Teng, N. Y. (2000). A cognitive analysis of the Chinese room argument.

Philosophical Psychology, 13.3, 313-324

A Cognitive Analysis of the Chinese Room Argument

Abstract

Searle's Chinese room argument is analyzed from a cognitive point of view. The analysis is based on a newly developed model of conceptual integration, the many space model proposed by Fauconnier and Turner. The main point of the analysis is that the central inference constructed in the Chinese room scenario is a result of a dynamic, cognitive activity of conceptual blending, with metaphor defining the basic features of the blending. Two important consequences follow: (1) Searle's recent contention that syntax is not intrinsic to physics turns out to be a slightly modified version of the old Chinese room argument. (2) The argument itself is still open to debate. It's persuasive but not conclusive, and at bottom it is a topological mismatch in the metaphoric conceptual integration that is responsible for the non-conclusive character of the Chinese room argument.

A Cognitive Analysis of the Chinese Room Argument

Searle's (1980a, 1980b, 1982, 1984, chap. 2, 1990) Chinese room argument has been one of the most celebrated philosophical arguments in the contemporary philosophy of mind. It is simple and elegant, and is the best-known and most-cited philosophical argument against Strong AI. (Strong AI is a claim that implementing the right program in any hardware is sufficient for having a mind.) In the following discussion, I attempt to analyze the argument from a cognitive viewpoint, arguing that the central inference constructed in the Chinese room scenario is a result of a dynamic, cognitive activity of conceptual blending, with metaphor defining the basic features of the blending. I will base my analysis on a newly developed model of conceptual integration, the many space model proposed by Fauconnier and Turner (see Fauconnier, 1994, 1996, 1997; Fauconnier & Turner, 1996, 1998; Turner, 1996a, 1996b; Turner & Fauconnier, 1995, in press). I explain what the many space model is, apply the model to an analysis of the Chinese room argument, and show how the composition, completion, and elaboration of the Chinese room scenario lead one to reconsider or refute Strong AI.

My analysis of the Chinese room argument has important consequences. With the Chinese room argument, Searle contends that Strong AI is false, because syntax alone is not sufficient for semantics. Recently, he (1992, chap. 9, 1997a, 1997b, chap. 1) argues that Strong AI is in fact incoherent, because syntax is not intrinsic to physics but is in the eye of the beholder. Searle seems to hold that this second argument is independent from the first. My analysis, nonetheless, shows that it is nothing but a new version of the old argument, which preserves most of its original structure. An appropriate realignment of parts of the Chinese room scenario, backed up with the conclusion from the Chinese room argument, can lead to the second argument. A second point clarified by my analysis regards the validity of Searle's argument itself. I will focus on the debate between Chalmers (1992, 1994a, 1994b, 1996, chap. 9) and Searle (1997b). An examination of why the Chinese room argument is persuasive but not conclusive is then proffered.

Before I proceed, a preliminary remark is in order. Philosophers, when they confront an argument and attempt to refute it, often provide a logical reconstruction of the argument with a conceptual analysis of the key concepts used in it. In this way, either a defect in the structure of the argument is located, or a premise refuted, or a key concept shown to be ill-defined. This is a valuable practice, and the debate about the Chinese room argument is no exception to it. But, as will be shown, the persuasive power of the Chinese room scenario is not derived from its logical structure. Rather, its construction draws on the cognitive resources we have in tackling problematic situations. An adequate and illuminating analysis of the scenario should be informed by what we have learned about human cognition. It is worth mentioning that though usually philosophy is called on to clarify a point in cognitive science, it is pertinent in the case of the Chinese room argument to use cognitive science to clarify the philosophical debate.

The Many Space Model

Fauconnier and Turner (1998) have shown how conceptual blends occur in a wide variety of cognitive phenomena and have developed an elaborate theory explaining the dynamic, supple, and active cognitive processing of conceptual integration. In this theory, mental spaces are on-line cognitive constructs, built up for purposes of local understanding and action. They are very partial assemblies containing elements, structured by frames, scripts, image schemas, or cognitive models without our recognizing them. They are interconnected, and can be modified and elaborated as thought and discourse unfold. In blending, mental spaces are configured, forming a network of mental spaces, typically with two spaces as inputs, sharing an abstract structure that constitutes another space, the generic space, which indicates the counterpart connections between the two input spaces. The counterpart connections guide the blending, in that the elements in the input spaces that are counterparts to each other are projected into a fourth space, the blend, and are fit together through the mechanism of composition. The initial composition is very partial, but it provides a working space for further composition. Completion is the mechanism that brings additional structure to the blend. And elaboration develops the blend through imaginative mental simulation according to the principles and internal logic of the blend. New structures emerge from the elaboration of the blend, which is crucial to the performance of the cognitive task at hand.

Two further considerations need to be noted. The blend can recruit conventional metaphors to develop structure not provided by the inputs. The recruitment of metaphors to the blend is in general partial and selective, but it can provide novel and transforming insights into the knowledge used to build the input spaces. When metaphors are in use, one input space is typically called "the source input space" (sometimes "the source space," or simply "the source"), and another input space "the target input space" (sometimes "the target space," or simply "the target"). Fauconnier and Turner also discuss a number of optimality principles governing the construction of the blended space. One principle, the topology principle, will figure importantly in my discussion of Searle's Chinese room argument. The principle is about the connections between the inputs and the blend, and is formulated as follows: "For any input space and any element in that space projected into the blend, it is optimal for the relations of the element in the blend to match the relations of its counterpart." (Fauconnier & Turner, 1998, p.163) With the basic ideas in hand, let us now turn to an example to illustrate the points.

A commentary, "Time to end the Olympics," written by Tom Regan for Christian Science Monitor (1999, March 18) has the following passage:

"Athletes are a walking laboratory, and the Olympics have become a proving ground for scientists, chemists and unethical doctors," Dr. Robert Voy, the director of drug testing for the U.S. Olympic Committee (USOC) at the 1984 and '88 Games told Sports Illustrated in 1997. "The testers know that the [drug] gurus are smarter than they are. They know how to get in under the radar."

An array of conceptual blends are deployed in this passage: "a walking laboratory," "a proving ground for scientists, chemists and unethical doctors," "the [drug] guru," and "how to get in under the radar." Note, however, that deployment of conceptual blends is not unique to this passage. Instead, blends are pervasive, and whenever you start looking for them, you will find them in every aspect of our life. But for the present purposes let us just take one sample for analysis. In "walking laboratory," we have two different input spaces. One has the group of the Olympians as its element, structured internally in terms of the cognitive model we have about Olympics. The other has the laboratory as its element, structured internally in terms of the cognitive model we have about gadgets, devices, and various functioning components of experimental settings installed in a room for scientific experimentation or research. The group of the athletes and the laboratory are counterparts, sharing an abstract structure, which can be formulated as follows: x is an entity with many constituent parts functioning as an integrated whole. The generic information makes up the generic space, indicating the counterpart connections that guide our construction of the blended space. In blending, we project partial structure from the input spaces and compose that structure in the blend. The athletes and the laboratory are brought into the blend and fused. Note that the compositional task can be done by exploiting the conceptual resources that we use to create ways of understanding other things in terms of ourselves, that is, personification. To wit,

His theory explained to me the behavior of chickens raised in factories.

This fact argues against the standard theories.

Life has cheated me.

Inflation is eating up our profits.

His religion tells him that he cannot drink fine French wines.

The Michelson-Morley experiment gave birth to a new physical theory.

Cancer finally caught up with him. (Lakoff & Johnson, 1980, p. 33)

Personification is a metaphoric device that we use in composing the walking laboratory, with the space of laboratory as the source and the space of athletes as the target. Completion provides additional structure, in which the athletes correspond to the functioning parts assembled in the laboratory. Elaboration develops the blend through imaginative mental simulation according to principles and logic in the blend. In elaboration, our knowledge of the problems of drug use by top Olympians is imported into the blend. The athletes, especially those who carry the illegal substances in their bodies, become a walking laboratory that makes the Olympics into a proving ground for scientists, chemists, and unethical doctors to outsmart the testers. Notice that this construction has developed a new structure not provided by the inputs. The inferential pattern and the viewpoint embedded in the new structure can be mapped back onto the target space, and provide a new understanding of the rampant drug use by elite athletes at the Olympic Games.

"Walking laboratory" is a result of the cognitive construction just described. The blending and the use of the metaphoric device of personification are crucial to this construction. And the connections between the inputs and the blend have been kept to conform to the topology principle.

The Chinese Room Argument

Strong AI claims that implementing the right program is sufficient for having a mind. According to this view, cognition is computation, which is a purely syntactical set of operations, in the sense that the only features of the symbols that matter to the computation are the formal syntactical features. Notice that computation, so defined, is independent of the structure of the physical system that carries out the instructions of the computer program. Different machines with different architectures can implement the same program. To be sure, the system implementing the program has to be physically realized, but it is the implemented program that is constitutive of having a mind.

The Chinese room argument is an attempt to refute the thesis of Strong AI. It has a simple three-step structure:

1. Programs are entirely syntactical.
2. Minds have a semantics.

3. Syntax is not the same as, nor by itself sufficient for, semantics.

Therefore, programs are not minds. Q.E.D. (Searle, 1997b, p. 11)

Premises 1 and 2 are not in dispute. Proponents for Strong AI would grant them to Searle. The key point lies in premise 3. And the famous Chinese room scenario is constructed to corroborate premise 3, or at least to make it more plausible. Here is the scenario Searle recapped in his recent book The mystery of consciousness:

Imagine that you carry out the steps in a program for answering questions in a language you do not understand. I do not understand Chinese, so I imagine that I am locked in a room with a lot of boxes of Chinese symbols (the database). I get small bunches of Chinese symbols passed to me (questions in Chinese), and I look up in a rule book (the program) what I am supposed to do. I perform certain operations on the symbols in accordance with the rules (that is, I carry out the steps in the program) and give back small bunches of symbols (answers to the questions) to those outside the room. I am the computer implementing a program for answering questions in Chinese, but all the same I do not understand a word of Chinese. And this is the point: if I do not understand Chinese solely on the basis of implementing a computer program for understanding Chinese, then neither does any other digital computer solely on that basis, because no digital computer has anything I do not have. (1997b, p. 11; italics original)

Let me first say something about the overall structure of the Chinese room scenario. Remember that, according to Strong AI, implementing the right program is sufficient for having a mind, and it doesn't matter which physical systems carry out the instructions of the program as long as the program is implemented. Now put a person (a physical system) who does not understand Chinese in the Chinese room. Let the person carry out the instructions from a rule book as described in the above scenario. Presumably, the rule book contains only the instructions stipulating the procedures to be followed solely in terms of the syntactical features of the Chinese symbols. The person manipulates the symbols according to the right procedures; that is, the person is a computer implementing the right program. If Strong AI is correct, the person can understand Chinese via manipulating the symbols. But the person, Searle contends, does not understand Chinese at all. Therefore, Strong AI is false.

Viewed from the perspective of the many space model, the construction of the Chinese room scenario involves a dynamic, cognitive activity of conceptual blending, which exploits a number of cross-domain mappings. The generic space consists of the generic information: x is a physical system capable of manipulating the symbols solely on the basis of their syntactical features according to a definite procedure. The input spaces contain the more specific information about the physical systems: the source space has the person who understands English but does not understand Chinese; the target space has the computer which runs on the right program for manipulating the Chinese symbols. The generic space indicates the counterpart connection between the two input spaces: the person corresponds to the computer. The blend exploits the counterpart connection between the input spaces and makes use of a number of metaphorical projections to develop structure not provided by the inputs. Inferences, arguments, and ideas are developed in the blend by three mechanisms: composition, completion, and elaboration.

Composition. The counterparts in the input spaces are mapped onto the same entity in the blend. The mapping is fit for the present purposes, and is sanctioned by the Machine As Person metaphor, which shows up in a wide variety of expressions in our everyday language:

This vaccum cleaner can pick up the tiniest grains of sand. My new Lexus can choose the best route to the ski resort and take me there in the shortest possible time. My microwave can bake a cake in six minutes. My car refuses to start. This vending machine is a bit recalcitrant; I'd better give it a kick. (Lakoff & Johnson, 1999, p. 263) This metaphor allows us to conceptualize the computer in the target as a person via the following metaphorical entailment:

A Machine Is A Person.

A computer is a machine.

Therefore, A Computer Is A Person.

The mapping of the two counterparts onto the same entity in the blend is based on the Computer As Person metaphor. The minimal composition provides a working space for further composition. (See Lakoff & Johnson, 1999, chap. 12 for an elaborate account of the conventional metaphors used in the Chinese room argument.)

Completion. With the Computer As Person metaphor, additional structure can be provided by completion. The person who understands English but does not understand Chinese is the leading character in our construction of the blended space. But the person in the blend is also the computer implementing the right program for answering questions in Chinese. The blend combines the information from the source and the target, which allows us to construct a scenario in which the person is situated in a room (the outer shell of the computer), equipped with boxes of Chinese symbols (the database) and a rule book written in English (the computer program), and is able to manipulate the Chinese symbols according to the rules (implementing the computer program). Questions handed in count as inputs, and answers handed out count as outputs.

Elaboration. We can now run the blend through imaginative mental simulation according to the principles and internal logic of the blend. The person gets small bunches of symbols, and in response manipulates the symbols according to the rules and hands back more small bunches of symbols. The process can be repeated and applied to different situations as long as the script is preserved. The purpose of the mental exercises is to see whether or not the person in the Chinese room can come to understand Chinese via manipulating Chinese symbols according to the syntactical rules. We know from our everyday situation that following syntactical rules is not sufficient for understanding a language, and that we can follow the syntactical rules without understanding the symbols on which we operate. Importing this knowledge into the blend, we reach the conclusion that the person does not understand Chinese in carrying out the steps in the program. Note that the blend remains hooked up to the input spaces. So the inference developed in the blend can be mapped back onto the input spaces. In this case, we are interested in the target. Projecting the inference from the blend to the target leads us to the conclusion that the computer does not understand Chinese solely on the basis of implementing the right program. The result can be generalized: Implementing the right program is not sufficient for having a mind.

Syntax Is Not Intrinsic to Physics

With the Chinese room argument, Searle argues that Strong AI is false. Recently, he (1992, chap. 9, 1997a, 1997b, chap. 1) argues that Strong AI is in fact incoherent. One main reason for this position is that syntax is not intrinsic to physics but is assigned relative to observers. Searle thinks that this point is so obvious that he should have seen it long ago. Indeed, given the Chinese room argument, plus a bit of cognitive realignment of parts of the Chinese room scenario, one is ready to reach the point.

Here is how it goes. We switch our attention from the symbolic level to the physical level. The generic space now consists of the generic information: x is a physical system capable of transforming one physical pattern into another solely on the basis of their physical features in accordance with a definite procedure. The input spaces contain the more specific information about the physical systems: the source space has the person who understands English but has no ideas what those physical patterns are; the target space has the computer which runs on the right program for transforming the physical patterns. The generic space indicates the counterpart connection between the two input spaces: the person corresponds to the computer. The blend exploits the counterpart connection between the input spaces and develops structure not provided by the inputs. The counterparts are mapped onto the same entity in the blend via the Computer As Person metaphor. The person in the blend is equipped with the structural components of the computer as before, but this time those physical components are interpreted solely in terms of their physical features. But since the person knows English, he can interpret the physical patterns (words of English) in the rule book as specifying the procedure for transforming one physical pattern (unbeknownst to him, the physical pattern can be interpreted as a sequence of Chinese symbols) into another.

The blend now can be elaborated as follows. The person receives a set of physical tokens that contain certain physical patterns. He transforms the physical patterns according to the rule book and returns some other physical tokens that contain other physical patterns. The process can be repeated and applied to different situations as long as the script is preserved. We know from our everyday situation that the person can operate on the physical patterns without knowing that those physical patterns are symbols of some language. Projecting the inference from the blend into the target leads us to the conclusion that the computer can operate on the physical patterns without knowing that those physical patterns are symbols of some language. This conclusion, however, has to be modified in view of the first run of the Chinese room scenario, where we establish that implementing the right program is not sufficient for having a mind. The computer could merely be a physical device without the right causal organization for having a mind. If so, we could not really say that the computer knows or does not know that the physical patterns are symbols of some language. Implementing the program could merely be a physical process of transforming one physical pattern into another. The very possibility that the physical patterns could merely be physical patterns suggests that syntax is not intrinsic to physics. This conclusion has the consequence that computation, algorithm, and program are not intrinsic to physical systems. They are assigned relative to observers and users. And without the assignment, there is simply no fact of the matter about brain states that would make them computational states.

The Debate between Chalmers and Searle

Chalmers (1994a, 1994b, 1996) argues that the main problem of the Chinese room argument is that the argument does not respect the crucial role of implementation. For a physical system to be capable of implementing a program, it has to have the right causal power built into it. Certainly, programs are purely syntactic, and no program by itself is sufficient for having a mind. But implementations of programs are concrete systems with causal dynamics. The Chinese room argument does not rule out the possibility that semantics can arise from implementing the right program. We may consider a program to be an abstract, systematic description of the dynamic arrangement of the implementing system. Any physical system with the right physical structure that makes it capable of implementing the right program is capable of having a mind.

Searle's (1997b) reply is chiefly as follows. An implemented program, of course, has to be embodied, in that there has to be a physical system that has causal power to realize the abstract procedure of transforming on pattern into another. But the point is that programs are purely syntactic and observer-relative, and thus have to be accounted for in terms of the operations of the physical systems that are capable of having a mind. Don't reverse the order of explanation. Moreover, we know independently that brain processes bring about mental states. And since programs are observer-relative, there is no such thing as programs involved in the operations of brains that give rise to the basic features of minds. The dynamic arrangement of the processes of transforming one pattern into another depends crucially on specific physical structures of the underlying systems, which, so far as we know, could be restricted to biological systems.

Two notions figure importantly in this exchange: system and implementation. Let us review the Chinese room scenario in the Chinese room argument with respect to these two notions. Consider first the system view. The person in the blend is the computer carrying out the steps in the program. This blending is based on the Computer As Person metaphor. There is a hitch in this blending, however. The computer in the target is projected to the person in the blend. Both are integrated systems. But the structural components of the computer (the program and the database) are projected to the structural components that are not constituent parts of the person (the rule book and the boxes of Chinese symbols). This violates the topology principle, which, you may recall, says that for any input space and any element in that space projected into the blend, it is optimal for the relations of the element in the blend to match the relations of its counterpart. So even if we accept that the metaphor is apt, and that the person in the blend does not understand Chinese in implementing the program, we are not obliged to project the inference to the target and conclude that the computer does not understand Chinese.

One way to fix the mismatch between the blend and the target is to let the Chinese room as a whole be the computing system carrying out the steps in the program. But that will not do, either. For the counterpart connection between the person in the source and the computer in the target would then be disrupted, and as a result the original space configuration is disintegrated and can no longer support the construction of the reasoning in the blend. Many opponents of the Chinese room argument argue that the Chinese room as whole is a computer, which can come to understand Chinese via manipulating Chinese symbols. However, it is highly counterintuitive to attribute any mental properties to the Chinese room, and Searle rightly rejects it. The problem, however, does not lie in whether or not the Chinese room understands Chinese. It is the disruption of the counterpart connection between the inputs that causes problems. Notice further that the original Chinese room scenario is constructed on the basis of the commonsense expectation that if you are the person in the room and operate on the symbols the way it describes, you will judge for yourself from the first-person point of view that you still do not understand Chinese via manipulating Chinese symbols. (If you already understand Chinese, replace Chinese with another language you do not understand and reconstruct the scenario accordingly.) For Searle, the first-person point of view is essential to one's understanding of a language, and indeed to one's understanding of anything at all. But the Chinese room, at bottom, lacks such a viewpoint. Therefore, it cannot judge for itself whether or not it understands Chinese.

Another way to fix the problem is to let the person in the room memorize the rules and the syntactical features of Chinese symbols, and implement the program internally. In practice, however, no one can really memorize the extremely enormous number of the rules and the symbols. To do so would require an extraordinary feat that no one can actually achieve. So we cannot directly import the knowledge of our everyday situation into the blend as we do in the first elaboration of the Chinese room scenario. Let us assume for the sake of argument that the person somehow internalizes the rules and the symbols, and implements the program internally. What can we say about this situation? One might argue that though the Chinese room scenario is not compelling, it is persuasive. We cannot directly import the knowledge of our everyday situation into the new blend, but the elaboration of the first run of the Chinese room scenario can serve as an anchoring point, with reference to which we may conjecture that the person who implements the program internally does not understand Chinese. After all, the scenario in the new blend is also an on-line dynamic construction, the processing of which is partial and is to be done on the basis of the ongoing episode of the way the person manipulates the symbols. The person's overall action consists of executing small stories: get a small bunch of symbols, look up in the rule book what he is supposed to do, perform the action in accordance with the rules, and give back another bunch of symbols. It seems quite natural that we may extrapolate the general shape of the imagined situation from what we know about the small stories, although there is always a risk in the extrapolation. The risk cannot be ignored, of course. But the extrapolation is so natural that one may still infer that the person who implements the program internally does not understand Chinese.

The above consideration leads us to a re-examination of the role of implementation. The on-line performance of the person described above seems to be merely the performance that conforms to the script of the small stories. But this one-level description of the performance cannot be empirically adequate, given our knowledge of how a biological system carries out its computational tasks. Indeed, the one-level description is quite misleading, and leads us to the wrong view that the task can be done just by moving the symbols around. There are many different levels of implementation, however; and the computational processes involved are massively parallel, with the high-dimensional vector of neuronal activation levels as the fundamental mode of representation and the vector-to-vector transformation as the fundamental mode of computation in the brain. Obviously, the one-level description is an oversimplification and cannot do justice to the rich and sophisticated causal organization of the implementing system situated in an environment (see Churchland & Churchland, 1990; see also Teng, 1999). It still remains a possibility that implementing the right program can give rise to the basic features of a mind. The following consideration may make this possibility into a live alternative. Take the brain's organization as an anchoring point. Abstract the brain's organization into a computational description, with the causal structure of the physical system mirroring the formal structure of the computation. If this can be done, we can have the right program for having a mind, in that any physical system implementing the program will have mental states as its emergent properties. Brains are such physical systems, so are any other physical systems that are capable of implementing the program. The key point is implementation, and it does not matter whether the implementation is carried out by a functioning brain or any other operating systems, as long as the right causal dynamics is secured. In this view, cognitive systems have their mental properties in virtue of their causal organization, which has to be studied scientifically. The risk of the extrapolation described above should be taken seriously, especially when the systems under investigation are so complicated and sophisticated.

The approach sketched above depends crucially on the assumption that the brain's causal organization can be abstracted into a computational description. Chalmers upholds it; Searle denies it. And, as has been shown, the Chinese room scenario does not provide us with sufficient cognitive resources to decide on which way to go. At bottom, it is the topological mismatch that is responsible for the non-conclusive character of the Chinese room argument.

Conclusion

The foregoing analysis establishes the following two points: (1) Searle's recent contention that syntax is not intrinsic to physics turns out to be a slightly modified version of the old Chinese room argument. (2) The argument itself is still open to debate. It's persuasive but not conclusive, and it is the topological mismatch that is responsible for its non-conclusive character. These two points, taken together, imply that Strong AI remains a contested issue, the solution of which is ultimately dependent on whether or not the brain's causal organization can be abstracted into a computational description. Our understanding of how the brain works matters deeply, and the philosophical debate about Strong AI can benefit from an analysis that is rooted in what we have learned about human cognition.

References

Chalmers, D. J. (1992). Subsymbolic computation and the Chinese room. In J. Dinsmore (Ed.), The symbolic and connectionist paradigms: Closing the gap (pp. 25-48). Hillsdale, NJ: Lawrence Erlbaum.

Chalmers, D. J. (1994a). On implementing a computation. Mind and Machines, 4, 391-402.

Chalmers, D. J. (1994b). A computational foundation for the study of cognition (PNP Technical Report, 94-03). St. Louis, MO: Washington University [Online]. Available from http://www.artsci.wustl.edu/~philos/pnp

Chalmers, D. J. (1996). The conscious mind: In search of a fundamental theory. Oxford: Oxford University Press.

Churchland, P. M. & Churchland, P. S. (1990). Could a machine think? Scientific American 262, 32-37.

Fauconnier, G. (1994). Mental spaces: Aspects of meaning construction in natural language. Cambridge: Cambridge University Press.

Fauconnier, G. (1996). Analogical Counterfactuals. In G. Fauconnier & E. Sweetser (Eds.), Spaces, worlds, and grammar (pp. 57-90). Chicago: University of Chicago Press.

Fauconnier, G. (1997). Mappings in thought and language. Cambridge: Cambridge University Press.

Fauconnier, G. & Turner, M. (1996). Blending as a central process of grammar. In a. Goldberg (Ed.), Conceptual structure, discourse, and language (pp. 113-129). Stanford: Center for the Study of Language and Information [distributed by Cambridge University Press].

Fauconnier, G. & Turner, M. (1998). Conceptual integration networks. Cognitive Science, 22, 133-187.

Lakoff, G. & Johnson, M. (1980). Metaphors we live by. Chicago: University of Chicago Press.

Lakoff, G. & Johnson, M. (1999). Philosophy in the flesh: The embodied mind and its challenge to western thought. New York: Basic Books.

Searle, J. (1980a). Minds, brains, and programs. Behavioral and Brain Sciences, 3, 417-424.

Searle, J. (1980b). Intrinsic intentionality: Reply to criticisms of minds, brains, and programs. Behavioral and Brain Sciences, 3, 450-457.

Searle, J. (1982). The Chinese room revisited: Response to further commentaries on minds, brains, and programs. Behavioral and Brain Sciences, 5, 345-348.

Searle, J. (1984). Mind, brain and science. Cambridge, MA: Harvard University Press.

Searle, J. (1990). Is the brain's mind a computer program? Scientific American, 262, 20-25.

Searle, J. (1992). The rediscovery of the mind. Cambridge, MA: MIT Press.

Searle, J. (1997a). The explanation of cognition. Philosophy: The Journal of Royal Institute of Philosophy, Supplement, 42, 103-126.

Searle, J. (1997b). The mystery of consciousness. New York: New York Review of Books.

Teng, N. Y. (1999). The language of thought and the embodied nature of language use. Philosophical Studies, 94, 237-251.

Turner, M. (1996a). Conceptual blending and counterfactual argument in the social and behavioral sciences. In P. Telock & A. Belkin (Eds.), Counterfactual thought experiments in world politics. Princeton: Princeton University Press.

Turner, M. (1996b). The literary mind: The origins of thought and language. New York: Oxford University Press.

Turner, M. & Fauconnier, G. (1995). Conceptual integration and formal expression. Metaphor and Symbolic Activity, 10, 183-204.

Turner, M. and Fauconnier, G. (in press). Metaphor, metonomy, and binding. In A. Barcelona (Ed.), Metonomy and metaphor. Berlin: Mouton de Gruyter [Online]. Available from

http://www.wam.umd.edu/~mturn

Acknowledgment

I thank Cees van Leeuwen for valuable advice and assistance in revising the manuscript.