, and given a large batch of Chinese writing" plus "a second batch of Chinese script" and "a set of rules" in English "for correlating the second batch with the first batch." The rules "correlate one set of formal symbols with another set of formal symbols"; "formal" (or "syntactic") meaning you "can identify the symbols entirely by their shapes." A third batch of Chinese symbols and more instructions in English enable you "to correlate elements of this third batch with elements of the first two batches" and instruct you, thereby, "to give back certain sorts of Chinese symbols with certain sorts of shapes in response." *Those giving you the symbols* "call the first batch 'a script' [a data structure with natural language processing applications], "they call the second batch 'a story', and they call the third batch 'questions'; the symbols you give back "they call . . . 'answers to the questions'"; "the set of rules in English . . . they call 'the program': *you yourself* know none of this. Nevertheless, you "get so good at following the instructions" that "from the point of view of someone outside the room" your responses are "absolutely indistinguishable from those of Chinese speakers." Just by looking at your answers, nobody can tell you "don't speak a word of Chinese." Producing answers "by manipulating uninterpreted formal symbols," it seems "[a]s far as the Chinese is concerned," you "simply behave like a computer"; specifically, like a computer running Schank and Abelson's (1977) "Script Applier Mechanism" story understanding program (SAM), which Searle's takes for his example. But in imagining *himself* to be the person in the room, Searle thinks it's "quite obvious . . . I do not understand a word of the Chinese stories. I have inputs and outputs that are indistinguishable from those of the native Chinese speaker, and I can have any formal program you like, but I still understand nothing." "For the same reasons," Searle concludes, "Schank's computer understands nothing of any stories" since "the computer has nothing more than I have in the case where I understand nothing" (1980a, p. 418). Furthermore, since in the thought experiment "nothing . . . depends on the details of Schank's programs," the same "would apply to any [computer] simulation" of any "human mental phenomenon" (1980a, p. 417); that's all it would be, simulation. Contrary to "strong AI", then, no matter how intelligent-seeming a computer *behaves* and no matter what *programming* makes it behave that way, since the symbols it processes are meaningless (lack semantics) *to it*, it's not really intelligent. It's not actually thinking. Its internal states and processes, being purely *syntactic*, lack semantics (meaning); so, it doesn't really have *intentional* (i.e., meaningful) *mental states*.

[Back to Table of Contents](#)

Replies and Rejoinders

Having laid out the example and drawn the aforesaid conclusion, Searle considers several replies offered when he "had the occasion to present this example to a number of workers in artificial intelligence" (1980a, p. 419). Searle offers rejoinders to these various replies.

The Systems Reply

The Systems Reply suggests that the Chinese room example encourages us to focus on the wrong agent: the thought experiment encourages us to mistake the would-be subject-possessed-of-mental-states for the person in the room. The systems reply grants that "the *individual who is locked in the room* does not understand the story" but maintains that "he is merely part of a whole system, and the *system* does understand the story" (1980a, p. 419: my emphases). Searle's main rejoinder to this is to "let the individual internalize all . . . of the system" by memorizing the rules and script and doing the lookups and other operations in their head. "All the same," Searle maintains, "he understands nothing of the Chinese, and . . . neither does the system, because there isn't anything in the system that isn't in him. If he doesn't understand then there is no way the system could understand because the system is just part of him" (1980a, p. 420). Searle also insists the systems reply would have the absurd consequence that "mind is everywhere." For instance, "there is a level of description at which my stomach does information processing" there being "nothing to prevent [describers] from treating the input and output of my digestive organs as information if they so desire." Besides, Searle contends, it's just ridiculous to say "that while [the] person doesn't understand Chinese, somehow the conjunction of that person and bits of paper might" (1980a, p. 420).

[Back to Table of Contents](#)

The Robot Reply

The Robot Reply - along lines favored by contemporary causal theories of reference - suggests what prevents the person in the Chinese room from attaching meanings to (and thus presents them from understanding) the Chinese ciphers is the sensory-motoric disconnection of the ciphers from the realities they are supposed to represent: to promote the "symbol" manipulation to genuine understanding, according to this causal-theoretic line of thought, the manipulation needs to be grounded in the outside world via the agent's causal relations to the things to which the ciphers, *as symbols*, apply. If we "put a computer inside a robot" so as to "operate the robot in such a way that the robot does something very much like perceiving, walking, moving about," however, then the "robot would," according to this line of thought, "unlike Schank's computer, have genuine understanding and other mental states" (1980a, p. 420). Against the Robot Reply Searle maintains "the same experiment applies" with only slight modification. Put the room, with Searle in it, inside the robot; imagine "some of the Chinese symbols come from a television camera attached to the robot" and that "other Chinese symbols that [Searle is] giving out serve to make the motors inside the robot move the robot's legs or arms." Still, Searle asserts, "I don't understand anything except the rules for symbol manipulation." He explains, "by instantiating the program I have no [mental] states of the relevant [meaningful, or intentional] type. All I do is follow formal instructions about manipulating formal symbols." Searle also charges that the robot reply "tacitly concedes that cognition is not solely a matter of formal symbol

manipulation" after all, as "strong AI" supposes, since it "adds a set of causal relation[s] to the outside world" (1980a, p. 420).

[Back to Table of Contents](#)

The Brain Simulator Reply

The Brain Simulator Reply asks us to imagine that the program implemented by the computer (or the person in the room) "doesn't represent information that we have about the world, such as the information in Schank's scripts, but simulates the actual sequence of neuron firings at the synapses of a Chinese speaker when he understands stories in Chinese and gives answers to them." Surely then "we would have to say that the machine understood the stories"; or else we would "also have to deny that native Chinese speakers understood the stories" since "[a]t the level of the synapses" there would be no difference between "the program of the computer and the program of the Chinese brain" (1980a, p. 420). Against this, Searle insists, "even getting this close to the operation of the brain is still not sufficient to produce understanding" as may be seen from the following variation on the Chinese room scenario. Instead of shuffling symbols, we "have the man operate an elaborate set of water pipes with valves connecting them." Given some Chinese symbols as input, the program now tells the man "which valves he has to turn off and on. Each water connection corresponds to synapse in the Chinese brain, and the whole system is rigged so that after . . . turning on all the right faucets, the Chinese answer pops out at the output end of the series of pipes." Yet, Searle thinks, obviously, "the man certainly doesn't understand Chinese, and neither do the water pipes." "The problem with the brain simulator," as Searle diagnoses it, is that it simulates "only the formal structure of the sequence of neuron firings": the insufficiency of this formal structure for producing meaning and mental states "is shown by the water pipe example" (1980a, p. 421).

[Back to Table of Contents](#)

The Combination Reply

The Combination Reply supposes all of the above: a computer lodged in a robot running a brain simulation program, considered as a unified system. Surely, now, "we would have to ascribe intentionality to the system" (1980a, p. 421). Searle responds, in effect, that since none of these replies, taken alone, has any tendency to overthrow his thought experimental result, neither do all of them taken together: zero times three is naught. Though it would be "rational and indeed irresistible," he concedes, "to accept the hypothesis that the robot had intentionality, as long as we knew nothing more about it" the acceptance would be simply based on the assumption that "if the robot looks and behaves sufficiently like us then we would suppose, until proven otherwise, that it must have mental states like ours that cause and are expressed by its behavior." However, "[i]f we knew independently how to account for its behavior without such assumptions," as with computers, "we would not attribute intentionality to it, especially if we knew it had a formal program" (1980a, p. 421).

The Other Minds Reply

The Other Minds Reply reminds us that how we "know other people understand Chinese or anything else" is "by their behavior." Consequently, "if the computer can pass the behavioral tests as well" as a person, then "if you are going to attribute cognition to other people you must in principle also attribute it to computers" (1980a, p. 421). Searle responds that this misses the point: it's "not. . . *how I know* that other people have cognitive states, but rather *what it is* that I am attributing when I attribute cognitive states to them. The thrust of the argument is that it couldn't be just computational processes and their output because the computational processes and their output can exist without the cognitive state" (1980a, p. 420-421: my emphases).

[Back to Table of Contents](#)

The Many Mansions Reply

The Many Mansions Reply suggests that even if Searle is right in his suggestion that programming cannot suffice to cause computers to have intentionality and cognitive states, other means besides programming might be devised such that computers may be imbued with whatever does suffice for intentionality by these other means. This too, Searle says, misses the point: it "trivializes the project of Strong AI by redefining it as whatever artificially produces and explains cognition" abandoning "the original claim made on behalf of artificial intelligence" that "mental processes are computational processes over formally defined elements." If AI is not identified with that "precise, well defined thesis," Searle says, "my objections no longer apply because there is no longer a testable hypothesis for them to apply to" (1980a, p. 422).

[Back to Table of Contents](#)

Searle's "Derivation from Axioms."

Besides the Chinese room thought experiment, Searle's more recent presentations of the Chinese room argument feature - with minor variations of wording and in the ordering of the premises - a formal "derivation from axioms" (1989, p. 701). The derivation, according to Searle's 1990 formulation proceeds from the following three axioms (1990, p. 27):

- (A1) Programs are formal (syntactic).
- (A2) Minds have mental contents (semantics).
- (A3) Syntax by itself is neither constitutive of nor sufficient for semantics.

to the conclusion:

(C1) Programs are neither constitutive of nor sufficient for minds.

Searle then adds a fourth axiom (p. 29):

(A4) Brains cause minds.

from which we are supposed to "immediately derive, trivially" the conclusion:

(C2) Any other system capable of causing minds would have to have causal powers (at least) equivalent to those of brains.

whence we are supposed to derive the further conclusions:

(C3) Any artifact that produced mental phenomena, any artificial brain, would have to be able to duplicate the specific causal powers of brains, and it could not do that just by running a formal program.

(C4) The way that human brains actually produce mental phenomena cannot be solely by virtue of running a computer program.

On the usual understanding, the Chinese room experiment subserves this derivation by "shoring up axiom 3" (Churchland & Churchland 1990, p. 34).

[Back to Table of Contents](#)

Continuing Dispute

To call the Chinese room controversial would be an understatement. Beginning with objections published along with Searle's original (1980a) presentation, opinions have drastically divided, not only about whether the Chinese room argument is cogent; but, among those who think it is, as to why it is; and, among those who think it is not, as to why not. This discussion includes several noteworthy threads.

[Back to Table of Contents](#)

Initial Objections & Replies

Initial Objections & Replies to the Chinese room argument besides filing new briefs on behalf of many of the forenamed replies(e.g., Fodor 1980 on behalf of "the Robot Reply") take, notably, two tacks. One

tack, taken by Daniel Dennett (1980), among others, decries the dualistic tendencies discernible, for instance, in Searle's methodological maxim "always insist on the first-person point of view" (Searle 1980b, p. 451). Another tack notices that the symbols Searle-in-the-room processes are *not* meaningless ciphers, they're Chinese inscriptions. So they *are* meaningful; and so is Searle's processing of them in the room; whether he knows it or not. In reply to this second sort of objection, Searle insists that what's at issue here is *intrinsic intentionality* in contrast to the merely *derived intentionality* of inscriptions and other linguistic signs. Whatever meaning Searle-in-the-room's computation might derive from the meaning of the Chinese symbols which he processes will not be intrinsic to the process or the processor but "observer relative," existing only in the minds of beholders such as the native Chinese speakers outside the room. "Observer-relative ascriptions of intentionality are always dependent on the intrinsic intentionality of the observers" (Searle 1980b, pp. 451-452). The nub of the experiment, according to Searle's attempted clarification, then, is this: "instantiating a program could not be constitutive of intentionality, because it would be possible for an agent [e.g., Searle-in-the-room] to instantiate the program and still not have the *right kind* of intentionality" (Searle 1980b, pp. 450-451: my emphasis); the *intrinsic* kind. Though Searle unapologetically identifies *intrinsic* intentionality with *conscious* intentionality, still he resists Dennett's and others' imputations of dualism. Given that what it is we're attributing in attributing mental states is conscious intentionality, Searle maintains, insistence on the "first-person point of view" is warranted; because "the ontology of the mind is a first-person ontology": "the mind consists of qualia [subjective conscious experiences] . . . right down to the ground" (1992, p. 20). This thesis of Ontological Subjectivity, as Searle calls it in more recent work, is *not*, he insists, some dualistic invocation of discredited "Cartesian apparatus" (Searle 1992, p. xii), as his critics charge; it simply reaffirms commonsensical intuitions that behavioristic views and their functionalistic progeny have, for too long, highhandedly, dismissed. This commonsense identification of thought with consciousness, Searle maintains, is readily reconcilable with thoroughgoing physicalism when we conceive of consciousness as both caused by and realized in underlying brain processes. Identification of thought with consciousness along these lines, Searle insists, is not dualism; it might more aptly be styled monist interactionism (1980b, p. 455-456) or (as he now prefers) "biological naturalism" (1992, p. 1).

[Back to Table of Contents](#)

The Connectionist Reply

The Connectionist Reply (as it might be called) is set forth - along with a recapitulation of the Chinese room argument and a rejoinder by Searle - by Paul and Patricia Churchland in a 1990 *Scientific American* piece. The Churchlands criticize the crucial third "axiom" of Searle's "derivation" by attacking his would-be supporting thought experimental result. This putative result, they contend, gets much if not all of its plausibility from the lack of neurophysiological verisimilitude in the thought-experimental setup. Instead of imagining Searle working alone with his pad of paper and lookup table, like the Central Processing Unit of a serial architecture machine, the Churchlands invite us to imagine a more brainlike *connectionist* architecture. Imagine Searle-in-the-room, then, to be just one of very many agents, all working in parallel, each doing their own small bit of processing (like the many neurons of

the brain). Since Searle-in-the-room, in this revised scenario, does only a very small portion of the total computational job of generating sensible Chinese replies in response to Chinese input, *naturally* he himself does not comprehend the whole process; so we should hardly expect *him* to grasp or to be conscious of the meanings of the communications he is involved, in such a minor way, in processing. Searle counters that this Connectionist Reply - incorporating, as it does, elements of both systems and brain-simulator replies - can, like these predecessors, be decisively defeated by appropriately tweaking the thought-experimental scenario. Imagine, if you will, a Chinese gymnasium, with many monolingual English speakers working in parallel, producing output indistinguishable from that of native Chinese speakers: each follows their own (more limited) set of instructions in English. Still, Searle insists, obviously, none of these *individuals* understands; and neither does the whole company of them collectively. It's intuitively utterly obvious, Searle maintains, that no one and nothing in the revised "Chinese gym" experiment understands a word of Chinese either individually or collectively. Both individually and collectively, nothing is being done in the Chinese gym except meaningless syntactic manipulations from which intentionality and consequently meaningful thought could not conceivably arise.

[Back to Table of Contents](#)

Summary Analysis

Searle's Chinese Room experiment parodies the Turing test, a test for artificial intelligence proposed by Alan Turing (1950) and echoing René Descartes' suggested means for distinguishing thinking souls from unthinking automata. Since "it is not conceivable," Descartes says, that a machine "should produce different arrangements of words so as to give an appropriately meaningful answer to whatever is said in its presence, as even the dullest of men can do" (1637, Part V), whatever *has* such ability evidently thinks. Turing embodies this conversation criterion in a would-be experimental test of machine intelligence; in effect, a "blind" interview. Not knowing which is which, a human interviewer addresses questions, on the one hand, to a computer, and, on the other, to a human being. If, after a decent interval, the questioner is unable to tell which interviewee is the computer on the basis of their answers, then, Turing concludes, we would be well warranted in concluding that the computer, like the person, actually thinks. Restricting himself to the epistemological claim that under the envisaged circumstances attribution of thought to the computer is *warranted*, Turing himself hazards no metaphysical guesses as to *what thought is* - proposing no definition or no conjecture as to the essential nature thereof. Nevertheless, his would-be experimental apparatus can be used to characterize the main competing metaphysical hypotheses here in terms their answers to the question of what else or what instead, if anything, is required to guarantee that intelligent-seeming behavior really is intelligent or evinces thought. Roughly speaking, we have four sorts of hypotheses here on offer. Behavioristic hypotheses deny that *anything* besides *acting* intelligent is required. Dualistic hypotheses hold that, besides (or instead of) intelligent-seeming behavior, thought requires having the right subjective conscious experiences. Identity theoretic hypotheses hold it to be essential that the intelligent-seeming performances proceed from the right underlying neurophysiological states. Functionalistic hypotheses hold that the intelligent-seeming behavior must be produced by the right procedures or computations.

The Chinese experiment, then, can be seen to take aim at Behaviorism and Functionalism as a would-be counterexample to both. Searle-in-the-room behaves as if he understands Chinese; yet doesn't understand: so, contrary to Behaviorism, acting (as-if) intelligent does not suffice for being so; something else is required. But, contrary to Functionalism this something else is not - or at least, not just - a matter of by what underlying procedures (or programming) the intelligent-seeming behavior is brought about: Searle-in-the-room, according to the thought-experiment, may be implementing whatever program you please, yet still be lacking the mental state (e.g., understanding Chinese) that his behavior would seem to evidence. Thus, Searle claims, Behaviorism and Functionalism are utterly refuted by this experiment; leaving dualistic and identity theoretic hypotheses in control of the field. Searle's own hypothesis of Biological Naturalism may be characterized sympathetically as an attempt to wed - or unsympathetically as an attempt to waffle between - the remaining dualistic and identity-theoretic alternatives.

[Back to Table of Contents](#)

Postscript

Debate over the Chinese room thought experiment - while generating considerable heat - has proven inconclusive. To the Chinese room's champions - as to Searle himself - the experiment and allied argument have often seemed so obviously cogent and decisively victorious that doubts professed by naysayers have seemed discreditable and disingenuous attempts to salvage "strong AI" at all costs. To the argument's detractors, on the other hand, the Chinese room has seemed more like "religious diatribe against AI, masquerading as a serious scientific argument" (Hofstadter 1980, p. 433) than a serious objection. Though I am with the masquerade party, a full dress criticism is, perhaps, out of place here (see Hauser 1993 and Hauser forthcoming). I offer, instead, the following (hopefully, not too tendentious) observations about the Chinese room and its neighborhood.

(1) Though Searle himself has consistently (since 1984) fronted the formal "derivation from axioms," general discussion continues to focus mainly on Searle's striking thought experiment. This is unfortunate, I think. Since intuitions about the experiment seem irremediably at loggerheads, perhaps closer attention to the *derivation* could shed some light on vagaries of the argument (see Hauser forthcoming).

(2) The Chinese room experiment, as Searle himself notices, is akin to "arbitrary realization" scenarios of the sort suggested first, perhaps, by Joseph Weizenbaum (1976, Ch. 2), who "shows in detail how to construct a computer using a roll of toilet paper and a pile of small stones" (Searle 1980a, p. 423). Such scenarios are also marshaled against Functionalism (and Behaviorism *en passant*) by others, perhaps most famously, by Ned Block (1978). Arbitrary realizations imagine would-be AI-programs to be implemented in outlandish ways: *collective implementations* (e.g., by the population of China coordinating their efforts via two-way radio communications), imagine programs implemented by groups; *Rube Goldberg implementations* (e.g., Searle's water pipes or Weizenbaum's toilet paper roll and stones), imagine programs implemented bizarrely, in "the wrong stuff." Such scenarios aim to

provoke intuitions that no such thing - no such collective or no such ridiculous contraption - could possibly be possessed of mental states. This, together with the premise - generally conceded by Functionalists - that programs might *well be so implemented*, yields the conclusion that computation, the "right programming" does not suffice for thought; the programming must be implemented in "the right stuff." Searle concludes similarly that what the Chinese room experiment shows is that "[w]hat matters about brain operations is not the formal shadow cast by the sequences of synapses but rather the actual properties of the synapses" (1980, p. 422), their "specific biochemistry" (1980, p. 424).

(3) Among those sympathetic to the Chinese room, it is mainly its negative claims - not Searle's positive doctrine - that garner assent. The positive doctrine - "biological naturalism," is either confused (waffling between identity theory and dualism) or else it *just is* identity theory or dualism.

(4) Since Searle argues against identity theory, on independent grounds, elsewhere (e.g., 1992, Ch. 5); and since he acknowledges the possibility that some "specific biochemistry" different than ours might suffice to produce conscious experiences and consequently intentionality (in Martians, say), and speaks unabashedly of "ontological subjectivity" (see, e.g., Searle 1992, p. 100); it seems most natural to construe Searle's positive doctrine as basically dualistic, specifically as a species of "property dualism" such as Thomas Nagel (1974, 1986) and Frank Jackson (1982) espouse. Nevertheless, Searle frequently and vigorously protests that he is not *any* sort of dualist. Perhaps he protests *too much*.

(5) If Searle's positive views *are* basically dualistic - as many believe - then the usual objections to dualism apply, other-minds troubles among them; so, the "other-minds" reply can hardly be said to "miss the point". Indeed, since the question of whether computers (can) think *just is* an other-minds question, if other minds questions "miss the point" it's hard to see how the Chinese room speaks to the issue of whether computers really (can) think at all.

(6) Confusion on the preceding point is fueled by Searle's seemingly equivocal use of the phrase "strong AI" to mean, on the one hand, *computers really do think*, and on the other hand, *thought is essentially just computation*. Even if thought is not essentially just computation, computers (even present-day ones), nevertheless, might really think. That their behavior seems to *evince* thought is why there is a problem about AI in the first place; and if Searle's argument merely discountenances theoretic or metaphysical identification of thought with computation, the behavioral evidence - and consequently Turing's point - remains unscathed. Since computers seem, on the face of things, to think, the conclusion that the essential nonidentity of thought with computation would seem to warrant is that whatever else thought essentially is, computers *have this too*; not, as Searle maintains, that computers' seeming thought-like performances are *bogus*. Alternately put, equivocation on "Strong AI" invalidates the would-be dilemma that Searle's initial contrast of "Strong AI" to "Weak AI" seems to pose:

Strong AI (they really do think) or Weak AI (it's just simulation).

Not Strong AI (by the Chinese room argument).

Therefore, Weak AI.

To show that *thought is not just computation* (what the Chinese room -- if it shows anything -- shows) is not to show that *computers' intelligent seeming performances are not real thought* (as the "strong" "weak" dichotomy suggests) .

[Back to Table of Contents](#)

Sources

- Block, Ned. 1978. "Troubles with Functionalism." In C. W. Savage, ed., *Perception and Cognition: Issues in the Foundations of Psychology, Minnesota Studies in the Philosophy of Science*, Vol. 9, 261-325. Minneapolis: University of Minnesota Press.
- Churchland, Paul, and Patricia Smith Churchland. 1990. "Could a machine think?" *Scientific American* 262(1, January): 32-39.
- Dennett, Daniel. 1980. "The milk of human intentionality." *Behavioral and Brain Sciences* 3: 429-430.
- Descartes, René. 1637. Discourse on method. Trans. John Cottingham, Robert Stoothoff and Dugald Murdoch. In *The philosophical writings of Descartes*, Vol. I, 109-151. New York: Cambridge University Press.
- Fodor, Jerry. 1980. "Searle on what only brains can do." *Behavioral and Brain Sciences* 3: 431-432.
- Hauser, Larry. 1993. [*Searle's Chinese Box: The Chinese Room Argument and Artificial Intelligence*](#). East Lansing, Michigan: Michigan State University (Doctoral Dissertation).
- Hauser, Larry. Forthcoming. [*Searle's Chinese Box: Debunking the Chinese Room Argument*](#). *Minds and Machines*, forthcoming.
- Jackson, Frank. 1982. "Epiphenomenal qualia." *Philosophical Quarterly* 32:127-136.
- Nagel, Thomas. 1974. What is it like to be a bat? *Philosophical Review* 83:435-450.
- Nagel, Thomas. 1986. *The View from Nowhere*. Oxford: Oxford University Press.
- Schank, Roger C., and Robert P. Abelson. 1977. *Scripts, Plans, Goals, and Understanding*. Hillsdale, NJ: Lawrence Erlbaum Press.
- Searle, John. 1980a. "Minds, Brains, and Programs." *Behavioral and Brain Sciences* 3, 417-424.
- Searle, John. 1980b. "Intrinsic Intentionality." *Behavioral and Brain Sciences* 3: 450-456.
- Searle, John. 1984. *Minds, Brains, and Science*. Cambridge: Harvard University Press.

- Searle, John. 1989. "Reply to Jacquette." *Philosophy and Phenomenological Research* XLIX: 701-708.
- Searle, John. 1990. "Is the Brain's Mind a Computer Program?" *Scientific American* 262: 26-31.
- Searle, John. 1992. *The Rediscovery of the Mind*, Cambridge, MA: MIT Press.
- Turing, Alan. 1950. "Computing Machinery and Intelligence." *Mind* LIX: 433-460.
- Weizenbaum, Joseph. 1976. *Computer Power and Human Reason*. San Francisco: W. H. Freeman.

Author Information:

Larry Hauser

Email: hauser@alma.edu

Alma College

Alma, MI

The Internet Encyclopedia of Philosophy

© 2001