

Piotr Konderak

Maria Curie-Skłodowska University in Lublin

ON A COGNITIVE MODEL OF SEMIOSIS

Abstract. What is the class of possible semiotic systems? What kinds of systems could count as such systems? The human mind is naturally considered the prototypical semiotic system. During years of research in semiotics the class has been broadened to include i.e. living systems (Zlatev, 2002) like animals, or even plants (Krampen, 1992). It is suggested in the literature on artificial intelligence that artificial agents are typical examples of symbol-processing entities. It also seems that (at least some) semiotic processes are in fact cognitive processes. In consequence, it is natural to ask the question about the relation between semiotic studies and research on artificial cognitive systems within cognitive science. Consequently, my main question concerns the problem of inclusion or exclusion from the semiotic spectrum at least some artificial (computational) systems. I would like to consider some arguments against the possibility of artificial semiotic systems and I will try to repeal them. Then I will present an existing natural-language using agent of the SNePS system and interpret it in terms of Peircean theory of signs. I would like also to show that some properties of semiotic systems in Peircean sense could be also found in a discussed artificial system. Finally, I will have some remarks on the status of semiotics in general.

Keywords: semiotic systems, sign, Peirce, artificial cognitive agents, semantic networks, GLAIR/SNePS architecture.

1. The problem

Woody Allen in one of his movies puts into the mouth of one of his characters the following words: “I took a speed reading course and read *War and Peace* in twenty minutes. It involves Russia”. The phrase shows the ability of a human reader to interpret words, to pick up the arbitrary aspect of a situation. It is a typical example of the role of interpretation in our, human, ordinary use of language. A human user of language signs is capable of such interpretation – constantly and usually effortlessly. My question is whether artificially created agents (such as SNePS agents) are capable of such interpretation – on their own, independently of (or without) a user’s interventions. If they are to be considered as semiotic systems, they will have to do so not because a programmer provided it, but because of the

systems knowledge, aims, etc. In my opinion, this is the main problem with alleged artificial semiotic systems. It is argued that the answer is in the positive: at least some kinds of artificial cognitive systems (properly construed) could be semiotic systems. The answers in the negative (such as Fetzer's) are based on the excessive requirements of semiotic systems. I will examine these issues from the point of view of a theory of signs proposed by Peirce. The choice is a consequence of my approach to semiotic theory in general and will be briefly justified in the final part of the paper. Peircean theory imposes some restrictions on the class of possible semiotic systems; in consequence – in my opinion – only selected artificial systems may play the role of semiotic systems. Such systems should be capable of – among others – interpreting and re-interpreting signs, learning, belief revision in the context of new experiences, etc. SNePS/GLAIR agents capable of using natural language utterances and representing knowledge in the form of semantic networks meet these requirements (Shapiro & Bona, 2010; Shapiro & Rapaport, 1987).¹

2. Semiotics as a theory

To answer fully Fetzer's question whether artificial cognitive agents **are** semiotic systems, we need to answer another question: how do we treat semiotics as a theory? There are (at least) two approaches to scientific theories in the philosophy of science: we can interpret them in a realistic way or we can treat them instrumentally. Scientific realism states that the objects and properties it postulates exist in the world. We will ultimately discover these objects and properties. Realism in semiotics would state that signs and their properties exist in some world (physical, social, ideal) and that we are discovering signs during our research. The realist stance requires choice of one, true theory. Then we can solve the problem, whether artificial agents are *really* semiotic systems.

Instrumentalism, on the other hand, regards a theory as a useful tool helping to understand the world around us. We do not make any ontological commitments concerning the existence of objects or properties as far as a theory predicts and explains certain phenomena. In this approach, semiotic theories (Peircean, Saussurean, Greimasian etc.) would be the only possible descriptions of the world around us. Taking the instrumentalist attitude towards semiotics, we are interested in the usefulness of artificial semiotic systems in explaining semiotic phenomena. The approach as presented below is somehow neutral in reference to the distinction.²

2. Peircean semiotics

The fundamental notion of the theory is semiosis as a process. Semiosis according to Peirce is: “the doctrine of the essential and fundamental nature of all varieties of possible semioses (CP 5.484). Semiosis, according to his definition, entails an irreducible relation between a sign, its object, and its interpretant (CP 1.363, 8.331, 7.537). Simultaneously, semiosis is understood as an interaction between the three elements. Additionally, semiotics describes and analyses the structure of semiotic processes without any concern over any basis of material support that serves to bring about such processes. Such an approach – emphasizing the processual character of semiosis independently of the underlying physical medium (or particular implementation) – encourages functionalist explanations of the process.

Peirce conceives “a ‘Sign’ or ‘Representamen’ as a ‘First’ (S) which stands in a genuine triadic relation with a Second called its object (O) which is in the process of determining a ‘third’, called its Interpretant (I) which assumes the same triadic relation with that object” (CP 2.274)

Interpreting the above quotation in the framework of the Peircean semiotics, we can state that:

- the sign (Firstness): it is the book *War and Peace*,
- its object (Secondness): it may be the situation in Russia during the Napoleonic Wars (or the model of it in the mind of Leo Tolstoi),
- there is also a place for an interpretation (Thirdness): the respect in which the sign stands for its object (an aspect of the object, quite limited in this case): it involves Russia.

The main concern of the paper is whether artificial systems are capable of an interpretation (Thirdness), i.e. whether such systems could choose an aspect of the situation on their own, independently of a programmer/user and whether such interpretation could change together with the development of the system (in a process of consequent interpretations of interpretations, just as Peirce devised it).

3. Semiotic systems and artificial agents

James Fetzer in several papers argues that artificial agents (“machines”) could not be semiotic systems. A semiotic system in the sense of Fetzer is “a system that has the capacity to create or to utilize signs, where this capability may be either naturally produced or artificially contrived” (Fet-

zer, 1990, p. 31). Fetzer precises the definition, appealing to the Peircean notion of the sign: “a system for which something can stand for something else in some respect or other, where such a something (sign) can affect the (actual or potential) behavior of that system.” (Fetzer, 2001, p. 49) There are at least two problems to be considered: what does it mean to stand for something else *for the system* in the case of artificial semiotic systems? and, secondly, why is a particular respect (or aspect) taken into account? The first problem could be dubbed *the intrinsicity problem*, and the second the *grounding problem*. I will discuss both of them below.

An artificial agent (or system) is understood here as a system that is situated within an environment that perceives and acts in that environment to achieve some goals. The system is capable of learning and of using its knowledge (cf. Russell & Norvig, 2003). Imagine an artificial, natural language using a system (such as the system described below) that is “reading” the text of the *War and Peace* novel (i.e. takes it as input and process). It transforms the words of natural language into some symbols of its internal knowledge representation (KR) formalism (e.g. a semantic network as presented in Figure 1).




Figure 1. Simple semantic network representing the belief that Russia is a big country

According to Fetzer, language expressions (“Russia”, “country”) are the symbols in a Peircean sense when appropriate symbols of the formalism (B1 or M4) are meaningful *for* the artificial system, not for a user of the system.

Fetzer (2001, pp. 56–57) invokes a pocket calculator as an example of a symbolic, but not semiotic system: the numerals it displays are interpreted by users as numbers, the operator symbols are interpreted as operations on

numbers; signs are not meaningful to a calculator, but to the user of the device. In consequence, two possibilities should be distinguished (taking into account the semiotic perspective):

- sets of symbols are significant for users of machines (users of programs, designers of artificial agents)
- sets of symbols are significant for machines (user = machine)

We need a criterion that allows us to distinguish between symbols that stand for something else for a system and symbols that stand for something else for a user of a system.

Fetzer (2001, p. 50) proposes such a criterion: something is functioning as a sign for the system when:

- that something can affect the behavior of the system, and
- that standing must make a difference to (the actual or potential behavior of) that system.

Changes in behavior are an obvious criterion, but after longstanding debate on behaviorism in philosophy of mind (it is pointless to repeat the arguments against it) we should seek a more adequate solution. There are some proposals modifying the initial claim: directing our attention to the relation between a symbol and its meaning instead of the symbol alone: “what makes a system ‘semiotic’ thus becomes that the behavior of the system is causally affected by the presence of signals because it stands for something else iconically, indexically, or symbolically for the system” (Fetzer, 1997, p. 358) or suggesting that “some object has significance to the mind only as a result of the causal consequences invoked by the interpretation of a representation of the object” (Whobrey, 2001, p. 311).

As long as the only criterion is affecting behavior, such modifications are not sufficient. Fetzer tries to explain this in his subsequent papers in terms of dispositions to behave: “this difference can be specified in terms of the various ways that such a system would behave” (Fetzer, 2001, p. 50). In other words, changes in external behavior are sufficient but not necessary. The key criterion are changes in internal behavior that are both sufficient and necessary. It seems that Fetzer understands “internal behavior” as the dispositions to behave mentioned above. The notion of dispositions, again, has been severely criticized (cf. Putnam, 1963). I think that we can explain meaningfulness of the sign for the system in terms of changes in the external behavior and *the internal structure* of the agent.

Fetzer concludes that physical symbol systems (in the sense of Newell and Simon) do not qualify as semiotic systems: “...there is a crucial difference between sign-using systems and other systems that... do not qualify as semiotic systems. Ordinary digital computers are systems of this [latter]

kind, because the marks by means of which they operate are meaningless for systems of that kind. While they may be meaningful for users of those systems, they are not meaningful for those systems themselves.” (Fetzer, 2001, p. 118). In my opinion the conclusion is premature: some kinds of artificial cognitive systems (such as SNePS systems described below) fulfill Fetzer’s requirements: language signs cause changes in their knowledge base (restructure their semantic networks) and can cause changes in behavior.

Fetzer sees the explanation for this difference in “the absence of a grounding relationship between these marks and that for which they stand” (Fetzer, 2001, p. 118): a system is a semiotic system only if its behavior is brought about by virtue of the grounding relations (resemblance, cause-effect, or habitual association). In other words: the difference is further explained in terms of (absence of) a grounding relationship between these marks and that for which they stand. Fetzer states that artificial cognitive systems lack ground relations required from semiotic systems. There are, however, at least 2 notions of grounding: “The sign stands for something, its *object*. It stands for that object, not in all respects, but in reference to a sort of idea, which I have sometimes called the *ground* of the representamen” (CP 2.228). In other words, a ground is understood as the nature of the semiotic relation between a dynamic object and its representation in the system; ground conveys the respect by which the entity is represented to the interpretant. A ground relation transforms an aspect of dynamical object (independent of mind) into a representation called an immediate object.

In discussions on Artificial Intelligence, on the other hand, the grounding problem is the problem of connecting symbols with their meaning without mediation of any external interpretation (Harnad, 1994). And it seems that Fetzer has in mind this latter meaning. At the moment I would like to cover the latter sense of the term “grounding”; I will discuss the former below (see Section 6).

There are two possibilities: either an artificial system uses pre-encoded symbols (devised by a programmer) or it can create new symbols in answer to incoming stimuli from its environment. Let’s consider the first possibility, when representations are pre-encoded by a programmer. A good place to start are Putnam’s (1975) or Burge’s (1979) thought experiments. Imagine that an evil scientist created an exact physical replica of you. At the moment of creation the replica’s words do not refer, they are “inserted”, hardcoded into the replica’s mind. After some time, due to interactions with the environment, the words actually would refer. In other words, they may be post-grounded. If you agree that representations inserted into a system di-

rectly, not causally (via the world), would not be properly grounded, maybe you would allow such post-grounding? And if we allow post-grounding in this case, why not with artificial systems?

Although some AI systems have hardcoded representations in the form of symbols, full-fledged systems will ultimately ground their symbols due to interactions with the environment via perceptual/motor modules. Most contemporary advanced cognitive architectures (such as the GLAIR/SNePS architecture described below) use sensory input and learning and get most of its representations via such means. Would they be semiotic then?

4. Artificial cognitive agent

To pin down the above claims, I would like to investigate an example of the cognitive model of the language faculty implemented in the GLAIR/SNePS cognitive architecture. GLAIR (Grounded Layered Architecture with Integrated Reasoning) with SNePS (Semantic Network Processing System) (Shapiro & Bona, 2010) as a knowledge representation formalism is a multi-layered cognitive architecture for embodied agents acting in real, virtual, or simulated environments containing other agents. The architecture provides 3 layers (cf. Shapiro & Bona, 2010, pp. 308–309):

- Knowledge Layer (KL) containing beliefs of the system represented in the form of propositional semantic network³ knowledge representation formalism. It is the layer in which reasoning, planning, and act selection is performed. Knowledge layer consists of terms for every mental entity that an artificial agent has conceived of (including individual entities, categories, colors, shapes, and other properties of entities etc.)
- Perceptuo-Motor Layer (PML) grounding KL symbols in perceptual structures and actions. Each KL term for a perceivable entity (category, property) is grounded by aligning it with a PML description
- Sensori-Actuator Layer (SAL) contains controllers of the sensors and effectors of the robot (“embodied” artificial agent)

The simplest behavior cycle of a GLAIR/SNePS agent consists of 3 stages: (1) the agent perceives a natural language utterance (parses it and transforms into the form of the semantic network representation formalism), (2) the system analyses the utterance in the context of the current beliefs (triggering reasoning and/or adding new beliefs to the KL). (3) The stages may be followed by an “acting stage”, i.e. answering a question or performance of an indicated act.

Propositions may be asserted in KL because they have entered from the environment (told by some agent, or as a result of perception). Thanks to the above structure of the architecture, the symbols that an agent uses (in particular language signs) are not inserted into the “mind” of the agent by the programmer, but are acquired in the process of interaction with the environment and other agents. The PML and SAL layers provide *grounding* of symbols in Fetzer’s sense.

Let’s have a look at the “symbols” of the semantic network knowledge representation formalism (cf. Figure 1 again). A semantic network consists of nodes connected by directed, labeled arcs. There are two kinds of nodes (cf. Shapiro & Rapaport, 1987, pp. 275–276):

- atomic nodes that have no arcs emanating from them:
 - sensory nodes representing interfaces with the external world: input (e.g. lexemes) and output (actions), e.g. “Russia” or “big” nodes
 - base nodes representing individual concepts and properties. They have no descending arcs emanating from them, except LEX arcs (M1, M3, M5 nodes).

Atomic nodes have no structure. Only these nodes can be considered “atomistic symbols” (Emmeche, 2007, p. 466);

- molecular nodes, that have at least one arc emanating from them (except LEX arc):
 - structured individual nodes and properties – their structure is exhibited and comes from the part of the network connected to the node (B1 node)
 - structured proposition nodes that have no incoming arcs; they are beliefs of the system (M4 node)

Most nodes of a semantic network are molecular, i.e. structured (non-atomistic) nodes.

The only primitives in semantic networks are arc labels, which are non-conceptual (they are syntactic and have no meaning); The arcs that connect nodes contribute no conceptual semantics; they are punctuation only. Finally, LEX arcs point from concept nodes to words; concept nodes are not propositional nodes.

5. Representation in GLAIR/SNePS agents

There is an agreement that artificial systems instantiate at least dyadic relationships: “The representationalism of classical cognitive science, like the physical symbol system hypothesis has a dyadic notion of representation”.

(Emmeche, 2007, p. 466) This is the third element: an interpretant that is crucial in the discussion. I will argue that this notion of symbol is misguided with respect to contemporary artificial cognitive agents.

Emmeche reproaches cognitive scientists for an inadequate notion of representation (cf. Emmeche, 2007, p. 458). He seems to rest on the typical, outdated image of an artificial natural language using system, namely that of fixed set of symbols (such as the node M6), representing in a mind some entities in the world (such as the country, Russia) related or connected with a set of natural language utterances, such as “Russia”. The nodes and connections are pre-set by a programmer. This is, however, an overly simplified view and at least two aspects should be highlighted:

- there may be some initial set of symbols of a given knowledge representation system (such as a set of nodes). The size of that set is an answer to the question: what is the minimal knowledge so that learning and communication could take place?
- there are also some structural constraints (including the form of knowledge representation system such as a semantic network).

The majority of symbols, however, is acquired through processes of learning, not provided by a designer of the system.

In the case of an atomistic concept of symbol – in contrast to the Peircean, triadic notion – using language signs (and their interpretation) consists in recalling from a system memory a symbol connected with a natural language utterance. The situation looks differently in the case of SNePS agents: the structure of the semantic network is to be considered as a whole. The only atomistic symbols are atomic (sensory and base) nodes (see Section 4). Any other node can be regarded only as a part of the semantic network. Bearing in mind this fact, we regard elements (symbols) of semantic networks as structured entities: the structure of any node comes from the network. This is the general property of any semantic network knowledge representation formalism. As Quillian (1968, p. 239) states: “there are no word concepts as such that are ‘primitive’. Everything is simply defined in terms of some ordered configuration of other things in memory.” Speaking in terms of semantic networks: the single node never corresponds (refers) to any object in any domain, it corresponds only as a part of the entire network.

Secondly – in such a simplicistic view – an interpretation takes the form of the process of connecting an utterance (“Russia”) with a node (M6). Once the connection is created, it becomes a part of a fixed structure. A dictionary of a cognitive agent consists of a list of such pairs. Obviously, the antecedent sentence describes a dyadic notion of representation. Most of the contemporary artificial cognitive agents (such as SNePS/GLAIR agents),

however, do not function that way. They are systems capable of learning.⁴ In consequence, most of the nodes are created (or deleted) in response to natural language utterances (or – more generally – to perceptions). An artificial agent of the SNePS system can not only match parsed utterance with the knowledge base it possesses, but it can also (and usually does) create new nodes corresponding to natural language utterances, place them in the network, and erase nodes from the network.

In this point appears the third argument against the possibility of artificial semiotic systems, which could be called the *dynamicity problem*. Artificial systems are usually considered as static, “tending to disregard the open-ended and processual character of interpretation ... in strong contrast to the triadic, general and evolutionary nature of sign processes” (Emmeche, 2007, p. 466), whereas “semiosis implies process”. Peirce argued that “all thinking is dialogic in form. Your self of one instant appeals to your deeper self for his assent” (CP 6.338). The argument could be accepted as long as we talk about e.g. knowledge representation formalisms or programs considered as chains of rules waiting for application. By artificial cognitive/semiotic agents I mean, however, processes together with the structures they operate on, namely: not only semantic networks as directed, labeled graphs, but processes of inserting, extracting, and modifying the information represented in a semantic network. There is an inherent dynamic of constant re-interpretation of utterances by artificial agents, as will be described below.

6. Interpretation in terms of Peircean semiotics

I argue that a cognitive architecture – such as the GLAIR/SNePS and artificial cognitive agents designed using the architectures – could be interpreted in terms of Peircean semiotics. The knowledge of the agent, as stated above, is encoded in the form of a semantic network. An agent must interpret any linguistic expression (sign) in terms of an existing knowledge base to understand it. It means that such an agent should find a corresponding node in his semantic network (if there is such a node) or should create a new node.

Let us assume that an artificial agent perceives the utterance “Russia”. The word itself can be considered as Peircean Firstness – the sign itself. The utterance of natural language indicates some object – an entity in a real, imaginary, or fictional world represented as a node of a semantic network. The node (such as node B1) could be interpreted as Secondness. If our agent has no node attached to the “Russia” node, it has to create one.

The relationship between an utterance (attached to a semantic network via LEX node) and some node is not enough for understanding of the utterance. This is because the node is not a meaning of the utterance. The meaning or interpretation of any node of a semantic network is its *role* in that network. In consequence, the meaning of a sign is some part of the semantic network connected to that node. In our example the (limited) meaning of the B1 node is: “something that has the property of being a big country”. As Peirce states: “a sign, or representamen is something that stands to somebody for something in some respect or capacity ... The sign stands for something, its object. It stands not in all respects, but in reference to a sort of idea which I have sometimes called the ground of the representamen” (CP 2.228). In the context of this definition, the scope of that part of the network is the Peircean “respect or capacity” the representamen stands to somebody or something in. In other words: an Interpretant (Thirdness) is the sense made of the sign; in our case: (a part of) the rest of the semantic network connected to the node representing an object. Interpretation is the place of the node in the semantic network together with the part of that network surrounding the node and specifying the meaning of the node (and representing the respect or capacity). The meaning would and will change due to e.g. the needs of the system, learning, reasoning etc. It is that aspect that is a basis for the grounding relation *sensu* Peirce. If cognitive agents are capable of a choice of an aspect on their own (i.e. due to their aims, reasoning needs, etc.), the symbols they use are grounded. Less important is whether the symbols are grounded in Fetzer’s (1997) or Harnad’s (1994) sense (we can imagine agents using signs, functioning in e.g. a world of their dreams). This is the reason for stressing the role of Peircean *respect* or *aspect*.

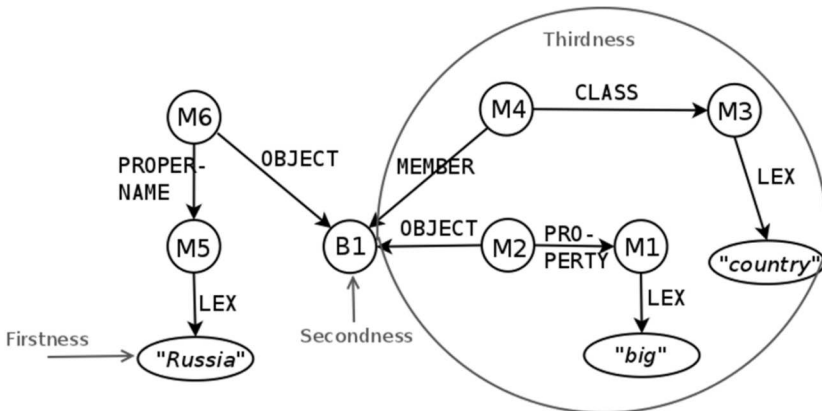


Figure 2. Semantic network representing the content of the sentence “Russia is a big country”

Figure 2 illustrates the correspondences between the notions of Peircean semiotics and semantic networks:

- The Representamen (Firstness): the form which the sign takes; in our case: the word “Russia”, represented as a LEX node.
- An Object to which the sign refers (Secondness), the concept (proposition), an element of the language of thought of an artificial agent, in our example: the B1 node in the semantic network to which a LEX node is attached.
- Finally, an interpretation (Thirdness) is (a part of) a semantic network connected to the B1 node.

7. Consequences: between Peircean semiotics and semantic networks

Recapitulating the previous paragraph: the *full* meaning of a node is determined by the whole semantic network representing an agent’s knowledge. Such an approach constitutes an example of the holistic theory of meaning: the meaning of any node is (in principle) the rest of the network. As Quillian (1968, p. 238) states: “a word’s full concept (i.e. node – PK) is defined to be all the nodes that can be reached by an exhaustive tracing process, originating at its initial, patriarchal type node.” Any new node added to the semantic network changes (again: in principle) the meaning of any node in this network. Accordingly, the meaning of any node (and any word) instantly changes along with the cognitive development of an agent. In consequence, the meaning depends on the whole network and its structure and the scope of the part of the network taken into account (i.e. used in reasoning). It should be emphasized that no single node has meaning (just like a single sign in Peircean theory).

The meaning of a node includes also the entire history of the acquisition of the concept; in fact, a connection is always made to nodes (concepts) in the context of learning – the set of nodes existing in a network will also depend on the order of presentation (order dependency). If an agent hears for the first time about Russia, e.g.: “Russia is a big country” it builds a node (M1) and asserts of it that it is called “Russia”. Hearing later “The country that the novel *War and Peace* is about” and not knowing that the novel is about Russia, it would build another node (M2). If the agent already knows that *War and Peace* is about Russia, it will connect the phrase to an existing node. A cognitive system creates intensions (nodes) only as needed for storing information about them. (cf. Rapaport et al., 1997)

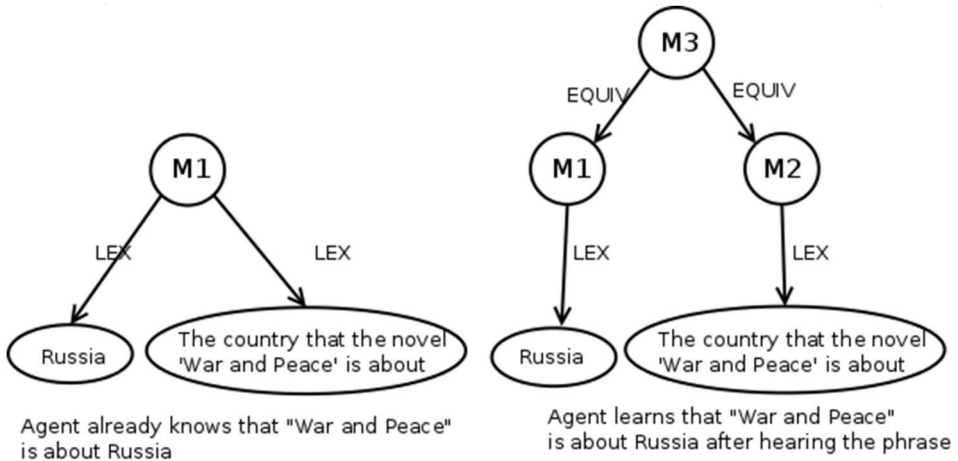


Figure 3. Adding new information to an existing knowledge base. The system does not create a new node (left) or has to create a new node for the utterance (right)

The meaning of a node can be theoretically taken as the full network surrounding the node; pragmatic considerations, however, require that we explicitly acknowledge that processing of such a network to derive such meaning must be limited. A SCOPE feature implemented in the GLAIR architecture would provide a measurement of processing extent and a measurement of a difference in meaning between two concepts. Incorporation of the scope feature would alleviate the problem of logical omniscience – the problem that cognitive agents do not know or believe all the logical consequences of their beliefs.

The two important properties of Peircean sign systems are fallibilism and synechism. Fallibilism, the process of constant reinterpretation of any sign and the subsequent, systematic “misunderstanding”, corresponds to dynamic revision of the knowledge base of the system due to interactions with external resources. As the meaning of a sign may be (and is) constantly changing, resulting in successive re-interpretations, similarly meanings of nodes in a semantic network are dynamic: they change along with adding new nodes to the network (or acquisition of new concepts). The meaning of the sign is not contained within it, but rather arises in its interpretation: once again, meanings are created in the course of language-usage, not discovered (conceptualism).

An isolated thought (and, in consequence, sign) has – in Peircean theory – no meaning. As Peirce states “it is a fundamental mistake to suppose that an idea which stands isolated can be otherwise than perfectly

blind” (CP 4.71). Meaning, according to Peirce, comes from other signs (synechism) – there is no absolute, “first” meaning. The synechism of the sign is reflected in a network structure of knowledge representation and implies the holism of the meaning of any symbol in the discussed cognitive system. An isolated node has no meaning; the meaning of any node comes from – as noted earlier – the surrounding network, other nodes connected with it.

8. Conclusions

Semiotic processes, according to Peirce, are in fact cognitive processes. I leave open for now the question whether all the semiotic processes are cognitive processes or only some of them. Even if only some of them are cognitive, we can still gain knowledge about the nature of semiosis by studying appropriate cognitive activities. Cognitive modeling may be a useful method of discovering properties of semiotic systems or processes. In other words, by creating artificial cognitive systems we are modeling cognition and looking for semiotic processes there. The similarities (presented in the paper) or dissimilarities can teach us something about the human mind as a semiotic system. Even within the Peircean approach, the processes of creation, use, and (re-)interpretation of signs are – from the point of view of cognitive science – underspecified. One of the advantages of the suggested approach is that it can fill in the gaps in traditional semiotic theories: suggesting specific computational mechanisms responsible for representation and (re-)interpretations of signs – in particular, the ability to interpret (partially) natural language utterances in terms of semantic networks, to revise beliefs in the context of new data.

I’ve tried to demonstrate that some artificial cognitive systems that are dynamic and use structured symbols, capable of natural language processing, could be examples of semiotic systems. In particular, artificial agents of the GLAIR/SNePS system using semantic networks as knowledge representation formalism fulfill the requirements posed by Fetzer. The systems could interpret natural language utterances in terms of their knowledge bases, finding a place for new nodes in a network structure. They also pick up an aspect of the described phenomenon or situation – namely they use only the part of the semantic network connected with a node. The scope of that part depends on the systems’ aims, problems to be solved, and the systems’ reasoning capabilities.

In general, excluding artificial systems from the semiotic spectrum, delimiting the spectrum in some way, is at least unreasonable. Emmeche (2007,

p. 457) defines semiotics in the following words: “concerning sign production, information transmission, communication and interpretation within some system”. It is astonishing that systems that gave us – in fact – the idea of information and information processing (i.e. artificial, computational systems), are explicitly precluded from semiotic research (cf. Zlatev, 2002).

N O T E S

¹ Another cognitive architecture which can be taken into account in the context is e.g. Anderson’s ACT-R architecture (<http://act-r.psy.cmu.edu>). However, the architecture is focused rather on modeling psychological processes.

² To present my own stance in terms of Peircean semiotics: our claims aspire to describe objective reality (or the domain of dynamical objects, noumena). As our theories are statements using signs, they represent the reality only in some respect or capacity. In consequence, our claims are fallible (false) – as they are only partial interpretations. Accordingly, the claims and theories are necessarily useful approximations and tools for prediction or action.

³ It is worth noting that Peirce is considered the creator of one kind of semantic network, namely assertional semantic networks. As Sowa (1992) notices: “In 1897, Peirce made a simple, but brilliant discovery that solved all the problems at once: he introduced an oval that could enclose and negate an arbitrarily large graph or subgraph. Then combinations of ovals with conjunction and the existential quantifier could express all the logical operators used in the algebraic notation (Peirce, 1909).”

⁴ The role of learning has been emphasized since the very beginning of research on artificial intelligence and artificial cognitive systems (cf. Turing, 1950). SNePS references on learning include: Martins and Shapiro, 1981; Choi, 1993; Ehrlich and Rapaport, 2005.

R E F E R E N C E S

- Burge, T. (1979). Individualism and the mental. *Midwest Studies in Philosophy*, 4, 73–121.
- Choi, J. (1993). Experience-based learning in deductive reasoning systems. *Technical Report*, 93–20. Buffalo, NY: State University of New York at Buffalo, Department of Computer Science.
- Ehrlich, K. & Rapaport, W. J. (2005). A cycle of learning: Human & artificial contextual vocabulary acquisition. In K. Forbus, D. Gentner & T. Reigier (Eds.), *Proceedings of the Twenty-Sixth Annual Conference of the Cognitive Science Society* (p. 1555). Mahwah, NJ: Lawrence Erlbaum Assoc.
- Emmeche, C. (2007). A biosemiotic note on organisms, animals, machines, cyborgs, and the quasi-autonomy of robots. *Pragmatics and Cognition*, 5(3), 455–483.
- Fetzer, J. (1990). *Artificial Intelligence: Its scope and limits*. Dordrecht: Kluwer Academic Publishers.
- Fetzer, J. (1997). Thinking and computing: Computers as special kinds of signs. *Minds and Machines*, 7, 345–364.

- Fetzer, J. (2001). *Computers and cognition: Why minds are not machines*. Dordrecht: Kluwer Academic Publishers.
- Harnad, S. (1994). Computation is just interpretable symbol manipulation: cognition isn't. *Minds and Machines*, 4, 379–390.
- Krampen, M. (1992). Phytosemiotics revisited. In T. A. Sebeok & J. Umiker-Sebeok (Eds.), *Biosemiotics: The Semiotic Web 1991* (pp. 213–219). Berlin: Mouton de Gruyter.
- Martins, J. P., & Shapiro, S. C. (1981). A belief revision system based on relevance logic and heterarchical contexts. *Technical Report, 175*, Buffalo, NY: State University of New York at Buffalo, Department of Computer Science.
- Peirce, C. S. (1931–1958). *Collected papers of Charles Sanders Peirce*. Volumes 1–6. C. Hartshorne & P. Weiss (Eds.), Vols 7–8, Arthur W. Burks (Ed). Cambridge, MA: Harvard University Press.
- Putnam, H. (1963). Brains and behaviour. In R. J. Butler (ed.), *Analytical philosophy: Second series*. Oxford: Blackwell.
- Putnam, H. (1975). The meaning of “meaning”. In K. Gunderson (Ed.), *Language, mind and knowledge* (pp. 358–398). Minneapolis: University of Minnesota Press.
- Rapaport, W. J., Shapiro, S. C., & Wiebe, J. M. (1997). Quasi-indexicals and knowledge reports, *Cognitive Science*, 21(1), 63–107.
- Russell, S. J., & Norvig, P. (2003). *Artificial Intelligence: A modern approach* (2nd ed.). New Jersey: Prentice Hall.
- Shapiro, S. C., & Bona, J. P. (2010). The GLAIR cognitive architecture. *International Journal of Machine Consciousness*, 2(2), 307–332.
- Shapiro, S. C., & Rapaport W. J. (1987). SNePS considered as a fully intensional propositional semantic network. In N. Cercone & G. McCalla (Eds.), *The knowledge frontier* (pp. 263–315). New York: Springer-Verlag.
- Sowa, J. (1992). Semantic networks. In S. C. Shapiro (Ed.), *Encyclopedia of Artificial Intelligence* (2nd ed.). New York: Wiley.
- Turing, A. M. (1950). Computing machinery and intelligence. *Mind*, 59, 433–460.
- Whobrey, D. (2001). Machine mentality and the nature of the ground relation. *Minds and Machines*, 11(3), 307–346.
- Zlatev, J. (2002). Meaning = Life (+ Culture): An outline of a unified biocultural theory of meaning. *Evolution of Communication*, 4, 253–296.