A Framework for Scalable Cognition

Propagation of challenges, towards the implementation of Global Brain models

David R. Weinbaum (Weaver)
Global Brain Institute
Vrije Universiteit Brussel

1. Introduction

1.1 Brains

This paper aims to introduce a conceptual framework for a scalable model of a brain. What is a brain? In biology a brain is an organ that evolved as a specialized means of survival (Ashby, 1960). The brain's general function can be summarized in the famous 3 Cs namely: Communication, Command, and Control of a more or less complex organism. Communication generally means the collection and distribution of information mostly via the central nervous system. Control, generally means the regulation and coordination of the organism's various subsystems aiming to keep a set of systemic parameters within their prescribed boundaries thus ensuring the survival of the organism in the presence of perturbations exerted by the organism's environment. The control function is also critical in the sense of establishing the distinction organism –environment. Basically anything which is not within the immediate and direct control of the organism can be considered environment[1].

Command, generally means the organism's executive faculty that is responsible for its goal oriented behavior. The functions under the title of command include predicting future situations, identifying goals, planning, supervising, learning and more, all of these in the service of the organism's primary imperative drives: survival and procreation. Brains, like any other biological system, evolved because they seem to provide increased fitness in the context of complex organisms. Communication, command and control are of course not distinct functions but rather three aspects of one complex function. There is no command without proper control and no control in the absence of communication.

Brain, the central nervous system of an organism, has become a metaphor to the control (actually the 3C structure) structure of systems in general but more in specific of autonomous agents, that is, systems capable of purposeful autonomous behavior. The general investigation of brains as complex control structures falls under the disciplines of general system theory and cybernetics on one hand and cognitive science on the other hand (Ashby, 1960) (Simon, 1980). While cognitive science was initially focusing on understanding the function of biological brains, the advent of machine intelligence and

[1] Organisms however do change the environment and even regulate it to some extent. But generally organisms do not have enough variety to fully regulate their environment see next section.

robotics much expanded this focus so the above mentioned research fields are in a process of convergence.

The following work is based on a fundamental premise inspired by knowledge of biological systems: Complex control structures emerge from the coordinated interaction of simple control structures to create hierarchies of control (Heylighen, 1991). For example: populations of cells organize into multi cellular organisms that in turn organize into societies, hives, schools and other super organism formations. The emergence of a new level of control is called a meta-system transition (MST) (Turchin, 1977)(Turchin, 1995) because it brings forth a new kind of autonomous agency. In specific, we know for a fact that the complex cognitive functions demonstrated by biological brains are realized by vast populations of interconnected neurons, each with very simple behavior. Such control structures are a product of self organization and not of a top down design and appear in a wide spectrum of structural complexity in various organisms (von der Malsburg, 1995). It is our working assumption that the emergence of control hierarchies (Simon, 1962) and subsequent metasystem transitions is, under the proper conditions, a scalable process bound only by physical constraints such as the availability of matter/energy or the speed of light. A scalable model of brain aims to provide a viable theoretical framework of how highly complex control structures emerge in populations of relatively simple interacting agents and what are the necessary conditions and specific mechanisms that may bring about such emergence.

1.2 Brains and the Global Brain

What is the Global Brain (GB)? A Global Brain is the projected product of the next metasystem transition of life on planet earth and possibly beyond (Heylighen, 1997)(Heylighen, 2008)(Heylighen, 2011). It is the projected emergence of a new control structure from the coordinated interactions of human and machine agents (and possibly other biological agents as well). The Global Brain will facilitate communication, control and command on the planetary level and will be in fact a new kind of autonomous agency with as yet unpredictable intelligent competences[2].

The challenge of bringing about the GB is not a matter of explicit design but rather in creating the proper circumstances for its emergence as a continuous process of self organization. Towards such end, there is much that can be learned from biological brains. The very paradigm of populations of interacting agents that bring forth complex cognitive functions is taken from neuro-anatomy, neuro-physiology and research into neural and connectionist models of cognition (Bechtel & Abrahamsen, 1991). However, the problem of a scalable model of cognition is at large an uncharted territory. First, the capabilities demonstrated by naturally evolved brains, even very rudimentary ones, far exceed the capabilities demonstrated by the most sophisticated manmade neural machines. But, moreover, there are

[2] Though it is a matter of speculation the Global Brain will be the next evolutionary phase of humanity and life on this planet. It will display vast computational resources and a superior intelligence capable of effectively tackling humanity's and the planet's hardest problems. It will be able to regulate the planetary system to ensure the necessary and even optimal conditions for life. It will be able to intelligently manage planetary resources and prioritize planetary level goals. It will be capable of scientific discovery and self improvement that will further accelerate its intelligence and fitness (Heylighen, 2002).

conceptual challenges involved: Biological brains are made from neurons which are more or less similar and communicate using more or less the same means and the same protocols. The GB as the convergence of human and machine agents must coordinate and connect a huge variety of different agents from simple sensors, to powerful domain-specific processors, to human agents of different capabilities and understanding, to AI agents, semi autonomous or autonomous robots and probably more. There is a vast variety and difference in competences, there are differences in the means of communication, in language, in understanding, in goals, drives, etc. How are we going to address such a formidable problem? Even if we did have a much deeper understanding of the human brain, such understanding would fill in only part of the puzzle, probably not even a major part.

The problem of the GB is first and foremost a conceptual problem that needs a conceptual approach i.e. extracting concepts and methods that apply across the many varied aspects mentioned above. The concept of machine or automata have brought the industrial revolution, the concepts of information and computation have brought the computer revolution. What are the concepts that will bring the next revolution of scalable cognition?

This paper suggests the concept of attention in conjunction with the theory of challenge propagation (Heylighen, 2012a) as one of the fundamental building blocks of a scalable brain.

2. Cognitive Agents

A cognitive agent is an agent characterized by displaying cognitive activity. Cognitive activity in the broadest sense may be defined as a non-trivial derivation of actions in response to events in the agent's environment(Franklin, Kelemen, & McCauley, 1998)(Franklin, 2008). Non-trivial here means that the derivation of actions is influenced by the environment, by the situation of the agent and follows a goal or a fitness criteria. Cognitive activity may also include adaptation/learning of future derivations of actions based on the success or failure of previous actions. For example: given a predefined goal or a set of fitness criteria, if the current action was a successful response of the agent to the current given event, the association between the event and the selected action is reinforced, otherwise the association is weakened.

Generally, an event is any difference in the environment that affects the situation of at least one agent. An action is any effect an agent may produce in its environment. Actions therefore produce events in the environment. An agent is identified by the events that affect it, by the events it is capable to produce and by the manner the latter are associated with the first. Of course the manner of association encodes a semantic structure; though it may include random or probabilistic elements, it cannot be entirely random. In other words, an agent must have structure. Similarly, the environment of the agent must have some structure otherwise there is no meaningful way for the agent to associate its actions to events since an environment without structure will necessarily respond randomly to the agents actions. The realization of cognitive activity must assume a structural coupling (H. R. Maturana, 1975) between the agent and the environment: differences in the environment induce differences in the agent and vice versa. Cognitive activity can be understood therefore as the manner by which an agent is affecting and being affected by the environment. Such activity has consequences on the dynamic structures of both agent and environment.

From the standpoint of an agent, the environment is the medium of events (i.e. differences) that affect it and are produced by it. When we think about a population of interacting agents, the actions of agents produce events that affect other agents and induce them to produce further actions. Agents share an environment to the extent that they interact through events, i.e. differences they exert on each other, but for each agent, all other agents are basically just environment.

Agents, in general, are not necessarily cognitive agents. A cognitive agent is distinguished from the general case in that that it derives actions in response to events in a non-trivial manner. If an agent's response to events involves no selection and no dependence on a state (i.e. no memory) than it can be considered trivial. Cognition therefore necessarily involves at least a situation or a state unique to the agent at a given moment and some selection mechanism. In other words, if all the information determining an agent's immediate action can be solely derived from the immediate event it is affected by, than, no cognitive activity is involved on the side of the agent and it would not be considered a cognitive agent. This is also consistent with the understanding of structural coupling because it means that the structure of the agent carries no relevant information affecting its response to events, i.e. it is unchanging. If the structure of the agent is constant, it is neither affected nor affecting the structure of the environment, i.e. no structural coupling. For cognitive agents, in contrast, there is extra information which is necessary such as the agent's state, goals, knowledge, tendencies, etc.

Following this rationale, the humble thermostat that features in so many examples in the cybernetic literature is indeed a rudimentary cognitive agent because the knowledge of its current state is necessary to correctly predict its response to an incoming event (change in temperature). A simple calculator which is immensely more complex than the thermostat is not a cognitive agent because its input sequence alone (buttons pushed) determines its actions (displaying alpha numeric characters on its display).

A simple yet a very general working definition of cognition can be given now: Cognition is the iterative coordinated processes of:

1. Selecting from the incoming stream of events which events are relevant and which are not. Relevance need not necessarily be a binary value. Events can be prioritized with varying levels of relevance according to the selection mechanism involved.
2. Given the current (most) relevant event, selecting from the available options of response what is the most appropriate action to execute next. An action may produce an event, change the agent's state or do nothing.

According to this working definition cognition is basically a selective process. Importantly cognition as selection is not deterministic as it may usually include probabilistic elements in the selective processes, or, dependency on a (possibly indefinite) number of previous events not all of which can be known. As such, cognition can be very simple or immensely complex. It can be more or less intelligent, adaptive, learning etc. Any one of the sub processes mentioned here can be redundant, but not both of them. If selection for relevance is redundant (i.e. no selection), every event is relevant so the agent is maximally sensitive. If selection for action is redundant (no selection), the relevant event solely determines the action so the agent is maximally instinctive. When an agent is both maximally sensitive and maximally instinctive it is not a cognitive agent anymore.

Implicit in this definition is that selection is made according to some set of criteria and possibly according to an internal state that may encode knowledge of performance in past events, goals, drives, knowledge the agents may have of its environment, predictions of future events and more. These implicit elements constitute together what may be called the context of the cognitive process. In the absence of context there is no cognition[3]. Relevance, the mark of the agent's intelligent interaction with its environment is if so context sensitive. It is the agent's dynamic situation which guides its cognitive activity. Consequently the agent's actions affect its situation closing a cybernetic loop through the environment.

Another important aspect of our working definition of cognition is that in information theoretic terms a selection is an operation that reduces the information contained in the agent's state. In other words, it is an irreversible operation. Once part of the information (the information content of irrelevant events for example) is lost, it is impossible to reconstruct the agent's state prior to selection using only the information present after the selection. More in specific, the selections the agent makes in response to events are constraining the variety (Ashby, 1958) of the effects of these events on certain parameters of its state (say its distance from a goal state).

Reducing information (variety) is how different agents gain different 'perspectives' in relation to the same events. Each agent reduces the incoming information in a manner that is specific to its state and goals while selecting. Selection for relevance and selection for action, therefore, render cognition an irreversible informational process. Clearly, this poses a criterion that differentiates cognitive agents from general agents: cognitive activity is characterized by reduction of information and the development of context sensitive perspectives. Following Ashby's discussion of requisite variety, a population of cognitive agents with diverse perspectives may achieve with coordinated activity higher fitness than is achievable by any single agent because they can together reduce more variety from the effects of the environment[4].

From yet another angle, cognition and most clearly the selection for relevance are symmetry breaking operations: prior to selection an input set of events may seem equivalent but post selection events can be regrouped according to relevance into smaller equivalence sets. The initial symmetry is broken (Heylighen, 1999). This notion of symmetry breaking is important in order to understand how the selections of the agent impart structure on the environment (see next section).

---

[3] There are of course borderline cases where the selective variety of the agent is very limited and the context therefore adds very little information. In such cases the presence of a cognitive process is a matter of arbitrary threshold.

[4] This is a good starting point to establish why coalitions of diverse cognitive agents can achieve more than any of the participating agents. It also alludes on how the measure the advantage of coalitions (see the following sections).

3. A framework for scalable cognition

The working definition of cognition suggested in section 2 above is inspired by Bernard Baars' global workspace theory of consciousness (Baars, 1993) (Baars, 2005) and the Stan Franklin's application of the theory in his work on the ontology of cognition (S. Franklin, 2006) (S. Franklin, 2008) and artificial minds (S. Franklin, 1995). Yet this definition aims to highlight different aspects of the cognitive process in order to prepare the ground for a scalable framework for cognition. In essence, the most important point of similarity to Baars' theory is that incoming information is scanned for relevant items and those relevant items gain for a while the global resources of the cognitive system in order to produce the appropriate response.

The global workspace theory is a theory that aims to explain functional consciousness (in distinction from phenomenal consciousness). It starts from a more or less given cognitive system which presents certain complex behaviors and tries to explain how these behaviors are realized. Very briefly, the Global workspace model of consciousness operates as follows: many highly specialist and relatively simple cognitive functional modules are working in parallel, processing incoming information and competing on grabbing the central stage of the agent's cognitive process. Once an item of information wins the competition it is globally broadcasted to all modules, recruiting a great portion of computational resources of the agent to further attend to the relevant piece of information while other items are being suppressed. This grabbing of the central stage means the item was 'brought to consciousness'. But the glory of each such item is fleeting as importance decays in time and soon the whole sequence of competition, and global broadcast repeats itself.

The starting point of our framework is fundamentally different. First it is synthetic and not analytic i.e. it does not aim to explain an existing system (i.e. human cognition) but to construct a general framework for an artificial cognitive process. There is no a priori given shape to the system. In fact we aim for open ended extensibility. Second, as we aim to describe a scalable cognitive process, we need to address structures which are self similar at various scales. This is not a requirement of the global workspace model. Third, our framework aims to describe cognition as a distributed within a diverse population of agents, while the original global workspace model, though utilizing massive parallelism at some stages, basically converges to a single stream of processing – the stream of functional consciousness.

We adapt if so from the global workspace model the basic idea as described above: the selection for relevance is the key for accessing resources in cognition. This process or the manner it is realized will be referred to from here on as the attention mechanism. Attention here is somewhat similar to functional consciousness in Baars' and Franklin's models in the sense that it indeed recruits the agent's resources but it differs from it in some important points. The difference lies in the constitution of the selection mechanisms. The mechanisms we need for our framework must allow for a distributed operation across many scales.

The simple competition and global broadcast model in Baars' model is replaced by a more general concept of ad hoc workspaces called coalitions. Coalitions are groups of interacting agents analogous to Baars' specialized modules. A coalition implies some kind of coordination and sharing of information among its participants that facilitates a collective cognitive function (i.e. specialized selection for relevance and selection for action preformed collectively). Coalitions are consolidated by means of spreading activation and are constituted from the resources and knowhow of the participating agents. Already Franklin (S. Franklin, 2006, pp. 14–15) mentions the forming of coalitions between competing modules in order to gain access to the global workspace. In our framework this is taken a few steps further: items of relevance rarely, if at all, gain global attention. Attention is distributed among coalitions which are ad hoc 'workspaces' with some specific capabilities. These as hoc workspaces are a product of self organization within populations of interacting agents. The self organization brings forth jointly both the item of relevance and a temporary coalition of coordinated activity that attends to it. In other words: items of relevance and the coalitions that attend to them co-emerge. This is a model more reminiscent of the way bacteria colonies coordinate feats of collective cognition (Ben-Jacob, 2003)(Ben-Jacob, Aharonov, & Shapira, 2004)(Ben-Jacob & Levine, 2006). In our global brain framework the workspace from which actions ensue will always be multiple, distributed and dynamic. It will not be confined anymore to the constraints imposed by how brains evolved and developed in higher animals with central nervous system.

When a single agent is considered, cognition is basically local and amounts to a recurrent selection for relevance and selection of how to act upon the selected relevance. What we aim at is to achieve an emergent distributed cognition within a diversified population of cognitive agents that self organizes through their interactions. Self organization however necessary is not sufficient. In order to achieve true scalability, the self organized system (a coalition) need to demonstrate an emergent cognitive function (distributed over the population of participating agents) that complies with the functional definition of section 2 but at a higher scale. It follows therefore that the definition of section 2 serves two purposes in our framework: 1. It serves as a conceptual definition for a generic cognitive agent. 2. It describes the self-similar theme of our framework: at each level the cognitive process can be described with the same conceptual tools.

How would such distributed cognition be described? Given a population of interacting cognitive agents[5]:

1. A subset of the population of agents identifies jointly a compound event as an item of relevance. A compound event is a set of contingent simple events that may or may not be causally connected or correlated.
2. The subset of agents forms a temporary coalition based on their shared 'interest'. Agents are reinforced to join coalitions they already joined in the past (Heylighen, 2012b). This is an abstracted version of the Hebbian rule (Hebb, 1968): Agents that cooperated once will tend to cooperate again i.e. reinforce their connections. The recurrence of a coalition follows the recurrence of its activating

[5] The population need not be homogeneous. Agents may vary in their structure and function as long as they share the same platform of communicating events to each other.

compound event, when the coalition is reinforced and becomes more stable, it also binds, so to speak, the set of contingent events that stimulated its formation into a distinct gestalt. This is how compound events are consolidated into 'compound items of relevance'. The cognitive process therefore takes an active role in structuring the information flow from the environment.

3. The members of the coalition jointly select and coordinate their actions in response to the item of relevance that brought them together. While a coalition is active, the resources of the member agents are committed to the coalition and the connections among them stays stable.
4. Upon completion of the coordinated action the coalition is dismantled. The agents become again free to seek opportunistic coalitions but they 'remember' the coalitions they participated in and their strength of association to each, so they can join them again even without being specifically activated. According to this description, distributive cognitive activity is opportunistic as coalitions may form only once. But the reinforcement mechanisms mentioned in 2 are necessary in order to ensure an ongoing tendency towards forming stable coalitions. Without a slight bias towards stability, the emergence of a hierarchy of more complex and capable super agents will be impossible. Following is an outline of such mechanism: the tendency to participate in a coalition is dynamic. It starts with a slight bias in favor of joining any coalition. Once a specific coalition has been joined, the tendency to participate in that specific coalition is reinforced for all participating agents every time it is activated. However, all specific tendencies decay in time back to their initial strength. If a coalition is not activated frequently it is eventually forgotten. The rationale behind this destabilization is double: first, a coalition that is not useful anymore because the circumstances of the environment have changed will slowly dissipate (be forgotten) releasing the constituent agents to form novel coalitions[6]. Second, the increase in scale necessarily implies an eventual decrease in the number of agents. If coalitions become too stable, the decrease in the number of agents necessarily implies a decrease in the variety of agents that can emerge at that scale. Decrease in variety will make the whole cognitive system less adaptable (with less options of configuration) at the high levels, again because the principle of requisite variety. But if we take care to preserve plasticity at all scales by not letting coalitions to become unnecessarily stable, the variety at each level will not be bound anymore to the number of agents at that level as coalitions can reconfigure according to need[7]. The structure we propose therefore, never rigidify beyond a certain threshold and never loses adaptability as it becomes more structurally complex.
5. Before being dismantled, the probability of an agent entering to a similar coalition in the future is updated according to the procedure outlined in 4. Reinforcements can be further adjusted positively or negatively depending on the relative success or failure of the coordinated action selected by the coalition. Such adjustments may allow an adaptive or even evolutionary process to take place modifying the structure of the coalition but as a rule, unsuccessful coalitions will dissipate quicker while successful and influential [8]coalitions will increase in stability.

---

[6] It is assumed that every agent can effectively participate in a limited number of coalitions. If the number of coalition exceeds the limit, the rate of conflicts and malfunctions will increase and the whole process may collapse.
[7] In short this ensures that there can be more kinds of agents than the number of actual agents at a given moment.
[8] Influential means that the coalition is either highly connected or incorporated in many other coalitions. See following sections.

6. Agents can participate in many coalitions, because usually their function in one coalition does not necessarily coincide temporally with their function in other coalitions. At those cases that there is a conflict, an agent will be bound to the coalition it is most committed to (according to the level of consolidation of the said coalition as reflected by the connection strength of its participant agents and possibly additional parameters). Since coalitions will also have functional redundancy like the one characteristic in neural nets, they will be able to perform considerably well also in the absence of a few agents as discussed in (Bechtel & Abrahamsen, 1991, pp. 60–62).

A recurring coalition is in fact a cognitive agent too. The exact mechanisms of internal joining and coordinating are not specified at this level. As long as the behavior of the coalition can effectively be described in terms of selection for relevance and selection of action, a cognitive process is taking place. When we want to emphasize the fact that an agent is constituted from a coalition of other agents we use the term super agent. As super agents emerge, they together select and respond to complex items of relevance in the environment. They discover patterns of events and form concepts of such compound events. The self organization process that brings forth super agents, also bestows structure on the environment shaping domains within which these super agents interact (for a discussion of cognitive domains see (von Glasersfeld, 1997)). In the same manner emergent super agents are capable of acting in a complex manner by producing mutually organized sequences of compound events. Super agents as coalitions of many simpler agents are not entirely predictable systems. Their characteristics may be given only in probabilistic terms and their inner structure i.e. the network of their constituent agents might itself be dynamic. In short, super agents in this framework are characteristically plastic and complex.

To achieve a scalable cognitive process it is necessarily required that super agents at every level will preserve the tendency and technical capacity to cooperate with other agents. In other words, the selective mechanisms of every super agent must be constructed in such a manner as to facilitate its further incorporation into yet higher coalitions. This requirement emphasizes the necessary structural self similarity of the framework.

To summarize, the framework proposed here combines the following characteristics:

1. Scalability – Hierarchies of agents that emerge from populations of simpler agents.
2. Plasticity – The tendency to consolidate coalitions is balanced by the freedom of every agent to form ad hoc opportunistic coalitions and the ‘forgetting’ of infrequent coalitions.
3. Self organization – no top down design is involved. Super agents emerge opportunistically and consolidate according to recurrent relevance (i.e. per repetitive occurrence of sequences of compound events in the environment). The self organizing nature of distributed cognitive activity is actually necessary to allow any kind of general problem solving capability. Also, such system is highly adaptive: if the nature of the environment changes radically, the consolidated coalitions that are not relevant will tend to fade out and new coalitions will consolidate in response to the new conditions.
4. Heterogeneity – Agents of different kinds and function are capable of interaction and coordination.

The realization of such a framework involves of course many difficult problems that will not be discussed here but its basic prospects of success involve finding solutions for two conceptual problems. The first which is implicit in the above understanding of cognition has to do with the very mechanism of attention: Given the situation (state) of an agent and its goal(s), how the selection of relevant events takes place? Our hypothesis is that the large variety of possible attention mechanisms is based on a relatively small number of selection mechanisms such as competition, consensus forming, co-inhibition, alignment and self organization. All these are in fact, as mentioned already above, mechanisms of reducing the information of an input set, or, in other words mechanisms of symmetry breaking as was already mentioned in section 2.

The second problem addresses more in specific the case of distributed cognition. How groups of simple cognitive agents align their (possibly diverse) local attention mechanisms into a collective attention of the group? This emergent mechanism need to be capable to attend holistically to complex items of relevance, where each such item may be heterogeneously distributed over the population of agents[9] as well as spread over the temporal dimension. Such alignment, according to our framework, culminates in the formation of a coalition of agents and implies further coordination of goals and sharing of resources as steps towards a coordinated selection of action. This problem is in fact a version the well known binding problem in cognition (Wikipedia contributors, 2012). Important aspects of it are also discussed in (Heylighen, 2012b).

The two problems involving the attention mechanism are of course strongly coupled. In fact they can be considered as two aspects of the same problem. The first aspect focuses on modeling the cognitive agents in functional terms without highlighting how the functional components are implemented. The second aspect is inspired by a connectionist approach that highlights the emergence of functionality, at any level, from the interactions of elements one level lower. The first perspective highlights function while the second highlights structure. It is the synthesis of these two aspects which is necessary for the realization of a scalable cognitive process.

The rest of this paper will address these two approaches to attention towards a possible synthesis. Here is the place to introduce the concept of challenge.

4. The distributed attention mechanism and propagation of challenges

According to the challenge propagation paradigm (Heylighen, 2012a), a challenge is any event that invites an agent to act. A difference in sugar gradient for a bacteria, the sight of a predator for a deer, the ring of a telephone for someone waiting for an important call, a significant continuous increase in human produced greenhouse effecting the global climate (for whom is this a challenge?), these are all events that can be considered as challenges. It is clear that a challenge must be a challenge for someone or something and therefore not every event is automatically a challenge. For an event to become a challenge it is necessary that it will be selected by an agent to be acted upon. An event may be a challenge for one agent and not for another. Similarly an event may be a challenge for an agent at a

[9] For example: the touch, the scent and sight of a flower are communicated via different agencies. To construct a concept of 'flower' these agencies must somehow interact and exchange information. For another important example see (Seeley et al., 2012) describing a decision making process in swarms where the relevant information is distributed among many agents.

given point in time and not be a challenge at another point etc. Events that are challenges at a given scale (bacterial) are not even registered on other scales (human).

The dependency of challenge formation on scale may explain why certain events or processes like climate change do not gain the resources necessary to address them effectively. Simply, there no cognitive agency in the proper scale that is enough consolidated to attend to them. Since we have argued above that the consolidation of an agency (as a coalition) and the consolidation of an item of relevance (a challenge) are co-dependent, it seems that in the case of climate change, such consolidation has yet to take place in order to allow effective coordinated response.

The concept of challenge clearly fits the framework developed above. We propose that challenges are the products of attention mechanisms. As such, a challenge is an item of relevance; it is not a mere difference, not even a difference that merely triggers an action. A challenge is the bringing forth of a *context sensitive relevance* which sets the ground for the selection of action. This context sensitive relevance may include value(s) (negative or positive in relation to given goals and state), measure of importance of the item i.e. (how strong are its effects on the agent), priority (how fast it should be responded to), meaning (its effects on the general state of the agent), and so on, all in relation to other concomitant circumstances.

The notion of challenge is quite abstract: agents need not have the same goals or the same attention discrimination (what they decide to be relevant) in order to respond to challenges in a cooperative manner or even coordinated manner. Yet if challenges are context sensitive and therefore particular to an agent, how are we to understand the notion of the propagation of challenges among agents? It is by answering this question that the significance of challenges to our framework is exposed. In the following subsections we explore two important conceptual descriptions of how challenges propagate among agents.

4.1 Influence networks

We first consider the meaning of propagation of challenges between two agents. For this we need to define the concept of influence. Given a network of interacting agents (not all of which are necessarily cognitive), let us consider two distinct cognitive agents A and B. We will say that B is influenced by A if and only if the selection items of relevance at the locality of agent B are caused, directly or indirectly, by the actions of agent A. Since all the actions of A are selected based on items of relevance for A, it follows that the selection for relevance by B is influenced by the selection of relevance by A.

A slightly different case is when both A and B select their items of relevance in a correlated manner. This correlation is caused either by the fact that both A and B respond to the same event(s), or, they may respond to distinct events say E_A and E_B that are somehow causally connected (without such causal connection being in anyway given or explicit). In such case, we would say that both A and B are influenced by agency S (for source). S can be, in the most general case, any structured phenomena in the environment, that is, a phenomenon that produces causally connected events. For example: wood fire produces simultaneously heat, light, crackling sounds and distinct smoky smell. These can become items of relevance to different agents – in this case, specialized sensory modules in the brain of an organism.

All the said effects of the wood fire are causally connected. Their corresponding events E_{Heat}, E_{Light}, $E_{Crackling}$ and E_{Smoky} are nevertheless distinct events whose causal connection is not given besides the fact that they may produce a correlated influence on a number of distinct agents. The case of correlated influence will be further discussed in the next section.

In a population of interacting cognitive agents and possibly a set of sources, we can conceive of deriving an influence network that represents how challenges are propagated among agents and sources. If two agents are not directly interacting and yet influence each other, this influence must be mediated by other agents. An influence network will indicate the paths by which such influence is mediated, or, in other words the topology of challenge propagation. In terms of attention mechanism, the influence network embeds the information of how local challenges may or may not propagate among agents to produce widespread or global items of relevance. In other words it may indicate trends of interest and relevance that characterize the attention dynamics of the whole population in relation to the environment.

Influence networks is a field of study in social sciences that investigates the spread of ideas, opinions and attitudes in various social contexts (Friedkin & Johnsen, 1999)(Kempe, Kleinberg, & Tardos, 2003). Though a model of propagation of challenges as influence addresses different problems as it is focused more on the general relations between challenges and actions our research can merit from the mathematical models that were already developed to tackle propagation of influence (see also (Montgomery James, 2011, chap. 8: Influence Networks)).

When it comes to cognition in our framework, influence precisely means what events attract the attention of agents and consequently recruit their resources towards action. It seems that in the large picture the flow of influence is highly predictive to the flow of resources. Cognitive agents can use the information embedded in the influence network in at least two important manners:

1. To adjust their own attention mechanisms to be more discriminative. For example: an agent can learn which of its neighbors is most influential, or, which events of relevance tend to gain global attention and adjust accordingly. This can become a basis for division of labor and specialization among agents.
2. To adjust their actions in order to attract the attention of other agents to certain events (and by that recruit their resources), or, alternatively, inhibit the attention mechanism of other agents in regard to certain events. This can become a basis for alignment, sharing and cooperation among agents (for further discussion see (Heylighen, 2012b).

If we consider agents that are capable to adapt their attention mechanisms and/or their selection for action mechanisms, and base such adaptations on the information embedded in the influence network, we can conceive that such adaptations may increase cognitive fitness. Generally, without referring to any specific kind of cognitive system, the cognitive fitness of an agent in our framework can be expressed in terms of the quality of its attention mechanism and the selection of appropriate (successful) responses thereof. Becoming more discriminative in attention and more influential in action are clearly correlated with increased cognitive fitness of an individual agent in the population[10].

To achieve a global increase of cognitive fitness, we need to assume that cognitive agents operate with tendency to regulate the usage of their own resources (self preservation), and the resources available in their environment (exploiting other agents' resources). More specifically, agents will tend to become more discriminative in their attention and therefore less prone to be exploited by their peers and on the other hand become more influential in order to better exploit their peers. Influence and discriminative attention are, therefore, qualities that balance each other in the ongoing dynamic interactions among cognitive agents. The more agents will try to exert influence over each other, the more they will need to become discriminative in their attention in order to counter balance increasing influence. If, on the other hand, influence diminishes, agents will tend to become more sensitive and less discriminative. In the context of the overall population, these adaptive tendencies are sufficient for a process of self organization to take place reconfiguring the influence network into more or less stable formations of local optimum. A self organizing process of optimizing attention and influence seems to be the conceptual method of achieving increased levels of cognitive fitness of the overall population. Such adaptive process may take place in parallel to the ongoing cognitive process and in a much slower timeframe.

Influence networks of the kind described here can be modeled as economic systems[11]. Events are assigned with value in terms of influence. Every selection for relevance by an agent is an economic transaction of selling and buying, the higher the value of an event for an agent the higher its probability to be selected as relevant (buy). Influence as value is of course local in the sense that the same event may have different influences on different agents. However, there is a sense to think about emergent global values of events of a specific kind. If an event becomes relevant to many agents, its influence increases and with such increase, it may become relevant to yet more agents. If the influence of an event diminishes the same mechanism of positive feedback will make it even less relevant for the overall population. A cognitive agent that produces events of high influence can be said to have more availability of resources because besides its own resources it is capable to mobilize to action many other agents. If an agent is uniquely capable to produce some special event, it is analogous to an owner of a resource of great scarcity and will be therefore highly discriminating in its selection for relevance, or, in other words only highly influential events may be selected as relevant and gain access to that resource. Moreover, if influence is analogous to currency, actions are products and the value of products is measured in terms of the resources they demand. In principle, influence as currency buys actions as products. The relations between influence and product value may define the economic dynamics of such economic model.

---

[10] Normally, cognitive fitness will be the measured capability to perform benchmark cognitive tasks of various kinds. In the discussion here cognitive fitness means the agent's qualitative capability to produce successful effects with fewer resources. Fitness, again, translates to selection: a fit attention mechanism ensures that the agent will not waste time and resources on non-relevant events; a fit selection for action ensures that the spending of resources is effective.

[11] The general idea of attention as a commodity that needs to be managed appears in (Davenport & Beck, 2001), though in a much more specific and business oriented context.

The most interesting aspect of such economic model is that the propagation of challenges, or alternatively the flow of influence, carries resemblance to a monetary flow in a financial market. In a market the flow of money abstracts away many of details of the specific transactions taking place. Nevertheless, a global market dynamics evolves and can be studied. In influence network, many of the details of local selections of agents are abstracted away too while the whole population of agents may present emergent global patterns of influence. A model of interacting agents with an attention economy is yet to be worked out in detail. It may present an interesting avenue of researching methods to implement self organizing cognitive systems with a capacity to increase cognitive fitness by way of self organization.

4.2 Vertical propagation

In vertical propagation of challenges we mean to indicate a process by which one or more challenges bring about the emergence of a compound higher scale challenge which is attended by the coordinated attention mechanism of many agents. Since every challenge is the outcome of selection for relevance by an agent, vertical propagation is a concept dealing with how local items of relevance may coalesce into a global item of relevance within a group of interacting cognitive agents. In our framework, vertical propagation is the first stage of the formation (and later the consolidation) of a coalition of agents and the coordination of their actions. Vertical propagation is therefore critical to the whole framework of scalable cognition. It is dubbed vertical because it is the process that propagates relevance from one scale of cognition into the next scale. Clearly, vertical propagation must be facilitated by influence networks. It is the flow of influence among agents that may result in their coordinated activity which is necessary for vertical propagation.

There are quite a few conceivable mechanisms that may facilitate vertical propagation. Here are a few examples that are certainly non-exhaustive:

1. Simultaneity based propagation- A recurrent set of events that are simultaneously selected as relevant by many agents may be a trigger for these agents to form a coalition and align their subsequent actions in response to it. Such alignment will be selected on the basis that the overall coordinated (and probably timed) response to the set of events is more successful than a non-coordinated response. The alignment must be a gradual learning process and therefore demands the recurrence of the triggering set of events.
2. Interest based propagation – Agents that select the same events or types of events as relevant, even at different timings, may start to exchange information on the best actions to be selected in response. They form a kind of coalition based on shared interests and agenda and bring forth a kind of collective intelligence where all the participants share the learnt experience of each. Such coalitions may be based on collaborative decision making, specialization and more. Such coalitions operate asynchronously.
3. Pattern based propagation - Pattern based propagation is based on the dynamics of influence networks from the previous section. When a set of agents activate one another through causally connected events (A influences B, B influences C, C influences D and E, etc.) They may participate in responding to an initial challenge in a manner that none of them could realize alone. Coalitions that emerge based on patterns are the kind of cooperative processing which is best suited to realize

distributed algorithmic computations, and other deterministic kinds of processing that involve recurrent temporal series of events. In this sense pattern based propagation is an extension and generalization of simultaneity based propagation. A temporal pattern of simultaneous compound events may bring forth a coalition of coalitions. First there is a set of coalitions that emerge in response to each simultaneous event separately. Then if those compound events tend to appear in some recurrent order they may bring forth another pattern based coalition that coordinates the set of coalitions into a coordinated whole.

Again, the defining property of vertical propagation is that it links individual challenges at one cognitive scale to (collective) challenges of the next cognitive scale. Each cognitive scale is distinguished from its previous one in that that its cognitive agents are in fact coalitions of agents of the previous scale. The emergent collective challenge cannot be apparent to any of the individual coalition members; it is clearly a compound event that is *relevant* only for the group as a group. In our framework, vertical propagation of challenges is both the cause and product of the formation of coalitions. Complex items of relevance on one hand and coalitions of cognitive agents that attend to them, on the other hand, co-define each other. A complex item of relevance appears as a structured feature of the environment in as much as there is a cognitive agent for whom this item of relevance appears as a product of its collective attention mechanism. Agent and challenge not only co-define each other in the formal sense they also co-emerge in the functional sense.

A complementary aspect of vertical propagation of challenges exists in the way already established coalitions may activate their constituent agents and recruit their resources for a shared goal. Such activation can recruit an agent even if it was not individually activated by an event from the environment. In contrast to the bottom up propagation described above, this is a top down propagation from the super agent to the agents that constitute it. Notice that there is no specific activating entity, but an emergent mutual activation in response to a certain set of events. But it is sometime clearer to describe the happening as a downward propagation.

It is also possible to see a coalition as a kind of episodic memory if we understand episodes as recurrent compound events. The coalition structure 'remembers' a challenge as a structured event of relevance and associates it with a compound response activity. This memory persists as long as it is reinforced by the recurrence of the relevant compound event. Once the environmental circumstances will cease to produce the relevant compound event, or, there will be a change in goals or value system of the agent that will render the event irrelevant, the coalition – the structure reinforced by this item, will start to dissipate and eventually disappear. The resources committed to respond to the episodic event will not be bound anymore as well and the flow of cognitive processing will continue along other avenues.

The adaptive mechanism described here as dependent mostly on the frequency of events in the environment may of course be modulated by additional factors. Yet the frequency of events at different levels is an important characteristic. When we compare events of relevance at different scales, it is plausible that they exist in distinct temporal scales. For example: if we consider a population of multi-cellular organisms, we will be able to observe a few temporal scales of cognitive processes that are correlated to the levels of complexity of events. Molecular events within cells, cellular events within a single organism, metabolic and behavioral events at the level of a whole organism and events at the level of the specie (a few additional intermediate levels can be considered here too). The different levels have between them vertical propagation of challenges and therefore exemplify a scalable cognitive process that takes place across a few scales. Each level seem to have a more or less distinctive time scale defined by the average duration of processing events at that scale i.e. the time from the introduction of a (compound) event from the environment to the agent until the agent's response action is issued. Simpler agents are processing simpler events much faster than coalitions of agents are processing compound events. This is why each scale of the cognitive process can be considered to have a different span of memory. One can say that each scale is characterized by a more or less distinct temporal granularity.

Modeling processes of vertical propagation of challenges seems to be the most important enabling factor of a scalable cognitive process. We attend to this challenge next.

5. The challenge of modeling scalable cognition

The framework presented to this point, has a few significant advantages. It has very few assumptions, it provides a working definition for a cognitive process, it takes into account the diversity of agents in terms of their capabilities and agendas, and finally, perhaps most importantly, it suggests a scalable cognitive process. Yet, the framework suffers of a few obvious weaknesses as well. First, it is not obvious at all how this framework may solve problems or be usefully productive. Second, it does not specify formalism and does not address issues of implementation. We are going to address now these weaknesses and see what can be done about them at this level of abstraction. Our aim here will be to specify clearly what should be expected from formal modeling of scalable cognition and what shouldn't.

A few philosophical issues need to be considered when we try to formalize and then model something as abstract as a global brain. In order to formalize our framework we will have to decide on quite a few constraining assumptions. Ideally, the assumptions we will make, will get us closer to realizing our framework even making it do something useful but without losing generality and scalability. This is a very difficult thing to achieve. We expect to demonstrate something reminiscent of general intelligence. There are many debates regarding the meaning of intelligence that we will not address here. One commonly accepted working definition defines intelligence as the capability to solve problems. But this definition assumes that a problem is already given. General intelligence, in distinction, is capable not only to address already formed problems but to recognize and form them in the first place. It then can generate a specific problem solver that will produce solutions.

When we observe the effects of general intelligent activity, we usually see the manifestation of a specific kind of intelligence in the course of solving a specific problem. For example: we can see the intelligent activity involved in playing a strategy game like 'Go', in riding a bicycle, in designing a tool, in speaking and understanding language and more. We rarely attend to those situations where it is not obvious at all what need to be done or can be done. Imagine for example an intelligent agent figuring what bicycles are without ever seeing them used before (Giving up the assumption that the agent is human will expose how difficult the issue is). But if general intelligence is known only through particular manifestations of special purpose intelligence, trying to demonstrate the latter, we will probably fail to achieve the first. In the face of this difficulty we are lucky to have in our hands working specimens of systems that realize general intelligence, primary of which is human brains.

What is fascinating about human (and to some extent animal) intelligence is its outstanding capability to invent, cultivate and apply an impressive variety of special purpose intelligent activities. In this sense human intelligence is a 'universal' problem solver. Remarkably, as far as science grasps, the impressive feats of this general problem solving are based on nothing more than a large population of relatively primitive cognitive agents (neurons) that interact among themselves.

In the course of their interactions they self organize to form hierarchies of more or less stable coalitions that specialize in performing specific functions. We know that the complex functions that emerge from this system involve a large degree of diversity among agents (many kinds of neurons) and plasticity in the way they connect and interact (synaptic plasticity). Remarkably, plasticity is realized not only at the level of neurons and synapses. Neural modules develop in an evolutionary process to serve multiple functions (Edelman, 1987)(Edelman, 2004). While neurons seem relatively well located in terms of their connections to other neighbor neurons, when we switch to the functional level of description, every neuron normally participates in a few structures where it serves in implementing a variety of functions. This achievement of evolution yields a highly efficient and fault tolerant system. Every component has multiple uses and at the same time every function can be carried out by more than one functional unit.

The rationale behind our framework comes to capture these very characteristics with a few important exceptions that are introduced to simplify matters and release the framework from unnecessary constraints characteristic to organic systems:

1. The neuron as the fundamental building block of brains is abstracted: A cognitive agent is not a model of a neuron and need not necessarily resemble one. It is a functional unit that may appear at many scales. In fact, what is significant in our framework is not so much the structure of a cognitive agent but rather the preservation of functional isomorphism across scales.
2. The diversity of cognitive agents allowed by our framework far exceeds that of species of neurons.
3. The scalability allowed by the framework is not constrained neither by any anatomical plan nor by specific goals/problems a designer might have in mind. The emergence of higher levels of cognition through the formation of coalitions is theoretically unbound as long as there are enough matter and energy to facilitate structure[12].
4. The Hebbian reinforcement (Hebb, 1968) at the basis of neural plasticity is replaced by an abstract reinforcement mechanism that reinforces the commitment of unit cognitive agents to the coalitions they participate in. The major difference here is that while Hebbian reinforcement is basically local, and operates between connected agents, coalition reinforcement is quasi local and connects a subset of the population.

[12] There are of course other physical limitations that are beyond the scope of this paper.

With these exceptions our framework becomes a platform, where many brains both organic and artificial, with diverse capabilities, form coalitions capable to discover, form and solve together problems.

Addressing the first weakness mentioned above we note that it is not a weakness at all but a characteristic which is shared by every research project that aims to study and eventually demonstrate general intelligence or universal problem solving. A system designed to demonstrate general intelligence does not have any immediate obvious usefulness especially while it is developed. Its usefulness needs to be achieved through a multi-leveled self organizing process that takes place within an environment that specifies circumstances and constraints that will trigger the eventual emergence of specific problem solvers but without losing the system's capacity for general problem solving. Remarkably, our framework does not even specify system boundaries at this point since such boundaries are co-emergent with cognitive functions as already mentioned before. But even with so few assumptions there are still necessary requirements that can be derived and used as guidelines for modeling.

1. Capability to interact – Agents at any scale need to be able to interact with the environment and with each other through selection and response to events. All interactions will be facilitated via a message passing communication layer (medium) that will be given and its implementation is opaque to the agents.
2. Tendency to cooperate - Agents at any scale must be driven to cooperate i.e. to form coalitions with other agents. For the efficient operation of coalitions a layer of communication needs to be established that is dedicated to coordinate the mutual influences and actions of agents within coalitions. This is where influence networks may be the primary theme.
3. Tendency to persist - Coalitions (i.e. super agents) at any scale must be driven enough to hold together against a constant inherent tendency to dismantle (dissipation, entropy). This drive will be correlated to the relevant function(s) they carry out successfully. Persistence may be achieved by learning mechanisms similar to those used in neural nets where links among agents are made stronger in relation to the frequency of usage and possibly other parameters. For example: weights of links can represent the probability of activation of an event sent through the link. Reinforcing links will ensure a higher probability of coherent interaction converging into deterministic mechanisms as probabilities approach 1.00. Such reinforcement will always work against a constant but relatively slow decay of the strength of links that can be interpreted as 'forgetting'. Everything not useful enough will eventually be eliminated. Yet, useful configurations will easily overcome decay. Additionally, a measure of criticality will be assigned to each agent in proportion to the number of different coalitions it participates in and the degree of influence it has within each coalition. Critical assets (i.e. 'special' super agents) will be made to decay slower.
4. Capability for measurement - Agents at any scale must be capable to assess their own performance i.e. relevance and successful response according to a given set of values. Measuring performance is necessary for all feedback mechanisms, learning, adaptation and self organization. Again, information about performance might be distributed through a dedicated communication layer.
5. A generic design for a (primitive) cognitive agent - The design of a primitive agent is necessary not only because pragmatically we need to start somewhere, but also because the design of the

primitive agent is also the functional template for all the cognitive units recurrent across scales. A good model for a generic agent is one of the most critical aspects of modeling the framework. Any consolidated coalition of agents needs to have the capability to 'hide' its structure within a shell that behaves like a generic agent. In other words, coalitions need to have the capacity to present themselves externally in the form of a generic agent. Basically such a presentation will include an input event set, an output event set and a few other conduits that mediate more specific communication and control functions according to the specific design. The I/O sets will need to be designed such as to represent events of scalable degrees of structural complexity. Challenges will be the internal representations of such events that might be different for different agents. These internal representations will not be available but to the agent whose attention mechanism selected and produced them. The significance of incoming events to an agent will be known to other agents only through correlating such events with the events produced by the agent's actions.

Additional to these guidelines, there is another important perspective regarding cognitive agents and their interactions. A cognitive agent can always be described from two complementary points of view. The first point of view focuses on the structure and function of the cognitive agent in terms of input/output relations. From this point of view the agent is an independent dynamic system with more or less specific behavioral tendencies. Every input event, combined with the current state of the agent will make it follow a trajectory towards another state. Aspects of the trajectory and the new state will also determine possible output events. The second point of view focuses on the connections and coalitions a cognitive agent participates in. From this point of view the agent is described in terms of its capacities to affect and be affected by other agents in a heterogeneous population of interacting agents (De Landa, 2011, pp. 1–6), or, in other words, by the variety of challenges and responses to challenges. Following the understanding of challenges as events of relevance, any capacity to affect must involve another agent's capacity to be affected. The important point to note here is that these capacities cannot be predicted from the dynamics of any single agent. Moreover, the variety within the interacting population of agents gives rise to a combinatorial variety in capacities to affect and be affected which is translated to a variety of possible combined behaviors. It follows that the behavior of coalitions of agents cannot be determined solely from the behaviors of the component agents. It might seem trivial at first sight but this unpredictability of capacities to affect and be affected is the source of the nearly inexhaustible innovation potential of coalitions. Any new capacity that emerges expands the horizon of additional new capacities. New capacities depend of course on the properties emergent in the dynamics of coalitions. When a coalition consolidates the various capacities that are expressed in the interactions among its participating agents give rise to the dynamism associated with the independent super agent that emerges. The new agent's properties are the resource of the capacities of the next level.

The combinatorial alignment of capacities among interacting agents is the process that critically produces novelty via the formation of new coalitions (super agents) and stand therefore at the basis of the emergence of novel cognitive functions. From this perspective we can plausibly argue that our framework is capable, at least conceptually, to produce open ended innovation in cognitive capacities. This is of course a necessary and critical prerequisite for demonstrating general intelligence.

We are left now with the second issue having to do with the more pragmatic aspects of formalizing the framework and modeling it. It should be clear that the aim is demonstrating a universal problem solving capability. The framework can initially be constructed to address constrained families of challenges in a model environment in order to establish and fine tune its capabilities. The ultimate test however is to realize a model demonstrating an emergent intelligent response in diverse and changing environments. In other words, the model will be expected not only to solve problems presented to it, but to discover problems and goals (challenges) at a few levels and be able to respond to them effectively without the guidance of a designer.

What will be given below are preliminary ideas as to how various aspects of the framework might be modeled. Detailed formalizations are of course a rich prospect for further research.

5.1 Modeling agents

The cognitive agent is the basic building block of our framework. As it is apparent from the above a cognitive agent is constituted from two functional modules: the first implements the selection for relevance while the second implements the selection and production of action (actuator mechanism). Generally, this rather abstract distinction need not be explicit in an actual model. Since the most interesting aspect of the framework is the demonstration of an emerging scalable cognitive process, it is clear that beyond the most primitive level, agents will be networks of simpler agents. Still, the distinction between attention mechanism and actuator mechanism is meaningful in the sense that the structure of the attention mechanism encompasses the capacity of the agent to be affected, while the structure of the actuator mechanism encompasses the agent's capacity to affect. Keeping these capacities more or less decoupled also means they can undergo adaptation and learning independently which might render modeling simpler especially in how higher cognitive functions may be constructed from more simple ones. For example: a complex attention mechanism can be constructed from a collection of similar more primitive mechanisms (belonging to the lower level agents) organized in a more or less simple logical decision tree such as combining all their outputs by an AND gate, an OR gate etc. More complex decisions can be implemented by constructing a novel (internal) event from the outputs of all the constituent components and feeding it to another selection mechanism (A logical function, a classifier or pattern recognition mechanism etc). Similarly, a complex actuator can be constructed from more primitive elements (again belonging to the lower level agents). Of course the implementation of agents at any level must facilitate such constructions and keeping the attention and actuator modules independent is a significant design decision towards that end. In more advanced scenarios, agents within a population can be bred with each other to produce, by means of genetic algorithms, offspring that combines variations of the attention and actuator mechanisms of the parents.

A more difficult and interesting problem is how to model an agent's history, resources, values and goals in a manner that would allow agents to share resources and even more importantly the inner representations that comprise the agent's knowledge. Sharing resources and sharing knowledge seem to be two of the most fundamental mechanisms of collective intelligence. If we relate to the inner representation of an agent as a kind of a 'worldview' (Vidal, 2008), agents need to be capable of sharing at least some aspects of their 'worldview' and establish consensual domains among them(H. Maturana, 1978). Forming a consensual domain is more or less equivalent to sharing context dependent information, or, in other word, sharing meaning (see also in the next section).

An approach to model this aspect of the framework would be to define an ontology from which all objects such as events, values, goals, histories etc, are derived. Agents would be able to possess quite diverse inner representations but they will share ontology. Based on such ontology it will be possible for agents to create consensual domains and possibly join their resources, commensurate values and share goals. Moreover, as coalitions are formed to produce super agents, it will be possible, as part of the process, to construct new complex representations of the environment based on the already established representations of the constituent agents from the lower level. Such approach supports the understanding discussed above that the structure projected on the environment by the cognitive process becomes progressively more complex and rich as additional levels of the cognitive process emerge from previous levels[13].

5.2 Modeling events and challenges

It is very important to clarify at the beginning the conceptual differences between events and challenges. An event is basically a piece of information that presents a change or a difference in the environment. In our framework the designation of 'environment' is pretty general and open to further specification. Basically, from the standpoint of every cognitive agent, 'environment' is the source of events that may trigger its attention mechanism. A challenge is an event that has triggered the attention mechanism of at least one agent. The triggering itself does not change the informational content of the event; it rather puts this informational content in the context of the agent's goals, history, values, and resources. The selection of the attention mechanism turns an event into 'an item of relevance' that most probably will involve some kind of internal representation upon which the agent will compute its next action(s). In other words, the attention mechanism assigns meaning to the event. It can be said therefore that challenges are internal representations of events. Metaphorically speaking, while events are 'objective' agent independent pieces of information, challenges are 'subjective' in the sense that they are agent dependent and related to a specific context. Moreover, it is initially impossible to derive the character of the challenge only from the informational content of the event that triggered it. On the same token, events that are produced by an agent do not necessarily reflect the nature of the challenge that originated them. The modeling of the framework we propose need to be such as to represent the distinct planes of events and challenges.

Another interesting point that arises here is that in a population of cognitive agents similar (but not necessarily identical) agents may form consensual domains as was already mentioned above. A simple model of a consensual domain in our case is a subset of events that triggers a group of agents to produce a subset of challenges identical for all of them[14]. It can be said that a consensual domain allows a group of agents to relate to a certain aspect of their environment consensually i.e. to share a

[13] A complex process/structure in the environment can be cognized and represented only by a cognitive agent of the same or higher complexity. This is an hypothesis that requires further investigation.

[14] Consensual domains are in fact much more complex than the given example. See (H. Maturana, 1978) for a more realistic description.

representation of their environment. Shared representations allow higher level of sharing resources, division of labor and coordination of actions among agents[15]. Nevertheless consensual domains are not necessary for cooperation among agents or the formation of coordinated coalitions.

The propagation of events should not be confused with the propagation of challenges. While the first has to do with the flow of 'objective' information, the latter has to do with the flow of meaning and influence which involves also aspects that are particular to the involved agents such as their history, values, goals etc. When it comes to systems that may allow the convergence of human and machine agents, the flow of information is only part of the picture. It is the flow of meaning and influence that is important to understand and model. When coalitions among agents are formed, new relevance emerges. This relevance is not only a product of the information flow within the coalition members, it is also a product of the way agents affect and are being affected by other agents. These affective relations that involve relevancies local to the participant agents bring forth a semantic field [16] that is associated with the coalition and is more than the sum total of the local contexts of the participant agents.

Modeling events is in fact a problem of modeling the environment. In (Heylighen, 2012a) the distinction between events and challenges is made redundant. Events are vectors that communicate the state of the environment at a given location (while location itself is not explicit) the components of the vector are independent parameters within an event space. The agent interacts with the event by modifying the vector and by that actually modifying the environment. Such modifications are communicated to other agents that do the same until a point where no further modifications take place. The event/challenge then is said to be relaxed as it does not trigger further activations. The great advantage (and also the weakness) of vector representation of events lies in its simplicity. The basic assumption is that any complex dynamics can be encoded in an N-dimensional state space. The agent (or a population of agents) is engaging this complex dynamics in a kind of cybernetic control loop: Inputs from the environment (derivations of the system's state) are fed into a network of interacting agents. Outputs or responses are fed back to the environment. The environment's inputs are considered as perturbations against which the agent operates with the goal of optimizing some internal fitness function(s). A multitude of fitness functions for various agents may of course introduce further complexity into the interaction as non-linear dependencies may arise among the interacting agents and between them and the environment. A great number of models can be formed based on this initial modeling of the environment by introducing more constraints or inner structure into the vector representation of events. The weakness of this approach lies mostly in two points. The first point is that a flat vector representation is not optimal for the description of many kinds of environments. The second has to do with scalability. High level events will be merely collections of low level events. Agents, at any level, cannot create events that do not belong to the initial challenge/event space and therefore cannot create novel descriptions and abstractions of the environment. Flat vector representation may be

[15] The higher level of sharing allows an additional level of signification that agents can share. An agent can, in addition of exchanging information, to communicate certain meanings of this information – what the information means in relation to a shared representation or context. This is a prerequisite to the emergence of linguistic exchange.

[16] Semantic field is a shared context distributed over a number of interacting agents but cannot be localized to any of them. The concept of field comes to express the non-locality of meaning.

sufficient therefore to certain families of problems but might not be general enough for a fully scalable cognitive process where the problem space is not defined a-priori. These weaknesses however can be overcome by introducing a flexible and extensible structure into the vector.

A different approach to modeling events and challenges is that each event is a piece of software code (codelets or objects) that is interpreted by the agents. The process of interpretation is part of the attention mechanism. Running the code results in figuring the relevance because the interpreter is in fact the local context of the agent. The incoming code might directly change the state parameters of the agent. Afterwards if the event is found relevant it is further interpreted to select a proper action. The selected action eventually produces further events in the form of pieces of code that are sent to other agents. This approach allows of course a higher level of interactivity between agents and their environment and among agents. The pieces of code may represent an 'ontology' underlying the environment and expressing the happening of the environment and agents in terms of ontological elements and their modifications. The very structure of agents might constitute of the same ontological elements as pieces of code that are interpreting other pieces of code. The whole framework therefore can be modeled in terms of a modifiable distributed computation that based on pre-defined ontological elements. With the prospects of the internet of things, such model of a scalable cognitive process may map very well into a global layer of sensors, actuators and other devices that communicates 'the state of the planet' at any given moment.

5.3 Modeling topological aspects

The topological modeling of the framework comes to answer the question how interacting agents within a population are connected. A connection generally designates a relation of neighborhood and direct influence between two or more agents (in the case of hyper graphs). Linked agents can affect each other directly while non-linked agents cannot affect each other directly. Still, in the second case, an indirect affect may exist via the mediation of other agents along a path that connects them. A connectivity graph can be pre- given and more or less fixed, or be reconfigurable and emergent. Links may of course have weights that affect the level of influence between the linked agents. Variable weights are one of the most classic manners of learning in connectionist models of agents (Bechtel & Abrahamsen, 1991).

According to modeling considerations of the framework, a number of mutually exclusive networks can connect a population of agents and carry distinct flows of information. Dividing the overall communication among agents into layers may be helpful to facilitate scalability in the model as each new level in the scalable cognitive process will establish a new layer of communication.

It is clear that the topological aspect of modeling may encode important properties and capacities of the overall cognitive process. In neural networks, for example, most of the network's 'intelligence' is encoded in topology while the agents are quite simple and standard processing units. In the framework proposed here, agents are generally constituted from coalitions of connected agents. How much of the functionality is encoded in the agent units and how much in the connections among agents is one of the more challenging problems of realizing the proposed framework.

At this stage, we tend towards an approach of minimum assumptions and prescribed constraints regarding any aspect of the framework. It seems therefore that the connectivity among agents will be designed to be a product of an ongoing self organization constrained by the conditions of the environment. The initial conditions to such self organizing topology will probably be a random graph or a fully connected graph but other initial configurations might be found useful as well.

5.4 Modeling coalitions and influence networks

The formation of coalitions stands at the core of our framework. A coalition brings together a number of cognitive agents to perform a significant cognitive process that cannot be carried out by any of them alone. There is a large variety of possible mechanisms of forming coalitions. Since our framework includes also the participation of human agents, at least part of the suggested mechanisms can be products of human design. But the mechanisms of interest are of course those that will be products of self organization responding to environmental signals without the intervention of a human designer. Beyond a certain level of cognitive capacity that we hope to achieve in the Global Brain, the sharp distinction between design and self organization will eventually disappear.

We start with a scenario of a given population of cognitive agents situated within an environment that produces a stream of difference events. The population is varied i.e. not all agents are identical but they share at least some similar characteristics and the capacity to interact with which other. In the most general case the agents are organized in a reconfigurable network but there are no further assumptions about the network's initial topology. Additionally, the events from the environment are not accessible to all the agents in the same manner though some may be. Modeling the formation of coalitions can be divided into a few problems:

1. Grouping – How a subset of the population of agents, depending on their structure, history, goals and current events come to identify each other and become a distinct group of interacting agents?
2. Cohesion – What are the mechanisms by which a distinct group just formed may coalesce into a more or less stable (or recurrent) entity? What are the necessary conditions for such cohesion to persist or discontinue?
3. Organization – How a coalition self organizes in terms of topology and dynamics to achieve together higher fitness of its participant members against a mutual goal, or alternatively, against a variety of goals particular to the participating members but somehow complementary [Ref: eva's scenario].
4. Measurement – In what ways a measurement mechanism and a unified set of fitness criteria may emerge within a coalition. Such mechanism and the fitness criterion need somehow to be based on the measurement mechanisms of the component agents and their fitness criteria. Measurement mechanisms and a unified set of criteria will be used to produce feedback for the cohesion and organization mechanisms, and possibly for grouping as well.

The four problems are in fact different facets of the process of self organization. Clearly if a coalition emerges, the interactions among its participating agents are such that the overall dynamics convergences to an attractor within its respective state space which is a subspace of the state space of the whole initial population. Grouping can be understood therefore as dividing this global state space into one or more regions/subspaces where further self organization is highly probable. In practice this involves reinforcing or inhibiting connections among agents. In the presence of a certain ecology of challenges, the grouping and networking of a population of agents can be the subject of a genetic algorithm where link strengths (can be 0/1 or a probability in a more complex case) are encoded into genes and groups of agents are bred according to their success in further converging into coalitions. This convergence in turn already depends on the kind of challenges or that the environment presents.

The leading hypothesis here is that the formation of coalitions is a basically a product of some kind of evolutionary process. The environment provides, as always, the system of pressures as a flow of challenges and the population of agents provides the variation in both function and structure. The selection however is not for replication but rather for cooperation. Structures as networks of interacting agents will be selected for their swift and rewarding organization for cooperation in the face of existing challenges (Heylighen, 2012b). Techniques from the discipline of artificial life can be adapted to experiment with this approach.

Another approach to model the formation of coalitions is a version of Pandemonium model (Jackson, 1987)(S. Franklin, 1995). A population of agents with a variety of capacities responds to compound events in their environment in the following manner:

1. Each agent whose attention is raised by an event creates an activation which is an output event with some parameter R (for relevance) the greater R is, the higher the relevance of the event.
2. The N agents with highest activation are selected to become a candidate coalition. The connections among these agents are reinforced in such manner that the next time a subset of the group is activated, it will suffice to activate the whole group.
3. The result events produced by this initial coalition causes other agents from the population that were not activated before to be activated (the challenges are propagated). Meanwhile, members of the first coalition cease to be active because the events that activated them in the first place are not present anymore.
4. The coalition is being updated by releasing the $n<N$ least activated agents and by adding n new agents from the most activated ones. Each time frame another such update takes place.
5. The current coalition’s activity (i.e. its coordinated response) is tested against a global fitness criterion whether its activities are successful or not.
6. If the response was successful the links between the coalition’s members are reinforced. Additionally the links between the active agents and the agents that stimulated them but no longer in the coalition are reinforced as well. If the response is not successful, the said links are weakened instead (reinforcement means that the future probability of activation along the link is higher).
7. (Optional) Agents in the population that are never activated are weeded out. The most activated agents are mutated and bred by a genetic algorithm to produce more variations of them so the whole population evolves towards higher performance.

In this approach a coalition is a more dynamic entity that eventually converges to a network which is capable of performing a certain task in response to a sequence of compound events. Once the links and activations are more or less stable, further changes can be reduced or entirely suppressed. More than one coalition can emerge from a population of agents, and each agent can participate in a few coalitions given that there are no conflicting activations. Additionally, once the triggering events from the environment change or disappear, the links will decay after a while and the coalition responsible to respond to it will be disbanded. Thus, the cognitive resources of the system are never stuck in functions that become obsolete.

6. Summary

It is argued that the prospect of the emergence of a Global Brain as a planetary level Communication Command and Control system capable of demonstrating general intelligence depends on realizing a scalable cognitive process. We have described an agent based framework for scalable cognition by first defining cognition as the combination of two selective processes: Selection for relevance (attention mechanism) and selection for effective action. These selective processes are context sensitive and operate on events that mediate differences in the state of the agent's environment. The structure of cognitive agents and the structure of the environment co-define each other and therefore co-emerge.

The framework suggests that the up-scaling of the cognitive process is realized by the agent's tendency to form cooperative coalitions. Every such coalition is in fact a super agent constructed from simpler constituent agents operating together in a collective cognitive process. Coalitions are formed and dismantled according to their relevance. The relevance of a coalition or a super agent is derived from: 1. the existence of sufficient triggers from the environment to which they respond effectively according to a context sensitive set of criteria. 2. The extent by which they influence other coalitions and participate in higher level coalitions. 3. A Decay factor that basically represents entropy and the tendency of ordered systems to disintegrate in time.

The concept of challenge is introduced into the framework as synonymous with context sensitive items of relevance that are selected by the attention mechanism. These items are analogous to the items 'brought to consciousness' in Baars' global workspace theory. As attention is spreading among agents, we say that challenges propagate within the population of agents along paths of influence that together form a network of influence. The propagation of challenges is analogous, at least in some aspects, to the monetary flow within a market system. Such flow abstracts the local context sensitive transactions and highlights instead the flow of the currency of attention. This is based on the understanding that the currency of attention 'buys' the resources necessary for effective action. Effective action in turn gains influence that further draws attention. This is quite different from the flow of information among agents because we hypothesize that what drives the formation of higher cognitive structures is the spreading of challenges and not only the spreading of information. Of particular interest in this sense is that aspect of propagation we called vertical propagation of challenges. Vertical propagation is associated with the formation of higher cognitive functions and takes place as challenges at a certain level are combined through the interactions of agents to a challenge of a higher level.

The final part of the paper explores general considerations and problems of modeling the framework and outlining initial directions for implementation. The approach is modular and divides the framework into a few modeling problems: Modeling agents (the generic agent), modeling challenges and events (modeling the environment), modeling topological aspects and finally modeling coalitions and their formation which is the most critical aspect of the framework. All these are of course subjects to further research.

The ultimate test of implementing the framework is the demonstration of general intelligence i.e. the spontaneous discovery of problems in the environment and the emergence of specific problem solving capabilities without the guidance of a designer. This is of course a very hard problem to begin with but this paper makes some conceptual headway in figuring how to get there.

7. Bibliography

Ashby, W. R. (1958). Requisite variety and its implications for the control of complex systems. *Cybernetica*, *1*(2), 83–99.

Ashby, W. R. (1960). *Design for a Brain: The Origin of Adaptive Behavior*. Wiley.

Baars, B. J. (1993). *A cognitive theory of consciousness*. Cambridge Univ Pr.

Baars, B. J. (2005). Global workspace theory of consciousness: toward a cognitive neuroscience of human experience. *Progress in brain research*, *150*, 45–53.

Bechtel, W., & Abrahamsen, A. (1991). *Connectionism and the mind: An introduction to parallel processing in networks.* Basil Blackwell.

Ben-Jacob, E. (2003). Bacterial self–organization: co–enhancement of complexification and adaptability in a dynamic environment. *Philosophical Transactions of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences*, *361*(1807), 1283–1312.

Ben-Jacob, E., Aharonov, Y., & Shapira, Y. (2004). Bacteria harnessing complexity. *Biofilms*, *1*(04), 239–263.

Ben-Jacob, E., & Levine, H. (2006). Self-engineering capabilities of bacteria. *Journal of the Royal Society Interface*, *3*(6), 197–214.

Davenport, T. H., & Beck, J. C. (2001). *The Attention Economy: Understanding the New Currency of Business*. Harvard Business School Press.

De Landa, M. (2011). *Philosophy and simulation: the emergence of synthetic reason*. Continuum Intl Pub Group.

Edelman, G. M. (1987). *Neural Darwinism: The theory of neuronal group selection.* Basic Books.

Edelman, G. M. (2004). *Wider than the sky: The phenomenal gift of consciousness*. Yale Univ Pr.

Franklin, S. (1995). *Artificial Minds*. Cambridge, Massachusetts: MIT Press.

Franklin, S. (2006). A foundational architecture for artificial general intelligence. *Advances in artificial general intelligence: Concepts, architectures and algorithms, proceedings of the AGI workshop 2006* (pp. 36–57). Amsterdam: IOS Press.

Franklin, S. (2008). An Ontology For Cognition. Institute for Intelligent Systems, The University of Memphis. Retrieved from http://ccrg.cs.memphis.edu/tutorial/PDFs/AnOntologyforCognition.pdf

Franklin, S., Kelemen, A., & McCauley, L. (1998). IDA: A cognitive agent architecture. *Systems, Man, and Cybernetics, 1998. 1998 IEEE International Conference on* (Vol. 3, pp. 2646–2651).

Friedkin, N. E., & Johnsen, E. C. (1999). Social influence networks and opinion change. *Advances in Group Processes*, *16*(1), 1–29.

Hebb, D. (1968). *0.(1949) The Organization of Behavior*. Wiley, New York.

Heylighen, F. (1991). Cognitive Levels of Evolution: pre-rational to meta-rational. *The Cybernetics of Complex Systems-Self-organization, Evolution and Social Change*, 75–91.

Heylighen, F. (1997). Towards a global brain: integrating individuals into the world-wide electronic network. *Der Sinn der Sinne* (pp. 302–318.). Göttingen: Steidl Verlag. Retrieved from http://pespmc1.vub.ac.be/papers/GBrain-Bonn.html

Heylighen, F. (1999). The growth of structural and functional complexity during evolution. *The evolution of complexity*, 17–44.

Heylighen, F. (2002). The global brain as a new utopia. *Zukunftsfiguren*. Suhrkamp, Frankurt. Retrieved from http://pespmc1.vub.ac.be/papers/GB-Utopia.pdf

Heylighen, F. (2008). Accelerating socio-technological evolution: from ephemeralization and stigmergy to the global brain. *Globalization as evolutionary process: modeling global change* (Vol. 10, p. 284). Retrieved from http://pespmc1.vub.ac.be/Papers/AcceleratingEvolution.pdf

Heylighen, F. (2011). Conceptions of a Global Brain: an historical review. *Evolution: Cosmic, Biological, and Social, eds. Grinin, L. E., Carneiro, R. L., Korotayev A. V., Spier F.* (pp. 274 – 289). Uchitel Publishing. Retrieved from http://pespmc1.vub.ac.be/Papers/GBconceptions.pdf

Heylighen, F. (2012a). Challenge Propagation: a new paradigm for modeling distributed intelligence. Global Brain Institute. Retrieved from http://pespmc1.vub.ac.be/Papers/ChallengePropagation.pdf

Heylighen, F. (2012b). Self-organization in Communicating Groups: the emergence of coordination, shared references and collective intelligence. *Language and Complexity*. Barcelona University Press. Retrieved from http://pcp.vub.ac.be/Papers/Barcelona-LanguageSO.pdf

Jackson, J. V. (1987). Idea for a Mind. *ACM SIGART Bulletin*, *101*, 23–26.

Kempe, D., Kleinberg, J., & Tardos, É. (2003). Maximizing the spread of influence through a social network. *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining* (pp. 137–146).

Maturana, H. (1978). Cognition. *Wahrnehmung und Kommunikation*, 29–49.

Maturana, H. R. (1975). The organization of the living: a theory of the living organization. *International Journal of Man-Machine Studies*, *7*(3), 313–332.

Montgomery James. (2011, February 23). Mathematical Models of Social Systems. Retrieved from http://www.ssc.wisc.edu/~jmontgom/influencenets.pdf

Seeley, T. D., Visscher, P. K., Schlegel, T., Hogan, P. M., Franks, N. R., & Marshall, J. A. R. (2012). Stop signals provide cross inhibition in collective decision-making by honeybee swarms. *Science (New York, N.Y.)*, *335*(6064), 108–111. doi:10.1126/science.1210361

Simon, H. A. (1962). The architecture of complexity. *Proceedings of the American philosophical society*, *106*(6), 467–482.

Simon, H. A. (1980). Cognitive science: The newest science of the artificial. *Cognitive science*, *4*(1), 33–46.

Turchin, V. (1977). *The Phenomenon of Science. A Cybernetic Approach to Human Evolution*. New York: Columbia University. Retrieved from http://www.google.be/url?sa=t&rct=j&q=turchin%20phenomenon%20of%20science&source=web&cd=3&ved=0CDgQFjAC&url=http%3A%2F%2Fpespmc1.vub.ac.be%2FPoS%2FTurPOS-prev.pdf&ei=HddVT5fwDcrd8QPz9Yn7CA&usg=AFQjCNELmVRHRRqciv5xi7sbtRQzswa6bg

Turchin, V. (1995). A dialogue on metasystem transition. *World Futures*, *45*, 5–57.

Vidal, C. (2008). What is a worldview? *De wetenschappen en het creatieve aspect van de werkelijkheid*.

von der Malsburg, C. (1995). Self-organization and the brain. *Vis. Res*, *33*, 85–90.

von Glasersfeld, E. (1997). Distinguishing the observer: An attempt at interpreting Maturana. *Towards an Ecology of Mind.[Originally published as: Die Unterscheidung des Beobachters: Versuch einer Auslegung. In: Riegas, V., Vetter, C.(Eds.) Zur Biologie der Kognition. Suhrkamp, Frankfurt, 1990, pp. 281-295]. http://www. oikos. org/vonobserv. htm*.

Wikipedia contributors. (2012, January 14). Binding_problem. *Wikipedia, The Free Encyclopedia*. Wikimedia Foundation. Retrieved from http://en.wikipedia.org/w/index.php?title=Binding_problem&oldid=471301667