INTENTIONALITY, UNDERSTANDING, AND SYMBOL GROUNDING: SEARLE'S *CHINESE ROOM ARGUMENT* AND THE LIMITS OF COMPUTATIONALISM

MAYANK BORA

The computational theory of mind for a few decades has been the ruling paradigm in the Analytic Philosophy of Mind and Cognitive Sciences, and probably one of the most significant facets of the "cognitive revolution". While it is believed that this is for a great part for good reasons, there are certain aspects, perhaps the most significant ones, of Mind and Cognition that may be far from being well understood within this paradigm. Searle's Chinese Room Argument (henceforth CRA) may be seen as providing reasons why this might be the case when it comes to Intentionality and Consciousness.

The Chinese Room Argument:

Searle presented 'the Chinese Room argument', which he considered to have thoroughly refuted the 'strong claim' of artificial intelligence. The strong claim of artificial intelligence or 'Strong AI' as defined by Searle is that "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" (Searle 1980), as against 'Weak AI' which according to Searle stands for the view that "the principal value of the computer in the study of the mind is that it gives us a very powerful tool. For example, it enables us to formulate and test hypotheses in a more rigorous and precise fashion." (*Ibid*)

The thought experiment (as presented in Searle 1990) calls one to imagine a monolingual English speaker (one who has no knowledge of what Chinese symbols mean) in a closed room, given some batches of Chinese symbols. The room is to have a slit from where certain symbols of Chinese writing can be passed in by Chinese speakers outside and passed out by the person inside. A rulebook in English language containing instructions on how to match Chinese symbols with one another is also supplied to the person in the room. The idea is that the rulebook is designed such that if the people outside pass in question in Chinese the person inside can 'answer' them appropriately by forming strings of Chinese symbols on the basis of the instructions available to him and passing them out. The Chinese speakers outside would naturally ascribe understanding of Chinese to the person inside. To them they would be having a meaningful conversation in Chinese with a Chinese speaker (i.e. someone who understands Chinese). Yet, the monolingual English speaker has no idea of what has been going on. He understands not a bit of Chinese and is hopelessly ignorant of what he has been asked and what he has 'answered'.[1]

Searle directed the CRA against Strong AI. The Strong AI contention, as Searle introduces us to it, is that if a computer can be programmed in such a manner that its behavior in the limited sense under observation is indistinguishable from that of a being to which understanding is ascribed, in other words that it passes the Turing Test (Turing 1950), then it also understands. The Chinese room argument very neatly shows that the man in the Chinese room would pass the Turing test but does not, in fact, understand Chinese; thus, the Strong AI claim is false. However, the CRA has ramifications for more than just its intended target. The computational theory of mind or computationalism may be summarized as: Mental states are just implementations of (the right) computer program(s). (Otherwise put: Mental states are just computational states). (Harnad 2001)

Understood this way computationalism is also under attack by the CRA. If a mental state is just an implementation of the right computer program, then the Chinese room, presumably, running the right program for the understanding of Chinese, must result in the appropriate understanding. But, that it does not seem to. Thus, implementing the right sort of computer program is, evidently, not enough to have the relevant mental states; computation is not enough for cognition.

Content, Cognition and Symbol Grounding

While, in (Searle 1990) the CRA was used to argue that computation is not enough for cognition, in (Searle 1980) the argument was centered around the notion of intentionality. Just as it can be said on the basis of the CRA that implementing the

---

[1] The original presentations of the argument do differ in details with the one provided here, but to the best of my knowledge the above account is true to the spirit of the argument. See (Searle 1980, 1990).

relevant program does not seem to give rise to the understanding of Chinese, it can also be said, given that the Chinese conversation was occasionally about (or, intentional towards) some objects or states of affairs in the world, that implementing the program also did not give rise to any intentional states directed towards these objects or states of affairs that an actual Chinese speaker would have had. For example, if part of the Chinese conversation included the people outside telling the person inside that there has been an earthquake in China, a fact otherwise unknown to the person inside, then an actual Chinese speaker would be expected to now have the belief regarding such an earthquake, but the man running the program in the CRA would still be oblivious of it.

Intentionality, the property of mental states of being directed towards objects or states of affairs in the world, is a significant property of cognition, but not co-extensional with cognition or consciousness. As Searle (1983) points out, conscious cognitive states like feeling pain etc are not intentional towards an object, or state of affairs in the world. Searle (ibid) also points out that some intentional states like beliefs are not always conscious. It is no coincidence that the CRA apparatus can be used to argue with respect to both intentionality (as in Searle 1980) and cognition in general (as in Searle 1990).

When it comes to cognition, there are two significant aspects of concern, which loosely track the consciousness/intentionality distinction. There is the matter of the raw feel of the experience/conscious mental state; what is it for the experiencing being to have that experience/conscious mental state.[2] We may term this the *qualitative* aspect of cognition. On the other hand, as we have already discussed, intentional states have intentional content. This we may term the *content-related* aspect of cognition. In the CRA the man inside is obviously capable of having cognitive/intentional states. In fact that is what allows us to get this 'look inside of a computer', not otherwise possible, and see what states do or do not obtain. It is also not that the man inside lacks any cognitive states that correspond to the Chinese symbols. While he might not be aware that the symbols are Chinese, he has to be able

[2] Note that not all mental states need be conscious. As such this does not apply to cognition in general, but many if not most cognitive states do have this experiential/conscious aspect.

to perceive the distinct shapes of the symbols to manipulate them according to the program, and thus must have some cognitive states pertaining to the symbols (states where the intentional objects are the symbols themselves). It is not, thus, a lack of the qualitative aspects of cognition that we witness in the CRA. Instead, it is that none of the cognitive states of the man inside have much of the sort of content that a Chinese speaker in the same position would. The Chinese speaker would not only have states intentional towards the symbols but also states that are intentional towards what the symbols symbolize, or 'stand for'.[2] In cases of linguistic cognition such as in the CRA, the relevant mental states get the 'stand for' (semantic) content via the knowledge of the meaning of the symbols. A Chinese speaker other than having the perceptual states regarding the shapes of the symbols also have states with the content that the Chinese symbols being used 'stand for'.[3] Precisely because the man inside does not know what these symbols stand for, he has no cognitive or intentional states that correspond with their meanings.

In short, Searle's argument is based on the Chinese symbols, which the man inside or more generally, a computer, operates on not being meaningful for the computer. In technical terms there is a symbol string that is put as input into the computer which after certain formal operations outputs another symbol string. Searle's argument is centered on that these inputs and outputs, including any mediate arrangement of these or equivalent symbols, are not in any sense interpreted by the computer but only by those outside who attach some interpretation to the symbols. According to Searle,

> The fact that the programmer and the interpreter of the computer output use the symbols to stand for objects in the world is totally beyond the scope of the computer. ... if you type into the computer '2 plus 2 equals?' it will type out '-4.' But it has no idea that '-4' means 4 or that it means anything at all." (Searle 1980)

---

[2] Thus, we can see that the design of the CRA is such that it cannot show us that by purely manipulating symbols according to a program whether consciousness can arise or not, since the man inside is already (capable of) having conscious states. (Not that we are trying to show or even believe that that can happen, but just for proper exegesis). What it does show is that whatever conscious states the man may have lack the relevant content.

[3] We disregard for now the possible limitations of vocabulary. We suppose that all the conversation uses basic vocabulary an average Chinese speaker would know.

Thus, Searle seems to be pointing out that the symbols that a computer operates on do not have any grounding internal to the computer system but, that the semantics of the symbols for the computer are actually derived from the semantics of these symbols for the programmers and the interpreters. As such, what we see here is the notorious problem in computer science called the symbol-grounding problem. According to Searle,

> (T)he point is not that it lacks some second-order information about the interpretation of its first-order symbols, but rather that its first-order symbols don't have any interpretations as far as the computer is concerned. All the computer has is more symbols. (Searle 1980)

While here Searle is trying to show what goes wrong for the Strong AI/Computationalism claim in the Chinese room, this quote is a very good statement of the symbol grounding problem too. The problem is precisely about making the interpretation of the symbols available internally to the system. In the Chinese room the Chinese symbols would have been considered grounded if the English speaker could interpret them. One way of doing it would have been to provide the English speaker with a Chinese to English dictionary. The person could have used the dictionary to translate the Chinese symbols into their English equivalents thereby understanding what they mean. However, such a scenario is not achievable in an actual computer.

For a computer all symbols, whatever they may mean to an interpreter, are like what the Chinese symbols are to the monolingual English speaker; meaningless. As there is no language natural for the computer, to a computer there is no symbol or system of symbols intrinsically meaningful, in terms of which other symbols can be defined and thence made internally interpretable for it. The meta-language/object-language approach that may work fine for a constructed formal logical system has no hope of working for a computer as the meta-language symbols themselves would not have any interpretation for the computer. Searle's Chinese room argument makes use of this fact about computers to conclude that computers, even if they have the right sort of program and the right sorts of inputs and outputs, do not have mental states of understanding, or in Searle's terms they do not have intentional states. Intentionality is not achieved by symbol-manipulation, or to once again use Searle's terminology, "Syntax by itself is neither constitutive of nor sufficient for semantics." (Searle 1990)

Syntax and Semantics

The crux of the story is that symbol manipulation cannot establish (linguistic) understanding, since the lack of symbol grounding means the meanings of the symbols and therefore understanding of the conversation is not available. However, there are two ways in which symbols can be meaningful/significant and thus grounded. One way is to know what the symbols mean/stand for. Such knowledge of symbols is what tells us what objects the symbols, and what states of affairs the sentences built out of those symbols, stand for and thereby results in the relevant intentional states. The other way symbols can be significant/meaningful is their structural or 'grammatical significance'. And grammatical significance arises out of merely symbol manipulations and the properties of the symbols that are required for the symbol manipulations and must therefore be given as part of the 'shape'[4] of the symbols. As such syntax is enough for grounding symbols in the latter way and thereby the corresponding understanding can be made available by merely symbol manipulations.

Only certain combinations of symbols, or symbol strings, are interpretable. That is only a select few combinations of symbols are acceptable and then interpreted as to the ‘stand for’ meanings of the symbols. In common parlance we know this as the grammar of a language and the grammaticality of sentence. The grammaticality of a sentence is iteslf a matter of linguistic significance.[5]

Alec Marantz, as quoted in (Mukherji 2010: 78), shows that one can start with English words, put them in a form true to Japanese word order and the resulting “word salad” (*Man the book a women those to given has) would not be interpreted as an English sentence. Whereas, a skeletal sentence with a recognizable verb form of an unknown verb (X bobbed Y) is readily interpreted as an English sentence. 'X' and 'Y' being unknown symbols do not impign any grammatically significant properties on the string, and those of the symbol recognized as a verb due to is morphological

[4] The shape of a symbol must be understood as the totality of the properties of the symbol that the relevant symbol manipulations are sensitive too. In other words, (the totality of) the properties of the symbol that the symbol manipulations require to manipulate on them.

[5] In fact, the grammar of a language may be the only theoretically salient conception of what counts as a particular natural language.

suffix take over, forcing Noun Phrase interpretations on 'X' and 'Y'. Chomsky's example highlighting the grammatical significance of sentences, as the grammaticality of sentences may be called; (Colorless green ideas sleep furiously) is legendary. It shows how a sentence, that is obviously not meaningful in the way a sentence like 'Good faithful dogs guard incessantly' is meaningful, is yet not mere noise but comes across as significant or meaningful in a way, which is relevant to linguistic/syntactic theory.

One may also think of grammatical significance in terms of an analogy with music. In music only certain combinations of the various notes are musical. The others are discordant and come across as noise rather than pieces of music. As such, the musical pieces have a sense of significance about them that is lacking in the discordant strings. Similarly, grammatical sentences are grammatically significant even though at times the stand for meanings of the symbols do not result in any stand for interpretation for the sentence as in Chomsky's example.

While in communication there is high measure of looseness in terms of grammaticality of sentences, it cannot be completely disregarded as the word salad example above bears witness to.[6] Thus, in the Chinese room if the programmers have to make the responses of the man inside indistinguishable from those of a natural Chinese speaker one thing they would have to do is make the responses of the man grammatically correct for the most part. For this they will have to program a grammatical subroutine and include that in the rulebook. The grammatical subroutine will tell the man inside in which order to place the symbols selected for the response. The grammatical subroutine will also first receive the inputs and according to the grammatical structure discern what state of affairs is being referred to. (Did the man give the book to those women, or did the book give the man to those women, etc.) As the grammar of a language is sensitive to certain categorical and semantic properties of the lexical items which the symbols represent, some way would have to be found so as to attach some representations of such lexical properties to the symbols. One

[6] For now, for the sake of argument let us just consider that a string is grammatical if and only if it is interpretable. This does not make a difference since, we may just argue taking the notion of interpretability as central instead of grammaticality.

way to do it would be to classify symbols according to their categorical properties by putting symbols representing the same grammatical categories in the same boxes. The man inside would not have any knowledge of the basis of these classifications. He would just have instructions to the effect, “pick a certain symbol from Box A”, “pick a certain symbol from Box B”, “place the symbol from Box B before the symbol from Box A”, etc. The input-output strings to be grammatical would need to agree to specific patterns, as determined by Chinese grammar, so that out of the many possible arrangements of a set of symbols, only a few would ever be selected.

Now, while the program is programmed to not produce ungrammatical strings and therefore barring a failure on the part of the man inside to carry some step of the computation properly, given that the program itself has no faults, there would be no output strings which would be ungrammatical, it might be the case that a person outside makes a mistake and sends in a string that happens to be ungrammatical. Since the program would be required to compute on the inputs in order to determine the right outputs, and will have to mimic the inability of a Chinese Speaker of not being able to understand Chinese word salads like the English word salad above, the program would have to be such that for such a string un-interpretable to a Chinese speaker is also un-interpretable for the program. While it could be that in such cases the program just asks the man inside to output something to the effect “Sorry! I didn't get that” directly, the program could also be designed so that it simply crashes and outputs nothing. Then the rulebook might tell the man that whenever the program crashes and selects no output then to send a particular string as output, a string to the same effect of expressing the failure to understand as the one above. Then again it might not and the lack of output would to the people outside be similar to the puzzled silence which is a common response in such situations in real life. Now, consider the scenario where the input to the Chinese room is mixed with respect to grammatical and ungrammatical strings. Whenever the input is ungrammatical the computation for the grammatical subroutine will crash, and the person will be able to note that and send the output expressing failure to understand. Whenever the input is grammatical the computation for the grammatical subroutine will sail through, which again the man will be again be able to note. As such the person will be able to tell whether the input string is grammatical or not. In effect, some states of understanding pertaining to the Chinese symbols would be available to the man inside. These intuitions are strengthened if we consider variation of the CRA where it is not the man inside the Chinese room, but the reverse, i.e. when the man is supposed to have internalized all the CRA apparatus. In such a case whenever the man is presented with a string of Chinese symbols, he would in his head carry out the grammatical subroutine and if the string is proper the computation will go through and the man will know now he can proceed to choose an answer for it. If not the man will not be able to carry out the grammatical subroutine properly and will know that he can not go on computing with the string. Other than the fact that in actual cases the computations are unconscious processes this internalized CRA picture is quite akin to the what linguistic/syntactic theory tells us about the syntactical aspects of our own language use.

Even if the level of understanding that can be achieved in the above case would be miniscule, it is not a trivial matter. For one thing, there are some states of understanding pertaining to the Chinese symbols that the man inside will be have purely on the basis of symbol manipulation and that are not merely perceptual states directed towards the symbols but would qualify as linguistic cognition, not cognition of the stand for meaning or referents of the symbols but of their structural significance. As such, even if Searle's use of the CRA with respect to intentionality would not be undermined his case against cognition in general would be. But, even for intentionality grammaticality is significant. In the word salad example the string of symbols can not correspond to any state of affairs. Only if the string was grammatical in English would a(n) (English speaking) person be able to discern the relevant state of affairs, and have a though intentional towards them. Also, changing the position of the terms and “man” and “book” in a grammatical English form of the string[7] would signal distinct states of affairs. Thus, even though grammaticality by itself does not determine intentionality, the grammatical structure of a sentence is very significant in computing the right states of affairs that correspond to it. I believe one thing has been shown. The CRA cannot make a case against symbol manipuation giving rise to cognition where the symbol manipulations are being used to ground

---

[7] Which happens to be: “The man has a given a book to those women.”

grammatical understanding of linguistic strings or, by analogy, structural understanding of non-linguistic strings where such an understanding makes sense.

Computationalism and Snowflakes

At this point of time a computationalist might want to say well in cases like these where structural significance is in question, the mental states consist of just the relevant computations. But, this I believe can lead to certain absurd conclusions. To see this consider an analogous scenario to the thought experiment above. The understanding emerging in the experiment above is the understanding of the structure of the output string. There could be another room where the symbols instead of representing Chinese words, represent other things, having outputs that also have notable structural properties.

Snowflakes do have very significant structural properties. They can have highly complex structures but almost always maintain a high degree of symmetry around their geometrical centre. But perhaps, the most consistent feature of snowflakes is that they always display a hexagonal structure and the symmetry is always six-fold. Their final structure is a function of the physical-chemical properties determining the crystal structure of ice particles on the basis of the nature of bonding between water molecules, and the values of the atmospheric variables mainly temperature and humidity during various points of time in the development of the snow-crystal. The snowflakes begin as tiny micro sized droplets of water which freeze and start gathering more and more water molecules by the condensation of water vapor directly into solid form. Thereon the atmospheric determinants result in the shape of the snowflake at any given time since the beginning of its formation.[8] (Libbrecht 2004/5)

[8] This is quite analogous to the Universal Grammar notion in linguistics. The inherent combinatorial tendencies of the water molecules dependent upon the physical-chemical properties of water molecules play the role of the Universal Grammar and thereon the atmospheric determinants result in the particular structure of any given snowflake just as the Universal Grammar and the environmental stimuli result in the particular structural/grammatical properties of the natural language acquired with the categorial properties of the words then deciding the grammatical structure of individual sentences.

Now imagine that, the program being run in the room is not one for understanding Chinese with the symbols being Chinese symbols, but instead a computer model simulating the process of snowflake formation with the symbols representing water molecules. If the simulation is accurate then the program will never output anything that does not have a hexagonal structure. Also as Libbrecht says, “The exact shape of each of the six arms reflects the history of the crystal’s growth,” (I-vi) given a complete understanding of the physics involved, we can design the model to not only simulate the formation of snowflakes but also be run in reverse to take a representation of a given snowflake and trace back its plausible histories. Designed this way, if the model is presented with an anomalous snowflake representation, say one displaying an octagonal structure, the model will not be able to select any plausible histories as it will not be able to recognize the input string as valid and the computation will crash. Similar to as in the modified CRA above with the grammatical subroutine, here too the man inside will be able to tell between (the representation of a) proper and (that of) an anomalous snowflake.

Given this the computationalist would have to say that even the understanding of snowflakes can be had by merely running a computer program, a statement that is in fact not in itself problematic for the computationalist. But it becomes, or ought to become, problematic when we consider what is or can be considered as a computer performing some computation. In fact, any thing, or process any aspects of which can be understood in terms of symbols, inputs and outputs can be computationally described and anything that can be computationally described can be thought of as a computer. It can at the very least be thought of as computing the function that has been used to describe it. Something as usual as a walk can be seen as a computation. All that needs to be done is to imagine a 2-D grid on the ground that a person walks. The man’s positions on the grid can be taken down in terms of a set of (x, y) values. There would be a mathematical function that captures this set of values (in fact, there would be infinite such functions). The man’s walk can then be said to compute this function, with his x-positions as inputs and his y-positions as outputs. Thus anything can be seen as a computational system.[9]

Given this consideration what could better computationally simulate the process of snowflake formation than the process of snowflake formation itself seen computationally. All that is needed is to attach the notion of symbols appropriately, i.e. taking any water molecule as a token symbolic representation of the symbol type corresponding to water molecules. Similarly, any individual snowflake would be the token representation of the string[10] type corresponding to the class of snowflakes physically identical to that one. Now we could even use the natural process of snowflake formation as a computer implementing a computational simulation of itself, akin to the one in the thought experiment above. When we have to see what kind of a snowflake would be formed given a history certain atmospheric values, we simply locate natural conditions where that history obtains.[11] Similarly, given a snowflake type we calculate for it a plausible history by simply locating a corresponding snowflake and tracing back its atmospheric history, assuming for the sake of argument that we are able to maintain a record of that. And, if we were to try and find a plausible history for an imagined octagonal snowflake we would not be able to assign it a corresponding string type and thus the computation crashes and no history is obtained. Of course, most of these things are impractical and some even impossible to actually perform. But, the point is not that we actually be able to use the process of snowflake formation as its own computer simulation but only that we be able to see it as implementing one. The fact that we can see how in principle it can be used to that effect tells us that we can, in fact, see it as implementing its own computer simulation.

---

[9] Note that I am not saying that every thing is every computer but only that everything is some computer, even if a computer computing nothing more than its own computational description. As such this may escape Block's (1995) criticism of Searle's idea, presented in (*Ibid*) that every thing is every computer, since here we only consider that thing,/process as machine the computational description of which is to be computed. A wall is not really isomorphic to the computation 1+0=1, since not all aspects of the wall can be mapped on to it, but this obviously is not the case for a walk and a mathematical curve that describes it perfectly.

[10] A snowflake would be seen as a (3–D) string of water molecule symbols.

[11] Of course, this is easy only in principle but may be close to impossible in realit. But, we are only running a thought experiment here.

Any computationalist who thought that in cases where structural significance is in question computers can understand will have to admit that even the natural process of snowflake formation literally understands a snowflake from a non-snowflake, and thus has a mind. This I believe is a totally absurd thing to believe in. Searle (1980) says:

> If Strong AI is to be a branch of psychology, then it must be able to distinguish those systems that are genuinely mental from those that are not. It must be able to distinguish the principles on which the mind works from those on which nonmental systems work; otherwise it will offer us no explanations of what is specifically mental about the mental.

The same also goes for computationalism. What Searle had in mind above were machines, like thermostats and the normal computers. But, Searle's concern becomes much more relevant when even natural processes begin to be counted as mental. Human beings have a mind and so do dogs, cats, elephants and dolphins. Perhaps, all higher animals do. The boundary is not clearly drawn and developing understanding about cognition might make it clearer where it lies. It may even tell us some dubious cases are on the same side as us. But, natural processes forming snowflakes, ferntrees or quartz crystals are clear cases of non- mental and if you get a theory of cognition "that denies this point you have produced a counterexample to the theory and the theory is false." (Searle *Ibid*).[12]

Some Cursory Remarks on Symbols and Consciousness

One more thing needs to be considered about computation, and this might tell us why symbol-grounding might be, in both of its aspects problematic for computationalism. As we already saw, any thing or process can be seen as a computation. What is important is interpreting certain aspects, objects, as symbolic representations of something. In other words the notion of a symbol is paramount for computation. And while sometimes a symbol is interpreted as standing for something, sometimes all that is necessary is the shape of the symbol. But, whether seeing a symbol as a symbol for something or just as a symbol with a certain shape, what is necessary is *seeing* it as a symbol. Be it any (possible) computer, an intel chip or a

---

[12] Having said that, the computational theory of mind is not committed to saying that mental states just are computational states. That computation is necessary for (some) mental states though not sufficient are also a recognizable doctrine under the computational theory of mind.

natural process, there is no such thing as a computation without someone attaching the notion of a symbol to something. But, objects or physical processes do not see anything as a symbol, only cognizing beings like us can. Seeing a symbol as a symbol is itself a cognitive act, one that logically precedes the notion of computation. As such, computation simply cannot account for cognition in general. The computational theory of mind to make any sense of its main tool must itself rely on the notion of cognition. Symbol-grounding therefore might not only itself be a (somewhat) hard problem for the computational theory of mind but also make providing a full account of cognition problematic for it. Furthermore, if we consider what kind of cognition must seeing a symbol as a symbol be, in terms of the qualitative aspects of cognition, i.e. can merely unconscious mental states ground the notion of a symbol as a symbol, we might be inclined to answer in the negative.[13] In such a case, this problem attaches itself to the problem of consciousness and inherits its level of hardness. (cf. Chalmers 1995)

Bibliography

1. Block, N. (1995). The Mind as the Software of the Brain. In Simth, E. And Osherson, D. (eds.) *An Invitation to Cognitive Science, Vol. 3: Thinking*. Cambridge: MA, MIT Press, 377–426.
2. Chalmers, D. (1995). Facing Up to the Problem of Consciousness. *Journal of Consciousness Studies*, 2: 200–19.
3. Harnad, S. (2002). Mind, Machines and Searle 2: What's Wrong and Right About Searle's Chinese Room Argument. In: M. Bishop & J. Preston (eds.) *Views into the Chinese Room*. Oxford, Oxford University Press, 294–307.
4. Libbrecht, K. (2004/5). Snowflake Science: A rich Mix of Physics, Mathematics, Chemistry and Mystery. *American Educator*. 28: I-ii–I-viii.
5. Mukherji, N. (2010). *The Primacy of Grammar*. Cambridge: MA, MIT Press.
6. Searle, J. (1980). Mind, Brains and Programs. Behavioral and Brain Sciences*, 3: 417–57*.
7. Searle, J. (1983). *Intentionality: An Essay in the Philosophy of Mind*. Cambridge, CUP.
8. Searle, J. (1990). Is the Brain's Mind a Computer Program? *Scientific American*, 262: 26 -31.
9. Turing, A. (1950). Computing Machinery and Intelligence. *Mind*, 59: 433–60.

[13] I only speculate here about the possible relation between consciousness and meaning; space and time constraints do not allow a fuller inspection.