Sensemaking as a way to manage complexity: does it extend to artificial agent organizations?

Matteo Bonifacio, Diego Ponte
ITC-irst. - University of Trento (Italy)
Via Sommarive 18 Povo - 38050 (TN) - Italy
{bonifacio, ponte}@itc.it

Abstract

Designers develop agent-based organizations with the goal to implement systems able to operate in complex environments. The distinctive quality of these systems is that agents are autonomous i.e. they can face complexity through their decision capabilities in order to make it more flexible. This paper proposes that the concept of complexity currently used in artificial intelligence is weak if compared with that elaborated in organizational sciences. As a result, current agent organizations cannot completely face all the range of complex situations. Borrowing some seminal contributions from organizational studies, a broader definition of complexity is presented introducing the notion of ambiguity, and its impact both on reasoning models and organizational design. Moreover, we present some major implications in the design of artificial systems and organizations.

Introduction

The analysis and development of an artificial system starts from the definition of those requirements that are needed in order to match its goals. A major and preliminary one is the definition of the level of environmental complexity it has to face, since such level has a direct impact on the complexity of the system. In fact, according to the principle of requisite variety, in order to be adaptable to its environment, a system has to accept an internal level of complexity positively correlated to the external one (Ashby 58). From this perspective, the complexity of task implies the complexity of the capability.

Of course, such statement seems quite obvious if we think about complexity in quantitative terms. For instance, if we think about complexity as the variety and variability of events, then increasing levels of complexity require to design systems made up of more various and variable components. But such conclusion is not obvious, since, as we will see, complexity can be of different qualities requiring, as a consequence, qualitatively different system's modeling approaches.

A qualitative jump of this kind has already occurred in the design of artificial systems. Is well known, in fact, how the passage from simple environments to complex ones, required to qualitatively step from a hierarchical design principle, to one that focuses on delegation. That is, more complexity could not be managed with more complex hierarchies, while new organizational design principles where introduced, such as the one of autonomous and interacting components. In such case, the designer, instead of trying to recognize all the possible events that the system will face and implement ex-ante proper responses, he delegates their identification and management directly to each system's component. In the first approach the designer has the very tough mission to design an all-comprehensive solution foreseeing each possible state and behavior of each component. The second approach proposes that the system tries to manage complexity itself assigning goals to components, and delegating the modality of their execution. In this way, great importance assumes the concept of software agent and of goal oriented reasoning.

The intention of this new computational paradigm is to develop artificial organizations of software components able to face complex, dynamic and uncertain environments (Georgeff et al. 99). The distinctive quality of software agents is their autonomy, that is, the ability of each agent to perform an endogenous decisional process able to identify, run time, the best course of action (Barber and Martin 01). Thus, software agent organizations should be able to manage the gap between the complexity foreseen by the designer and the one faced in run time environments. Interestingly, such idea seems intuitively near to the notion of agency proposed in organization sciences by Arrow: the more the task is complex, the more it must be delegated to autonomous (human) agents (Arrow 85).

Again, if we where to think about complexity in quantitative terms, than if complexity increases, agents will have to be more autonomous in deciding courses of action. But, if we accept the possibility that there are qualitative different levels of complexity, then the question is whether there are other types of complex situations and, consequently, if current notions of agent, reasoning, and organization, are appropriate. Said differently, are artificial organizations designed on the principle of goal oriented agents able to qualitatively cope with all the possible ranges of complexity?

Going back to Arrow, such question seems to find a provoking answer if we consider that agency relationships are of different kinds depending on the complexity of the task. There are less complex tasks in which the principal can control the agent's goals and behaviors (for example, manual execution), while there more complex types of relationships in which the principal can control just the goal level and not the execution level (for example, commercial activity). It is easy to notice that the former seems quite near to the hierarchical design principle, while the latter to goal autonomous agents. However, as noted by Arrow, there are situations in which also the identification of goals must be delegated. In such cases, agents decide almost everything, putting under discussion the very notion of organization.

Such productive link between AI field and organization sciences, although originally vital in the famous work proposed by Herbert Simon, seems to be increasingly dry and forgotten (Simon 84). In fact, if we review AI literature, it seems quite evident that the last accepted notion of complexity refers to the one of uncertainty and, as a consequence, the last accepted model of goal oriented reasoning (or practical reasoning) refers to the one of bounded rationality and reasoning under uncertainty. Such circumstance is particularly unfortunate if we consider that after Simon, his major scholar, James March, started a whole new path in the analysis of complex situations and consequent rational models, aimed at overcoming the boundaries of uncertainty and address the qualitatively different notion of ambiguity (March and Olsen 1976; Zack 1999). Nonetheless, such amnesia is perhaps due to a reasonable epistemological explanation; in Kuhnian terms (Kuhn 96), the concept of agent, rationality, and organization that emerges from March's inquiry is seemingly incommensurable with that positivistic paradigm that still seems to entail AI research. In fact, as we will see afterwards, such passage is rather a qualitative jump than a quantitative improvement; it literary inverts our classic notion of what is a rational agent, and what an agents organization means.

As we believe that innovation lies in the explorative encounter of different systems of thinking, we're definitively convinced that reestablishing some bridge among challenging organization science developments and AI, could be somehow a quite provoking and interesting experiment. As in the case of Arrow, organizational studies can bring into the AI field new and even ingenuous perspectives that could suggest unforeseen alternative ways to interpret traditional problems. At first, as cleverly said by Bowker and Star (Bowker and Star 00), such exercise may generate "conceptual monsters" like the one here proposed of agents that before act, and afterwards seek for a plausible justification for their choice. Nonetheless, as it typically happens along scientific development (Di Trocchio 98), such "monsters" represent a boundary object upon which new concepts can be negotiated, developed, and than become "normal". Goal of this paper is exactly to through into the AI field some organization science concepts that, due to their distance from the current paradigm, may hopefully be seen as "conceptual monsters". As such, they are quite probably useless and unable to enhance concrete implementations. Nonetheless, if they are considered as they are, that is monsters, we believe they could generate some provocative reflection.

In particular, we will focus on the sensemaking approach to complex situations. The key issue of sensemaking is that it views complexity differently from the classic notion used in AI. In fact, it views complexity not only in the sense of information's lack (uncertainty), but rather in the sense of meaning's lack (ambiguity). As a method to face this last notion of complexity, sensemaking defines a decision making approach radically different from that used in classic software agents systems. As a consequence, we will show how sensemaking and ambiguity entail a different perspective on the concept of organization. Finally, we propose some ideas on how these concepts may contribute to the development of artificial systems.

Current agent-based organizations paradigm

Current research in AI is focusing on the production of software able to make decisions in complex situations. These systems are considered useful because, through their decision capabilities, they are able to handle that part of complexity that could not be foreseen by the human designer. Several examples are used to display such systems but perhaps the most known is that of the air traffic controllers, i.e. the software controller systems responsible for the take-off, the landing of airplanes and, more in general, the management of an airport (Rao and Georgeff 95).

To better understand the capabilities of software agents, it is useful to consider the notion of environmental complexity subsumed by such paradigmatic and so often cited example. Complex environments are those where there are (Russel and Norvig 95):

- Inaccessibility to updated information: there is a degree in which information is not available;
- Environmental indeterminacy: there is a degree in which actions have unpredictable effects;
- Interrelation of effects: different events can be influenced by other events;
- Environmental dynamicity: the environment change dynamically due to events independents form agent willingness;
- Discrete or continuous environment: the degree trough which environmental changes are considered is influenced by the quantity of information to consider.

Is to be underlined how, all these elements, entail a notion of environmental and organizational complexity in the sense of uncertainty. That is, information to make decision lacks or, on the contrary, there is so much information that the agent does not have the necessary storing and computational resources to effectively and efficiently process it. To face this kind of complexity, software agents are designed according to the so called practical reasoning process, which is one of the most important decision making processes used in artificial intelligence (Bratman 87). In order to keep alive our linkage to organization sciences, such process finds its root in one of the "ideal types" of rational behaviors stated by the famous sociologist Max Weber (Weber 78); that is, the ideal type of "instrumental rationality" now evolved into the notion of prospective rationality. This form of rationality notion, so similar to the practical reasoning model, was deeply analyzed in decision making theories, and it views choice as a process characterized as follows (March 94):

- There is a set of predefined (given) preferences. They state what is "good" and what is "bad" (evaluation criterion). Preferences can be viewed as values used by an agent to judge and rank facts and events

according to their desirability, and, as a consequence, to make decisions. Such preferences are by definition considered as given, that is, they are independent variables exogenous to the agent reasoning process (the agent doesn't reason about preferences).
- Goal to achieve. They represent the state of affairs that maximizes benefits and minimizes costs (min-max rule). Thus, they are selected according to the given set of preferences, and to the information collected from the environment. The goal, differently from preferences, is endogenous i.e. a dependant variable which is a function of available information and preferences.
- Means useful to reach goals. After the selection of goals a plan must be drawn. In AI means to achieve goals are usually called actions, and are rooted in capabilities. According to the prospective model, actions generate a cost that impacts on the net benefit of goals (although, costs are often not explicitly considered in practical reasoning)

This decision making process is definitely the correct in a complex environment characterized by uncertainty. In such situation, according to the bounded rationality model (Simon 82), information is limited and its acquisition, such as every other possible action, generates a cost. That is way the major problem of such model is to determine when information acquisition must be stopped, that is, configuring the choice as characterized by risk (not reducible informative gap) rather than uncertainty (reducible informative gap). In other words, the reduction of uncertainty must find a limit in some threshold that defines the boundary between what is the acceptable risk, and what is the uncertainty that needs to be managed. Such threshold, in economic studies, is typically called "risk propensity".

In agents terms, these principles leads to a model by which reasoning happens according to two main moments: given a rule that determines whether system's variations are acceptable or not, if the boundary is not overcome, the agent continues to commit to his current goal, while on the contrary, he has to revise his beliefs analyzing the environmental feedback in order to control if the plan must be changed or the goal must be reset. In such sense, the agent is goal autonomous. In bounded rationality terms, the agent "has" a risk propensity (the threshold) that determines the costs it is willing to pay (for example, in terms of decision speed) in order to acquire additional information and reduce risk. If the system's configuration passes such risk line, than the agent starts to reduce uncertainty recalculating plans and goals that is, reformulating courses of action on the base of additional information. One of the main applications of such model is the one implemented by the so called BDI agents (Belief-Desire-Intentions) (Rao and Georgeff 91).

As a corollary still derived from bounded rationality, achieved solutions may be local optima, given that the commitment to a course of action may lead to goal states that where not optimal if compared to those that would be set if more information would be available (that is, if the agent would be more risk taking). Of course, such possibility is acceptable by definition in practical reasoning, which exactly aims at balancing performance with cost. In the case of the air control system, this is intuitively clear if we think that a continuous re-planning of landing schedules may keep in the air planes for an indefinite amount of time. In agents terms, goal autonomy implies that, given the same environment and capabilities (and implicitly, preferences), two agents may pursue different local goals according to the information they have, and moreover, to their risk propensity.

In epistemological terms, the following features characterize environments in which agents should display such type of behavior (March 94):

- Objective reality: there is a reality that entails meaning, that is, the correct meaning of a fact "is", in theory, available in the environment. As a consequence, partial perspectives that agents can have on the same environment is due to their economic or computational limitations in retrieving all the needed information.
- Cause-effect chain: events are related to each other by objective cause effect relations. Each event has his predecessors and has its own consequences. Actions and effects in the system are exogenously determined.
- Probability: in uncertain situations, agents can predict evolutions of current actions through a probability calculus. Agents act in a way that is consistent to such predictions.
- Intentionality: decisions are tools to pursue intentions. In logical terms, agents perform decisions before acting, and actions are intended as means to achieve goals.

In organizational terms, the bounded rationality perspective leads to consequences that are again similar to those postulated by current multi agent systems. On the one hand, since complexity is resolved as a problem of information processing, it is useful to create artificial agents able to make this in place of humans as far as artificial agents are more capable to store and process information. On the other hand, bounded information processing capacity leads to approach complex tasks through the division of decisional labor. That is, agents are organized in information processing organizations that have the exact goal to reduce information overload on each actor through an efficient distribution of information processing tasks. Borrowing from the classic work of Hayek (von Hayek 48), when the task is complex and composed of context dependant sub tasks, is more efficient to locally delegate its resolution while providing shared mechanisms of coordination. As we will show in the following section, from a broader perspective such delegation is heavily partial since it assumes that agents can decide on everything besides what they believe is desirable. That is, they freely decide goals and actions, while delegate to some designer (human, in the case of artificial agents, or metaphysical in the case of humans) the criterion through which these are considered as worth to be pursued.

Complexity as ambiguity

While such notion of rationality is held as normatively correct from an AI perspective, the evolution of organizational studies has heavily revised this model and its underlying assumptions. In fact, the definition above proposed does not cover entirely the concept of complex situations used by a growing school of organizational scholars (Weick 95). These scholars consider the concept of complexity not only in the sense of uncertain situations. They include the concept of ambiguity as a quality describing environments in which the meaning of things, events and objects is not clear.

Such environments are those where, for instance, the interaction mechanisms between agents are extremely rich (individuals may act in response of their different roles), where interpretations of facts are local and contingent, and where an entity controlling the entire system does not exist (Cilliers 98). Often, in these cases "individual elements are ignorant of the behavior of the whole system in which they are embedded" (Lyotard 79). In really complex systems things happen simultaneously, technologies changes very frequently and, as stated by March, problems, solutions, opportunities, ideas are so much mixed that it is hard to interpret the connections between them (March 94).

As a consequence, ambiguous situations are those in which there is a lack of meaning since there is no sufficient information in order to formulate a unique frame according to which events can be interpreted and probabilities formulated. In such cases, an agent can formulate alternative and even conflicting interpretative frames and, according to these, information can be interpreted in different (and plausible) ways. For instance, if there is either an impossibility to establish a unique judgment standard to compare consequences, or cause-effect relationships (what leads to what) are not so clear, different interpretations of the data are then available. Hence, an apparently clear situation (in particular from an ex-post and external perspective) can have more than one possible interpretation and, therefore, more future scenarios become plausible.

It is important to underline that ambiguity, differently from uncertainty, presumes a different epistemological and ontological view. Here meaning, rather than lying in the structure of the world, resides in the structural coupling between the agent and the environment. Borrowing from constructivism (Berger and Luckman 66), and structuration theory (Giddens 91), the actor is both a product and a producer of reality, that is, the environment constraints the agent's behaviors while these are able to shape environmental constraints. From this perspective, ambiguity qualifies situations in which there "are" multiple plausible readings of the same context, in the sense that each reading can implement, through action, such possible configuration. In epistemological terms, knowing alternatives is neither a passive acquisition of information about reality (description) nor an idealistic production of structure (prescription), but rather an ongoing interaction between the agent and the environment. In ontological terms, there is not an a priori correct interpretation of the world, but rather a continuous accomplishment of a possible coupling between a decision maker and a decisional context. From a decision making perspective, while uncertainty presumes that there is an ideal interpretative frame (or utility function) through which alternatives can be ranked, ambiguity presumes that successful decisions are those that, because of interpretative framing, were able to shape a favorable decision context[1].

For our purposes it is to be underlined that, while uncertainty and risk assume an ontologically exogenous preference function, that is, a given frame of reference, ambiguity considers framing as endogenous to the decision making process. In other words, the agent cannot assume a given frame (such as, typically, the preference function or risk propensity) but rather the latter can be an outcome of the former. Therefore, a decision under ambiguity can be seen as doubled faced processes of interpretation (choosing a frame) and consequent selection of an alternative (choosing a course of action).

Following considerations of these scholars it is worth to revise the beforehand notion of complexity with other characteristics (March 94):

- Lack of a unique and predefined meaning of the environment. This implies that environmental facts need an interpretation. Data and feedback of the world are interpretable and are prone to possibly different explanations.
- Lack of a unique cause-effect relationship. The interpretability of the environment allows different cause-effect chains. In fact cause-effect chain changes according to the "point of view" used by the decision maker.
- Lack of intentionality: not all behaviours of agents are the result of a decision processes; they can be the consequence of rules, or the results of casual and not related events, etc.

While the effects of uncertainty and risk on decision making are well known in the AI domain, probably less attention has been paid on how decisions are affected by ambiguity.

Ambiguity and its influence on an agent's reasoning model

In ambiguous situations, the normative decision making process is unable to provide a viable solution due to the lack of the necessary conditions to make a clearly and *a priori* rational choice. As Weick declares, a possible

[1] As proposed by Zack, ambiguity, equivocality, uncertainty and risk can be seen as logically related situations. In particular, an ambiguous situation (infinite interpretations) is reduced to an equivocal one (finite interpretations) and, within a selected frame, uncertainty is managed acquiring information until risk is acceptable (Zack 99).

way to face and reduce ambiguity in decision making is through the sensemaking process:

> To talk about sensemaking is to talk about reality as an ongoing accomplishment that takes form when people make retrospective sense of the situations in which they find themselves. There is a strong reflexive quality to the process. People make sense of things by seeing a world on which they have already imposed what they believe. (Weick 95)

The concept of sensemaking contains two main logical moments: the retrospective process of interpretation (sense) and prospective one of enacting a new configuration of the environment (making). The starting assumption is that rationality, rather than being a process aimed at keeping consequences coherent to preferences, is more weakly oriented towards either keeping coherency between consequences and preferences, that is, adapting the former to the latter or the opposite. On this regards, the concept of rationality is the one proposed by March of appropriateness; agents are aimed at matching identities (expressed, for example, by preferences) and situations (such as decisional contexts), either modifying their identities or the situation the face (March 94). Consequently, when agents face a situation that interrupts the normal flow of significance (a fact that generates a cognitive dissonance (Festinger 57)) between expected and manifested consequences, they need to re-establish appropriateness between their identities (preferences) and situations (consequences). If the situation they face is ambiguous, before taking action in order to re-establish a correct environment (prospectively), agent try to interpret what they have done in the past (retrospectively) seeking for plausible reasons and motivations. In fact, since there is no predefined meaning "attached" to facts, the sensemaking process has the objective to create meaning and build context to understand unexpected situations.

Such process needs to handle to some independent variable; as proposed by Weick, since preferences are not fixed, this handle is provided by the past. That is, through the observation of the past, subjects can give an order to current events to give them significance. As clearly stated by Weick, what decision makers "see" is heavily influenced by what they believe[2], and such perception actively shapes the way an environment is understood in order to take consequent action (for instance the cause-effect relations or the judgment standard to use). From this perspective, sensing is an action since it's more a process of meaning "imposition" rather than a mere one of meaning perception; it aims at acting coherently instead of passively contemplating.

However, the very importance of the sensemaking process goes beyond the simple clarification of a situation. In fact, the sensemaking process influences what and how data is to be considered and judged by the subject in order to implement future actions. In this sense, Weick proposes the concept of enactment used to describe the way an agent alters and changes its environment through committed action (Weick 79). In ambiguous situations, once an interpretation is formulated, decision makers are able (have the opportunity) to enact an environment favorable to such interpretation. In fact, due to the ambiguity inherent the situation, an agent is able, through his commitment, to work upon the environment rather than simply adapt to it. In epistemological terms, once a decision maker has chosen an interpretative frame, he can influence and manipulate the environment (the reality) in order to establish, *a posteriori*, the truth conditions of his beliefs (self fulfilled prophecy):

> In enactment, people actively construct the environments which they attend to by bracketing, rearranging, and labeling portions of the experience, thereby converting raw data from the environment into equivocal data to be interpreted. In selection, people choose meanings that can be imposed on the equivocal data by overlaying past interpretations as templates to the current experience. Selection produces an enacted environment that provides cause-effect explanations of what is going on. In retention, the organization stores the products of successful sensemaking (enacted or meaningful interpretations) so that they may be retrieved in the future. (Choo and Bontis 02)

Summarizing, ambiguous situations are those in which different feasible interpretations are ontologically available, allowing the decision maker, therefore, to pursue different future scenarios. Since each interpretation implies a different evaluation about the decision context, there is no a unique self-evident rational action to follow. From an ex-ante perspective, such evaluations are neither correct nor wrong; each could become "true" depending on the actions that the decision maker will commit to. That is, the rationality of a decisional behavior can be judged only *a posteriori*, when consequences appear and manifest in the world. More precisely, when facing ambiguity, decision makers will implement a double-faced cognitive strategy: at first, they adopt a retrospective reasoning mode in order to generate a plausible interpretation of the decisional context (belief-driven sensemaking), and, in a second moment, they commit their courses of action in order to enact consequences consistent to such evaluation (action-driven sensemaking). That is, they seek for an explanation of what they did in the past and, consequently, commit their future actions in order to ontologically "fix" such explanation in the world.

[2] "Your beliefs are cause maps that you impose on the world, after which you see what you have already imposed." (Weick 79)

Changes of complexity management in organizations

Can we still talk about decisions when preferences are manipulated in order to justify past actions, evaluation standards are debatable and ill defined, cause-effect relationships are multiple and conflicting, consequences are opaque and interpretative frames intrinsically vague? Furthermore, if organizations are defined as forms to divide decisional labor, can we still talk about organizations?

Such view led some authors to a sort of post modernist criticism to the very notion of decision, stating that decisions are dead. In organizational terms, organizations are proposed as "garbage cans" in which solutions seek to find plausible problems, and decisions look for decision makers (Cohen et al. 72); moreover, such couplings are almost casual, often generated by the occasional circumstance of what issue is written on a meeting agenda, or which decision maker is available at the moment. In such sense, decisions represent our "western" way to celebrate modern identity, that is, the ideal belief that we are intentional actors aimed at achieving goals through the engine of rationality. The step that leads to state that, if decisions are dead than organizations are dead, seems quite short. Such step is done if we very intuitively consider organizations as built upon the division of decisional labor.

Interestingly, such critics have more recently given a chance to the survival of organizations (March 97), emphasizing how decisions are tools to collectively elaborate meaning, and celebrate a committed coordinated action. That is, decisions are primarily a moment through which organizations make sense of their environment and, once a plausible sense is produced, they build and celebrate that form of "collective trust" through which a socially constructive behavior can be sustained (Daft and Weick 84). In short, decisions are needed in order to sustain a collective faith in the successful implementation of retrospectively generated beliefs[3].

Such opportunity has been collected by Karl Weick and his concept of sensemaking, which brings the relation between decisions and meaning to its normative evaluation. In fact, it is important to note that the sensemaking process is not formless:

> The model also implies that the greater the interplay between the processes of sensemaking, knowledge creating, and decision making, the greater the organization's capacity to detect threats and opportunities, create valuable knowledge, and act on new knowledge. This interplay is necessarily fluid and open-ended, but it is not entirely random or without structure. (Choo and Bontis 02)

The sensemaking organization is the one in which multiple perspectives are seen as opportunities rather than threads, semantic redundancy is considered as a way to promote innovative behaviors able to capture spaces of environmental manipulation, retrospective reasoning is a form of learning that aims at exploiting value out of failures, and evaluation standards are discussion objects around which new notions of what is virtuous rather than vicious are negotiated:

> This means that we must define organization in terms of organizing. Organizing consists of the resolving of equivocality in an enacted environment by means of interlocked behaviors embedded in conditionally related processes. To summarize these components in a less terse manner, organizing is directed toward information processing in general, and more specifically, toward removing equivocality from informational inputs. (Weick 79)

Rather than being centrally controlled and coordinated, agents of a sense making organization have the opportunity to act autonomously creating meaning of the environment with other entities. Furthermore, in contrast with the classic notion of complexity, agents do not only decide "how to do something" (actions) or "what is to be done" (goals), but rather and more importantly, they define "what earns to be done" (preferences). Such process does not happen according to another decisional meta-level (such as a theory that determines preferences), but rather occurs as an inverted process through which, assuming actions and goals, an agent formulates appropriate preferences. As we show elsewhere (Bonifacio et al. 02), such apparently irrational behavior can lead to notice unforeseen opportunities from unexplainable failures, innovate the common sense notion of what is achievable, and discover new ways to reuse and combine existing capabilities.

Applications on artificial organizations

The question of whether is useful to design artificial systems according to a model such as the one here proposed, finds an answer in its very definition; that is, are there artificial systems characterized by ambiguity? Are there informative environments where meanings are opaque, where solutions are to be constructed, and evaluation standards are lacking?

In order to give a response, it could be useful to consider how one of the main challenges in systems design is exactly the one of providing a clear definition of meaning[4]. Moreover, clear semantics is exactly a

[3] Decision making shapes meanings even as it is shaped by them. A choice process provides an occasion for developing and diffusing interpretations of history and current conditions, as well as for mutual construction of theories of life. It is an occasion for defining virtue and truth, discovering or interpreting what is happening, what decision makers have been doing, and what justifies their actions. (March 97).

[4] Works on ontologies aims properly at define actors, objects and their relationships in a unique sense. This allows ontologies to be

prerequisite for artificial systems to work, and the designer usually does such work manually and off line. If this is true, the questions then become whether semantic disambiguation is needed because current artificial systems are able to cope only with uncertainty, or, on the contrary, if artificial systems that can cope with uncertainty are the only one needed since semantic disambiguation is a non problematic task. In the latter case, our conceptual "monster" is probably condemned to remain so, since it aims at managing problems that do not exist. In the former case, the answer is quite opposite; that is, our model can be seen as a plausible way to extend the boundary of what can be artificially handled while reducing the space of those tasks that are hand made, and thus, time consuming and imprecise.

However, that represent also moreover a qualitatively substantial limitation in the sense that current systems are inflexible by definition, since each time circumstances require a change in semantics, such change cannot be recomputed run time. In short, if semantic disambiguation and adaptability is an issue in artificial systems, either within the software environment or within the designer's head, than current models are exactly unable to cope with the primary goal they are designed for, that is, being able to properly manage the complexity implied in the task. So, if complexity, in the sense of ambiguity, exists, than at the moment, rather than being fulfilled by the artificial, is still managed by the human.

It clearly emerges that our view leads directly to think about artificial agents as supporting, at run time, meaning generation, revision and negotiation. Such view may seem grounded more in fantasy rather than in real life. And maybe it is. But before saying so, we think it could be worthwhile to consider whether ambiguity is an issue in the design of artificial systems, whether is reasonable to think that ambiguity is managed by humans (us) as a sense making process, and whether is plausible that a software can be, on the one hand, a retrospective decision maker, and, on the other, a manipulator of its social environment.

Between the others, AI researches could take advantage in exploiting this new perspective in different directions. First, the approach to decision making may be very useful to design more powerful decision making systems capable of recognizing different environmental characteristics (those related to the definition of ambiguity) that current systems do not consider. In particular planning and goal-oriented decision making analysis can be reconsidered to think decisions not only under the notion of "best response" to the environment, but as tools to manipulate the environment.

A major consequence is that agents can be considered really autonomous in the sense that they are able to decide not only action plans and goals, but also evaluation criterions (such as risk propensity, preference functions, and interpretative frames). Such feature, which is not allowed by definition in the classic paradigm and subsequent practical reasoning models, is here considered as a perfectly acceptable decisional momentum. Moreover, preference formulation is here explained as a process that can handle to past experiences, courses of action, and sunk costs (Bonifacio et al. 02). This type of autonomy is, of course, useless in uncertain environments while in ambiguous ones is the only chance to maintain flexible behaviours.

Finally, the way in which organizational ontologies[5] are created can be viewed as an endogenous process. Typically, in fact, ontologies are assumed as a premise for system operability. From our perspective, since ontologies are plausible and shared interpretations of the world, they can be better seen as the exact outcome of agents negotiating meaning in order to make sense of ambiguous environments. This allows the system to be real-time flexible and self-adaptable since its components are able to question themselves about the reasons why they are doing what they are doing, rather than just doing it.

Conclusion and future work

As obvious for an explorative contribution such as this, the type and quantity of future work is much more than what can be said in a final paragraph. To some extent, the article itself is a proposal for future work. Here we briefly trace some major research lines that can be fruitfully explored in the near future:

- Formalize a model for sensemaking agents, providing primitives and reasoning methods in order to sustain retrospective reasoning and enacting capabilities.
- Explore the semantic dimension of AI societies, that is, analyze if and to what extent interpretative processes can be artificially implemented in order to sustain run time semantic disambiguation.
- Detect whether artificial systems typically face ambiguity, analyzing how such ambiguity is currently managed through existing systems and methods, while producing paradigmatic cases that could drive research and design.

Finally, we believe it could be useful to pursue a continuous contamination process among AI and organization sciences, paying attention not to filter those contributions of the other that seem incompatible with current paradigmatic assumptions. In fact, we frequently notice organization science concepts that could fruitfully provoke insights and reflections to AI researchers (and of course, we presume the opposite). Such concepts should be thrown into the other one's arena, hopefully generating monsters able to capture, at least, fantasy and attention.

executed by software applications (Davies, Fensel, & Van Harmelen, 2002)

5 "ontologies are a consensual agreement on the concepts and relations characterizing the way in which knowledge in a domain is expressed" (Mentzas, 2002)

References

Arrow K. 1985. *The economics of agency*. In Pratt J. and Zeckhauser R. eds. Principals and agents: The Structure of Business. Harvard University Press.

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

Barber, K. S.; Martin C. E. 2001. *Autonomy as decision-making control*. In Intelligent Agents VI. Springer-Verlag.

Berger P.; Luckman T. 1966. *The Social Construction of Reality: A Treatise in the Sociology of Knowledge*. Anchor Books.

Bonifacio, M.; Bouquet, P.; Ferrario, R.; Ponte, D. 2002. *Rationality, Autonomy and Coordination: the Sunk Costs perspective*. In Proceeding of Third Int. Workshop on Engineering Societies in the Agents World. Springer.

Bowker, G. C.; Star, S. L. 2000. *Sorting Things Out: Classification and its Consequences*. The MIT Press.

Bratman, M. E. 1987. *Intentions, Plans and Practical Reason*. Harvard University Press.

Cilliers, P. 1998. *Complexity and Postmodernism: Understanding Complex Systems*. Routledge.

Choo C. W.; Bontis N. eds. 2002. *Strategic Management of Intellectual Capital and Organizational Knowledge*. Oxford University Press.

Cohen, M. D.; March, J. G.; Olsen, J. P. 1972. *A Garbage Can Model of Organizational Choice*. Administrative Science Quarterly 17(1):1-25.

Daft, R. L.; Weick, K. E. 1984. *Toward a model of organizations as interpretation systems*. Academy of Management Review 9(2):284-295.

Di Trocchio, F. *Il genio incompreso. Uomini e idee che la scienza non ha capito.* Mondadori. 1998.

Davies, J.; Fensel, D.; van Harmelen, F. 2002. eds. *Towards the Semantic Web: Ontology-Driven Knowledge Management.* John Wiley & Sons.

Festinger, L. 1957. *A Theory of Cognitive Dissonance*. Stanford University Press.

Georgeff, M.; Pell, B.; Pollack, M.; Tambe M. and Wooldridge, M. 1999. *The belief-desire-intention model of agency*. In Intelligent Agents V: Agent Theories, Architectures, and Languages. Springer.

Giddens, A. 1984. *The Constitution of Society*. Polity Press.

von Hayek, F. 1948. *Individualism and Economic Order*. Chicago University Press.

Kuhn, T. S. 1996. *The Structure of Scientific Revolutions*. University of Chicago Press. 3rd edition.

Lyotard, J.F. 1979. *The Postmodern Condition: A Report on Knowledge*. University of Minnesota Press.

March, J. G.; Olsen J. P. 1976. *Ambiguity and choice in organizations*. Bergen.

March, J. G. 1994. *A primer on decision making : how decisions happen*. The Free Press.

March, J. G. 1997. *Understanding how decisions happen in organizations*. In Shapira Z. eds. Organizational Decision Making. Cambridge University press.

Mentzas, G. 2002. *Knowledge asset management: beyond the process-centered and product-centered approaches*. Springer.

Rao, M. E.; Georgeff M. 1991. *Modeling rational agents within a BDI* architecture. KR 1991. 473-484

Rao A. S.; Georgeff M. P. 1995. *BDI agents: From theory to practice*. In Proceedings of the First Int. Conference on Multi-Agent Systems. Menlo Park, Calif.: AAAI Press.

Russel S.; Norwig P. 1995. *Artificial Intelligence. A modern approach.* Prentice All.

Simon, H. A. 1982. *Models of Bounded Rationality: Behavioral Economics and Business Organization. The* MIT Press.

Simon, H. A. 1984. *The sciences of the artificial*. *The* MIT press.

Weber, M. 1978. *Economy and Society: An interpretive Outline of Economics*. In eds. Roth C. and Wittich C. University of California Press.

Weick, E. K. 1979. *The social psychology of organizing*. McGraw-Hill Inc.

Weick, E. K. 1995. *Sensemaking in Organizations* SAGE.

Zack, M. H. 1999. *Managing Organizational Ignorance*. Knowledge Directions 1:36-49.