

Adaptivity in Agent-based Systems via Interplay between Action Selection and Norm Selection

Henry Hexmoor

Computer Science & Computer Engineering Department, Engineering Hall, Room 313,
Fayetteville, AR 72701
hexmoor@uark.edu

1 Introduction

Beyond the ability for adaptation to an environment, Self-Adaptive Software (SAS) embodies the capacity and the initiative to adapt its own behavior. The adaptation is due to a need to maintain a desired or a reference relationship between a set of input and output signals. We may loosely divide SAS into adaptor and adapted components. The adapted component has an ongoing relationship with the environment. The adaptor detects and evaluates the need for change in the operation of the adapted component. In anthropomorphic terms, this detection and evaluation involves the cognitive processes of introspection and assimilation. However, an artifact may only need a supervisory control module. In a hierarchical system, the idea of adapted and adaptor can be extended to several levels with the higher level adapting the function of lower level. Hierarchical architectural systems are well studied. However, self-adaptive software is more than hierarchical control or application of adaptive techniques such as neural nets or genetic approaches. SAS adaptation is at the mission level. To the extent it is adaptive, a SAS system perceives the environment, computes trends, and compares the system performance to the operationally defined reference metrics and not to concrete quantities. For instance, in a game of football, the coach monitors the game and each player very closely and decides on strategies and swapping players. The coach's decision is at the operational level. The coach as the adaptor component of the team understands nuances of the game and the overall state of the game at a global, functional level. But in its most general form, to understand the environment and to adjust to it is the wholly grail of AI.

We focus on agent-based systems, which represent an important class of potentially self-adaptive systems. An important paradigm in agent-based systems is to consider intentional notions of Belief, Desire, and Intention (BDI agents). BDI agents possess update/revision functions for each intentional component. Beliefs are adapted to the agent's current state of mind and changes in the environment. Desires or specific goals are adapted to the agent's beliefs and attuned with the changing environment. Intentions over specific actions are adapted agent' beliefs and desires. These adaptations create adaptivity for the agent at a behavior level called action selection. Action selection is a cognitive

process of finding the most appropriate and the most relevant action to perform [12, 18, 19]. Adaptivity at the action selection is limited to individual rationality. Agents that operate among other agents must reason about other agents and possess social rationality. Coordination with other agents provides yet another form of adaptivity. Coordination strategies are varied among domains. In this paper we will investigate selection of social norms as coordination strategy among agents.

Actions selected based on norms may be in conflict with action selected based entirely on individual concerns, i.e, based on action selection. To avoid such a conflict, we suggest treating *action selection* as the adapted component and the *norm selection* as the adaptor, Figure 1. Here norm selection will suggest modes of action selection as normative patterns.




Figure 1. Norm selection guides action selection

Interaction between layers is continuous. The following is a simplified outline of introducing norms to the agent's top-level decision-making loop. As the descriptions in the steps will illustrate, the interaction between layers is nontrivial. In general norm selection will be over a larger temporal horizon than action selection and in that respect, the process of action selection will cohere to adopted norms. Of course actions will need to be coherent but that is beyond this discussion. For a discussion of coherence see [21].

1. Perceive the environment and update beliefs in the world model. Both action selection and norm selection share the results of this process.

2. Consider options for individual goals, and norms. Adopt the most compatible goals and norms. I.e., processes of norm selection and action selection are concurrent and interacting but norm selection will be the dominant process.
3. Consider plans to implement adopted goals, and norms that lead to options for action selection. Both action selection and goal selection will inform planning. Subsequently, planning will determine actions to be performed.
4. Formulate intentions to act and perform planned actions.

The remainder of this paper begins with a discussion of norms and norm selection. This is followed by sections that take on utilitarian and intentional views of dealing with norms. We then feature teaming as a special class of norms. The paper ends with concluding remarks about self adaptivity in agent based systems.

2 Norm selection guides action selection

Selection of a normative pattern of interaction is often articulated by a person's discourse that cites predominant situated mental attitudes for choosing the normative pattern. For example, if you ask a person proctoring an exam "why are you monitoring your students so closely" the answer might be "I don't trust them they have cheated before". In this example, the person is engaged in a normative pattern of keeping a watchful eye on students and expresses that *trust* was the main reason for adopting this specific normative pattern. By exploring the connection between mental attitudes, normative patterns of interaction, and action selection we wish to develop agents that can self-adapt. We aim for a calculus over mental attitudes that will modify the agent's in-the-field decision making and lower the need for human supervision. Our model borrows from social human behavior. Human's ever-changing mental attitudes regulate selection of normative patterns needed for survival and success. The contribution to adaptive software is the understanding and algorithmic duplication of these mental attitudes. This paper is a preliminary report on developing computational framework that supports adaptation in software in similar ways that humans perform this task.

For a rational agent, action selection is guided by its situated attitude governing its relationship to its environment and other agents. But possessing social attitudes might be unnecessary for agents. In fact it might be argued that attitudes are absent in biological systems. For instance, it can be argued that a combination of behavioral pattern recognition, imitation, and instinctual attunement is sufficient to produce complex relationships such as teaming in Vervet monkeys.¹ We suggest that if an agent is allowed to have nontrivial autonomy about whether and how it will choose the complex relationship (such as teaming), it must be able to reason about mental ingredients that lead

¹ Chris Hill at U of Arkansas suggested this example.

it in its choice and use of those relationships. The monkeys do not have our level of requisite autonomy. Perhaps, they engage in teaming by default and not by choice.

Norms are “often viewed as emergent properties of utilitarian agents’ behavior” and “no theory of the acquisition of normative attitudes as grounded upon agents’ representations has yet been provided” [11]. Norms can be useful when two or more interacting agents are aware of each other’s norms, and use it in their individual deliberation. The general purpose of norms is in anticipation of expected mental and physical actions of other agents. There are many types of norms. Some norms govern intra-personal interactions such as “reactivity”². Some norms govern interpersonal interactions such as “communal laws”. Social norms can in general influence an agent’s behaviors and plans as well as the agent’s mental states with respect to goals, beliefs, and intentions. Some norms can either be a collection of disjoint rules of thumb as in “look before crossing” and are domain specific. Other norms might be a cluster of related norms and are adopted in entirety. For example, “friendship”, and “teaming” are well-developed clusters of social norms useful in an agent society [7]. A recent domain independent technique for collecting communication norms is a *conversation policy*, which is a model of the structure and content of a type of agent conversation for a certain communicative function [13]. Norms can be specified formally and systematically or informally and heuristically. Norms are not mutually exclusive and an agent may enact one or more of them concurrently. An external observer can often detect many of the active norms. In social situations most people are aware of effectiveness of their action selection under the influence of a norm. For instance, people know when a negotiation scheme (as a cluster of norms) has become fruitless, their role choice is ineffective, and their team membership is a hindrance. This can be used to assess related goal status and perhaps utilities.

Here are a few clusters of norms. In this paper we will focus only on teamwork.

- Teamwork is a class of interpersonal social conventions that bond a number of individuals to a goal. The bond might be loose or rigid, [17].
- Situated action and reaction is a class of intra-personal methods for coping in the world that connects sensing to acting with minimum of mental processing. In their repertoire, many agent architectures provide a tier for this kind of interaction [10].
- Control and monitoring is a class of systematic methods in which the agent monitors its environment for predetermined outcomes. The mechanisms can be interpersonal or intra-personal. For instance, an agent may select a master-slave role for agents in its surrounding. [8].
- Negotiation is an interpersonal template for interaction that an agent may use in order to bring about desired changes [22].
- Social role adoption and adaptivity is an interpersonal skill for choosing a role that best fits the agent’s capabilities and objectives [23].

² “reactivity” is a pattern of personal habits/skills/routine [12, 19].

It is very interesting but at present omitted from our consideration is: (a) emergence of norms in agents or in the society of agents, (b) changes in norms, (c) how an agent acquires a norm as a pattern of activity, or (d) how an agent violates or modifies a norm. Instead we are interested in the following: (a) how an agent instantiates (i.e., enacts) a norm, and (b) what mental ingredients of an agent are responsible for characterizing and sustaining a norm.

2.1 Identifying a norm

We suggest the following minimal list of norm characteristics.

- Involve two or more agents, each of whom shares the norm.
- Agents have power to not choose them.
- There need not be any direct rational account of the norms available to the agents.
- There are explicit or implicit sanctions and rewards for norm adoption.

While strategies might fit the above requirements, norms dwell in pragmatics of interaction whereas strategies are means-end techniques. Norms and values (i.e., principles) are both influences on behavior. The agent must be able to respond to changing conditions and this requires that the agent be able to drop current norms, obligations and intentions in favor of norms, obligations and intentions that better fit the changed environment. Norms therefore cannot be unconditionally adopted. Norms are not defaults or exceptions. Defaults and exceptions can be used for similar purposes as norms. e.g., in coordination. However, unlike defaults and exception, agents deliberate about norms. Obligations are strong norms and are derived from it. We consider interdictions, bans, and other similar terms as kinds of norms. Norms might produce goals for an agent but the norm is not the goal itself but the surrounding social context. Some researchers have treated norms as constraints. Constraints are a knowledge encoding method regulating an agent's actions. Norms that might resemble constraints preserve the teleological foundations of such encoding so the agent can perform reasoning. Norms are a subset of social laws and the agent's pragmatics of using them.

To recapitulate, we divide norms into two classes. The first class is a set of conventions or behaviors for a given situation. The second class of norms is valuations of behavior for deliberation and prediction. The second class of norms varies from societal or institutional to personal or individual. These norms also vary from absolute or global to relative and local.

2.2 Norm Selection

The problem of off-line norm selection has been studied in [14, 20]. Simplicity and minimality as criteria used for norm selection where alternative norms are available. In contrast to this approach, we focus on the process of norm formation and its relation to mental attitudes. This will allow us to account for both the choice and dynamics of social norms.

Here, like [1], and in contrast to a collective decision, we take an individualistic view of the choice of norm that will affect the agent itself and others. A possible approach to norm selection by an agent can be in terms of rule-utilitarianism where an agent chooses a norm that maximizes social welfare of agents it cares about. This would mean a kind of machine learning that would use feedback to reinforce choices of norm. However attractive, such feedback is hard to compute as well as unintuitive.

Another basis for norm selection can be based on a system of mental attitudes with varying intensity. These models will prescribe the proper combination of mental attitudes as reasons to adopt a norm. An agent equipped with such a system and endowed with abilities to form mental attitudes over such things as power, autonomy, sociability, etc., will then instantiate the norms that comply with its accompanying system of decision.

In the next section we will discuss the relation between norms and utilities.

3 Utilitarian view of Norms

We limit our consideration of norms relative to goal achievement. Furthermore, we assume agents are all self-interested. We assume the agent has a goal that it can accomplish. However, when it chooses to consider interaction with another agent, it will consider applicable norms, which in turn affects its action selection.

Let's assume an agent's norms with respect to a goal can be grouped into disjoint sets. Assume norms in a set are consistent and the agent must pick one set from among its norm sets. Examples of norm sets are collaboration or competition.

It is conceivable to reason about norms explicitly. However, we believe that reasoning about norms is often indirect and through mental attitudes. The indirect nature of reasoning about norms is explored in section 3. By using utilities, we can model arbitration among norms explicitly. I.e., the agent can weigh one norm against another. This is the focus of this section.. A common objection to this approach is unavailability of utilities. Other objections are along the lines of inadequacies of decision theory and its unrealism in human decision-making. However, this kind of explicit reasoning with utilities can be used in problem solving and planning. Let's define a norm.

Definition: Norm = (O, R, G, U).

- O is the content of the norm set which is being considered for norms selection. This is independent of G, which is being considered for action selection.

- R is a social marginal utility function that considers the social gains and losses with respect to G.
- G is the agent's goal that invokes the norm set. Norms are indexed by goals. This is the connection between the agent's individual and social rationality.
- U is an individual marginal utility function that considers the agent's gains and losses with respect to G. The gains and losses might be rewards, penalties, and other valuable s such as reputation.

Rules 1 (below) is suggested for selecting a norm at the norm selection layer when a goal is considered at the action selection layer.

Rule 1: Norm selection for one goal

An agent considering a goal G and norm sets $N_1 \dots N_m$ will pick the norm set N_i , which maximizes its goal achievement utility. If all norms for G will produce a negative utility, no norm is selected and the goal is reassessed.

Rules 2 (below) is suggested for selecting a goal at the action selection layer when a norm is considered at the norms selection layer.

Rule 2: Goal selection for one norm

An agent considering a norm N with content as goal G and other goals $G_1 \dots G_m$ will pick the goal G iff the utility of the norm (and its corresponding content G) is higher than utilities of all $G_1 \dots G_m$.

Rule 3 (below) is suggested for selecting a norm at the norm selection layer when multiple goals are considered at the action selection layer.

Rule 3: Norm selection for multiple goals

An agent considering goals G_1, G_2, \dots, G_m and corresponding norm sets $N_{11} \dots N_{1m1}, N_{11} \dots N_{1m2}, \dots,$ and $N_{11} \dots N_{1mm}$ will pick one norm set for each goal (Norms N_1, \dots, N_m) with the constraint that $U_1 + \dots + U_m$ is maximized.

At times, utilities due to norm selection can have such an effect as to influence the goal selection. The agent may not be able to select a norm since the utility losses outweigh any gains it might have. This agent must be prepared to pay the sanction penalties for non-adoption of norm corresponding to the goal. If the sanction penalty is larger than the utility of the goal, the goal should be dropped. An example is robbery. The thief may think about the penalty of the sanction and when it is considers it to be higher than any gains, it should drop it. What happens if the agent has several goals and is faced with some sanctions that each outweighs utility of the corresponding goal? Naturally, the utility maximizing agent will not drop the goals and will be prepared to pay the sanctions.

4 BDI and deliberative view of norms

We now turn to defining norms indirectly in terms of mental attitudes. Since agents deliberate about their mental attitudes, norms based on mental attitudes are deliberative norms.

Mental attitudes arise out of the agent's deliberation on motivations, which we will take for granted in this paper. To use a production system metaphor, these are the agent's situated attitudes that the agent will use for conflict resolution among competing courses of action. We suggest that if a sufficiently sophisticated agent is questioned about its choice of a specific norm or a norm class, it should mirror typical human response and present the mental attitudes that led to its choice of template. These mental attitudes are sometimes colloquially referred to with phrases such as "in the spirit of ...", which imply a "mood" or a predominant feeling. This might suggest the notion of personality. However, unlike personality, which dictates an anthropomorphic emotional state of the agent such as fear and tenacity, for agents we will focus on self-evaluative qualities, which are derived from an agent's experience. We are focusing on the space of relationships among mental attitudes and how that leads to its choice of action selection template, Figure 2. Much of this diagram is a simplified relationship between the Beliefs, Desires, and Intentions of an agent from the popular BDI paradigm [24]. The most important parts for us are the *adopted norms* and *adopted intentions* along with accompanying motivational attitudes. Motivations are instantiations of an agent's high-level objectives and reasoning which produce such an instantiation. Motivations can be thought of as the agent's *desires* and might be amorphous and abstract. Goals are concrete and more readily achievable. "brf" is belief revision function and *percept and communications* are used as inputs to revise old beliefs. "Option generation" is a function that maps beliefs and old desires and old intentions to updated desires. In addition to functionality described in [24] we have this function take into account agent's deontological commitments as well. This is the main function for the action selection layer in Figure 1. "Selection function" maps beliefs, desires, prior deontological commitments, as well as agent's adopted roles to adopted norms. This is the main function performed in the norm selection layer in Figure 1. *Adopted roles* are a result of negotiating plans with other agents. *Negotiation* might be negligible or substantial and there are normative considerations as well. Fuller treatment of plans and roles is beyond our scope in this paper.

In social situations most people are adept at being conscious of their own mental attitudes and detect mental attitudes of others. People might even command or suggest adoption of these mental attitudes from other people as in "follow the rules" or "don't be impulsive". This is shown in Figure 2 as the environment's direct effect on beliefs. Unlike bi-valued logics, mental attitudes can be more accurately understood as having quantity as in a mass noun. For instance, 'sociability' is like 'water' in its measurability. There are metrics for some nouns like 'water' whereas others like 'cloud'; 'ink blot', and our attitudes lack such well-developed metrics. Common parlance such as "fairly

sociable” or “highly impulsive” motivate us to consider attitudes as having possibility values, as in fuzzy logic.




Figure 2. Relationship between norm selection and intention selection

The following is a partial list of mental attitudes of general interest to us: cooperation, autonomy, exploration, impulsiveness, commitment, responsibility, trust, conformance, sociability, dependence, and power. Defining these terms is difficult and can take much space that we don't have here. In this paper we will rely on the reader's intuitive understanding of these terms. Next we redefine norms in term of mental attitudes.

Definition: Norm = (O, R, G, A)

- O is the content of the norm set which is being considered for norms selection. This is independent of G, which is being considered for action selection.
- R is the marginal deontological effects of adoption/rejection of the norm.
- G is the agent's goal that invokes the norm set. Norms are indexed by goals. This is the connection between the agent's individual and social rationality.
- A is a set of mental attitudes along with degrees for each. The mental attitudes are characterized by using the notions of belief, desire, and intention. The degrees of mental attitudes form a required pattern against which we will match the agent's actual mental attitudes.

A is a set of mental attitudes the agent somehow chooses. The mechanisms of selecting a mental attitude are very much an area of research in Ecommerce and multiagency for which we will not provide an answer in this paper. Utilitarian approaches have been

applied in such mechanisms for mental attitude determination [4, 5, 6]. There are many discussions of mental attitudes found in the literature, e.g., [7, 8, 9].

An agent at any given time will have a pattern of mental pattern of mental attitudes with respect to an adopted goal. For instance, the agent will experience a degree of autonomy, a level of trust, and so on. Norm selection turns into matching the existing pattern of mental attitudes to required patterns of mental attitudes for each norm.

Rule 4 (below) revisits Rule 1 and is suggested for selecting a norm at the norm selection layer when a single goal is considered at the action selection layer.

Rule 4: Norm selection for one goal

An agent considering goal G , norm sets $N_1 \dots N_m$ each with required mental attitudes and their corresponding degrees and actual mental attitudes $A_1 \dots A_N$ will pick the norm set N_i with the closest match between actual mental attitudes and the required mental attitudes.

Rule 5 (below) revisits rule 2 and is suggested for selecting a goal at the action selection layer when a norm is considered at the norm selection layer.

Rule 5: Goal selection for one norm

An agent considering a norm N with content as goal G and other goals $G_1 \dots G_m$ and actual mental attitudes $A_1 \dots A_N$ will pick the goal G iff the actual mental attitudes match the required attitude of N more than required attitudes of any of the Norms corresponding to goals $G_1 \dots G_m$.

Rule 6 (below) revisits rule 3 and is suggested for selecting a norm at the norm selection layer when multiple goals are considered at the action selection layer.

Rule 6: Norm selection for multiple goals

An agent considering goals G_1, G_2, \dots, G_m , corresponding norm sets $N_{11} \dots N_{1m1}, N_{11} \dots N_{1m2}, \dots$, and $N_{11} \dots N_{1mm}$ and actual mental attitudes $A_1 \dots A_N$, will pick one norm set for each goal (Norms N_1, \dots, N_m) but with the constraint that maximizes the match between degrees of actual mental attitudes and degrees of required attitudes in norms.

With utility analysis of norms, norm selection was shown to influence the goal selection and at times lead the agent to abandon its goals. This is also possible in the attitude formulation of norms. The content of the norm can be the negation of the norm's goal. This is the case where the norm for a goal prohibits that goal. If the actual mental attitudes activate the norm strongly, then the goal and its norm counteract. The agent might abandon such a goal. If the goal is part of a series of related goals, as long as the number of counteractions is fewer than the number of goal, the agent might be inclined to execute the goals. We have been very sketchy here and there is a lot of room for formalization, which we leave for future work.

5 Teaming

Teaming is a cluster of norms for collaboration. We have developed a utilitarian approach for deciding to team [15, 16] and an attitude-based approach [2, 17]. Our model of selecting a team as a cluster of norms includes consideration of four conditions: (1) joint intentions, (2) group autonomy, (3) joint awareness, and (4) joint cooperation. We claim these conditions that an agent needs to consider to form a team. Once a team is formed there will be conditions beyond the ones we will review that the agent will use to remain in a team or exit it.

The central attitudes in considering to form a team are the notions of group autonomy and cooperation. Agents jointly intend a goal if and only if as a group they are jointly autonomous toward that goal plus they choose to attempt actions toward it. Group autonomy in a group of agents is when each member of the team believes in adequacy of joint abilities of the group and perceives adequate mutual (among group members) and external (from agents outside the group) permission to exercise their ability. Joint awareness in a group is when each agent in the group believes that together with the group it has joint ability toward the desired state and believes that it has joint intention to achieve that state. Joint cooperative condition is met when all member of the team believe that they will share in the benefits with the group and they have joint responsibility to achieve the group objective. Responsibility is in turn the condition that if an agent can do an action to facilitate the desired group goal will intend to do it.

Perhaps other attitudes such as trust toward potential and actual team members are relevant. It is shown that agents that learn to trust one another are more likely to cooperate [3]. It is intuitive that with more trust among agents, there will be more motivation for teamwork.

6 Conclusion

Norms affect action selection in social agents. We showed that reasoning about norms guide norm selection as well as affect choice of goals. Norms might be arbitrated by utilities. Alternatively, an agent may deliberate about mental attitudes to determine norms. We have sketched some relationships between norms and goals.

In many artificial systems, designers hard-code norms on agent systems. We believe by allowing an agent to select its own norm that fits its attitudes, we increase its adaptivity. Agent system engineering needs to be extended to account for norm classes so that agents can adapt to social contexts. Much remains to be done in extending the agent communication language and agent inference mechanisms to encode notions of mental attitudes such as autonomy, trust, power, etc.

Acknowledgements

This work is supported by AFOSR grant F49620-00-1-0302.

References

1. E. Alonso, 1998. How Individuals Negotiate Societies, Proc. of the Third International Conference on Multi-Agent Systems (ICMAS-98), pp. 18-25. July 3-8, Paris. IEEE Computer Society Press.
2. G. Beavers and H. Hexmoor, Teams of Agents, In the Proceedings of IEEE SMC, Arizona.
3. A. Birk, 2000. Boosting Cooperation by Evolving Trust, **Applied Artificial Intelligence Journal**, Taylor and Francis.
4. S. Brainov, 1996. Altruistic Cooperation between Self-Interested Agents. In Proceedings of **European Conference on Artificial Intelligence**, 519-523, John Wiley & Sons.
5. S. Brainov, 2000. The role and the Impact of Preferences on Multiagent Interaction. In **Intelligent Agents VI**, Lesperance and Jennings (eds.), 349-363, Springer-Verlag.
6. S. Brainov and T. Sandholm, 1999. Power, Dependence and Stability in Multiagent Plans. In Proceedings of **AAAI'99**, pages 11-16, Orlando, Florida.
7. C. Castelfranchi, 1990. Social Power. In Demazeau Y. and Muller J.-P. eds. *Decentralized AI - Proceedings of the First European Workshop on Modeling Autonomous Agents in a Multi-Agent World*. 49-62. Elsevier Science Publishers.
8. C. Castelfranchi., and R. Falcone, 2000. Trust and Control: A Dialectic Link
9. C. Castelfranchi, F. Dignum, C. Konker, and J. Treur, 1999. Deliberative Normative Agents: Principles and Architectures, pp. 364-378, **ATAL-99**.
10. W. Clancey, 1997. **Situated Cognition: On Human Knowledge and Computer Representations**, Cambridge University Press.
11. R. Conte, R. Falcone, and G. Sartor, 1999. Introduction: Agents and Norms: How to fill the gap? **AI&Law Special Issue on Agents and Norms**. R. Conte, R. Falcone, and G. Sartor (Eds), vol 7, no 1, p 1-15.
12. R. Cooper, T. Shallice, and J. Farrington, 1995. Symbolic and continuous processes in the automatic selection of actions. In Hallam, J., (ed.) **Hybrid Problems, Hybrid Solutions**. IOS Press. Amsterdam. pp. 27-37.
13. M. Greaves, H. Holmback, and J. Bradshaw, (1999). What is Conversation Policy? In *Proceedings of the Autonomous Agents '99, Workshop on Specifying and Implementing Conversation Policies*, Seattle, WA, May 1999.
14. D. Fitoussi and M. Tennenholtz, 2000. Choosing social laws for multiagent systems: minimality and simplicity. *Artificial Intelligence*, 119(2000) 61-101.
15. H. Hexmoor and H. Duchscherer, 2000. Shared Autonomy and Teaming: A preliminary report. In Proceedings of **Workshop on Performance Metrics for Intelligent Systems**, NIST, Washington, DC.
16. H. Hexmoor and H. Duchscherer, 2001. Efficiency as a Motivation to Team. In Proceedings of **FLAIRS 2001**, FL.

17. H. Hexmoor and G. Beavers, 2001. Towards Teams of Agents, In Proceedings of the International Conference in Artificial Intelligence, H. R. Arabnia, (ed), (**IC-AI'2001**), Las Vegas, CSREA Press.
18. P. Maes. 1989. How to Do the Right Thing. **Connection Science Journal**, Vol. 1, No. 3, pp. 291-323.
- 19 T. Norman, and T. Shallice, 1986. Attention to action: Willed and automatic control of behavior. In Davidson, R., Schwartz, G., and Shapiro, D., (eds.) *Consciousness and Self Regulation: **Advances in Research and Theory***, Volume 4. Plenum, New York, NY. pp. 1-18.
- 20 Y. Shoham and M. Tennenholtz, 1995. Social laws for artificial agent societies: Off-line design. *Artificial Intelligence* 73(1995) 231-252.
- 21 P. Thagard, 2000. **Coherence in Thought and Action** (Life and Mind: Philosophical Issues in Biology and Psychology), MIT Press.
- 22 D. Walton, 1989. **Informal Logic: A Handbook for Critical Argumentation**, Cambridge, Cambridge University Press.
- 23 R. Werner, 1989. Cooperative Agents: A unified Theory of Communication and Social Structures. In L. Gasser, M. Huhns, editors, **Distributed Artificial Intelligence**, Volume II.
- 24 M. Wooldridge, 2000. **Reasoning about Rational Agents**. The MIT Press.