Social Laws for Agent Modeling

Amol Dattatraya Mali
Dept. of Mathematical and Computer Sciences
Colorado school of mines
Golden, CO 80401-1887
amali@mines.edu

Abstract

Multiagent domains emphasize that agents should be able to predict actions of other agents. A popular mechanism to achieve this is that of providing agents with models of other agents. Social laws have been proposed as a model of multiagent interaction with claims that they can get rid of perception, reduce communication cost and planning time. However the utility of social laws for agent modeling has not been explored. In this paper types of social laws are characterized and criteria for evaluating and comparing them have been defined. These criteria also serve as measures of the utility of social laws in modeling agents. It is shown here that to be easier to design and test, social laws should exploit other representations like potential fields to model interactions between agents. Social laws and belief, desire and intention based (BDI) agent modeling paradigm are compared. Use of laws to encode meta-knowledge, characteristics of laws and situations under which they can be used are discussed. Some unaddressed issues have been raised at the end.

1 Introduction

Multiple agents are proving to be useful for parallelization and decomposition of activities. In particular, they emphasize distributed perception and actuation. For successful completion of tasks, avoiding interference with the actions of other agents, it is necessary for agents to predict actions of other agents. [Tambe *et al* 95] point out that agent modeling (especially opponent modeling) is a key requirement for building automated pilots. [Hill & Johnson 94] mention that agent modeling capabilities are valuable for intelligent tutoring, natural language processing, expert consultation and tactical decision making. Traffic control, coordination of vehicular traffic or air traffic control illustrate situations where software agents make decisions based on communication and agreement with other agents.

Researchers like [Durfee 92] have claimed that social skills learnt in childhood like sharing information, playing fair, avoiding collisions with people, putting things back at right places and cleaning up our own mess provide necessary and sufficient skills for the components of a distributed system. The notion of social laws is similar to the notion of these social skills. [Shoham & Tennenholtz 95] define a social law as a set of constraints on actions available to an agent. [Nagao & Takeuchi 94] realize a social agent that hears human-to-human conversation and identifies the causes of misunderstanding. They point out that social rules like "avoid misunderstanding", "do not speak while other people are speaking", "silence is not good" and "contribute whenever possible" are highly useful in a conversation. Desirable behavior is always easier to specify than an optimal behavior and social laws provide a mechanism to express such preferences. Laws allow *agents to develop models of potential interactions* among their actions and those of other agents.

[Shoham & Tennenholtz 92] present a general model of a social law in a computational system and investigate some of its properties. This work does not address the issue of how to evaluate and compare social laws and use them to model agents.

Social laws have been proposed as a model of multi-agent interaction with claims that they can get rid of perception, reduce communication cost and planning time. For example, if there is an $N \times N$ grid with N robots in first row, with only one robot in a grid square and if each robot follows convention of staying in its column, collisions are automatically avoided. Such a team of robots can be useful for transporting material from first row to last row. In this case, the robots do not need to perceive each other and no communication is required among them. However the research lacks a formal approach to evaluate and compare the laws. In this paper types of social laws are characterized and criteria for evaluating and comparing them have been defined. These also serve as a measure of the utility of social laws in agent modeling. I argue that social laws for modeling agents can be designed without trial and error if representations like potential fields are used to express them. The paper points out the advantages of social laws over belief, desire and intention based (BDI) paradigm. The limitations of social laws in modeling agents have been discussed. It is proposed that an integration of social laws and the BDI paradigm will be a more efficient approach to modeling agents. Some issues not addressed by previous research have been raised.

This paper is organized as follows. *Section 2* characterizes the types of social laws and defines some criteria to evaluate and compare them. *Section 3* raises the issue of representations for social laws to effectively model agents. In *section 4* social laws and BDI based paradigm are compared. *Section 5* discusses role of social laws in encoding meta-knowledge, their characteristics and situations under which they can be used. *Section 6* presents conclusions.

2 Formalization

In this section I characterize the types of social laws and define criteria to evaluate and compare them.

2.1 Types of social laws

In a general model of DAI systems, multiple agents (distributed spatially, logically (softbots on the same machine) or temporally) fulfill coordinated tasks. Taking this as a basis, I identify two types of social laws here and discuss their utility in agent modeling.

K denotes the total knowledge required for a task T, L denotes the set of social laws designed for the task and A denotes the set of agents in the colony.

- **Laws based on knowledge decomposition**

Here a task is decomposed into sub-tasks such that each sub-task requires distinct knowledge and can be allocated to a set of agents. By distinct knowledge of different sets of agents, I mean different goal fulfilling plans or $< condition, action >$ rules. Fulfillment of a task is defined in terms of attainment of desired world state. All agents dedicated to a subtask have the same subset of the knowledge. Social laws can be designed for agents allocated to a particular subtask. Consider the scenario where a set of robots have laws for carrying out the task of moving unpainted cans to a painting station (e.g. "do not try to grasp a can that other robot is trying to"), other set has laws for painting the cans (e.g. "do not grab an unpainted can that is being painted") and third set has laws for drying the painted cans and putting them in a store (e.g. "do not pick up a can dropped at the store" which avoids infinite cycles). In a colony with laws based on decomposition of knowledge, we have

$$\exists A, T, K, L[Needs(T, K) \bigwedge K = \bigcup k_j \bigwedge A = \bigcup a_i \bigwedge T = \bigcup t_j \bigwedge L = \bigcup l_j \bigwedge Allocated(a_i, t_j) \Longrightarrow Needs(t_j, k_j) \bigwedge Needs(t_j, l_j) \bigwedge Has(a_i, l_j)]$$

- **Laws based on spatial decomposition**

Here same task is carried out in different regions of a workspace in parallel and the task is fulfilled when individual subtasks in different regions of the workspace are fulfilled. All agents are controlled by the same set of laws. Examples of this type of laws are those for searching an area (e.g. "maintain a distance d from other robots" so that no two robots search the same area), extinguishing fire etc. Let S denote the workspace. Hence we have

$$\exists A, T, S, L[Needs(T, S) \bigwedge S = \bigcup s_j \bigwedge A = \bigcup a_i \bigwedge T = \bigcup t_j \bigwedge Allocated(a_i, t_j) \Longrightarrow Needs(t_j, s_j) \bigwedge Needs(t_j, L) \bigwedge Has(a_i, L)]$$

If each agent models other agents in terms of the laws that control them, it becomes possible for it to evaluate the impact of its actions on other agents and vice versa e.g. if the law of maintaining a fixed distance from other robots is used then every robot knows that its movement would cause its neighbours to move by a certain distance which would be propagated through the colony.

2.2 Criteria for evaluating laws

Here a number of criteria for evaluating and comparing social laws have been defined. Here it is assumed that an agent is modeled by the laws it obeys. Hence by an agent A_1 having a model of agent A_2 I mean A_1 knowing the laws that A_2 obeys. Since successful agent modeling is a prerequisite for fulfillment of tasks in multiagent domains and it is assumed here that laws are used to model agents, these criteria also serve as measures of success of laws in modeling agents.

- **Flexibility**

Let s_i denote the number of situations under which a law l_i works. (A situation is a well formed formula describing the state of an environment.) Then a law l_1 is more flexible than a law l_2 if and only if $s_1 > s_2$. A law incorporating more constraints is likely to be less flexible. This notion of flexibility is different from that in [Briggs & Cook 95] because they add more laws for more flexibility and I define flexibility for a single law.

- **Usefulness**

In real world, the frequency of occurrences of the situations under which a law works is more important, e.g. if a law l_1 works under 10 situations and a law l_2 works under only one situation (under which l_1 fails) that occurs all the time (infinite frequency) then l_2 is more useful than l_1. This is captured by *usefulness* of a law defined as the sum of the frequencies of situations under which it works,

$$\sum_{i=1}^{n} f_i$$

where f_i are the frequencies and n is the number of situations under which the law works.

- **Efficiency**

Let $t(l_i, T_i)$ denote the time required to fulfill a task T_i using a law l_i. Then a law l_i is at least as efficient as a law l_j if
$\forall T_k \in T, t(l_i, T_k) \leq t(l_j, T_k)$ where T is the space of tasks that can be fulfilled by using both l_i and l_j independently. We define a space of laws as a set of laws. We denote by $\tau_G(L)$ the greatest task space fulfillable using a space of laws L. Then a space of laws L_i is more efficient than a space of laws L_j if $\frac{\tau_G(L_i)}{|L_i|} > \frac{\tau_G(L_j)}{|L_j|}$. It should be noted that simple tasks like goal seeking and obstacle avoidance can be fulfilled by using a few laws without the need to use plans. For complicated tasks, laws will be used with plans, in that case the above definition assumes that different spaces of laws being compared are used with the same set of plans or the same set of actions available to agents in a colony.

- **Uniformity**

Using a law, agents will fulfill a task taking various times e.g. let us say that all agents in a colony have to reach a goal avoiding each other, using a law they can reach the goal but will take different times. Let t_i be the time taken by agent a_i to reach the goal using law l_k. Let t = $\frac{\Sigma t_i}{|A|}$ be the average time required to reach the goal in the colony A. Then I define the mean squared error $\frac{\Sigma(t-t_i)^2}{|A|}$ as a measure of uniformity of the law l_k. A more uniform law will have a lower mean squared error.

- **Speedup**

This is a measure of reduction in the time required to fulfill a task using a social law over some other agent modeling paradigm. Let $t(L, T_i)$ be the time required to fulfill a task T_i using a space of laws L alone or a combination of L and some other agent modeling paradigm p. and $t(p, T_i)$ be the time required to fulfill the same task using the other agent modeling paradigm p alone (e.g. BDI paradigm discussed in section 4). Then $\frac{t(p,T_i)}{t(L,T_i)}$ is the speedup.

3 Representations for social laws

[Shoham & Tennenholtz 95] propose some traffic laws to ensure that the robots in a colony reach their destinations and avoid collisions. However these laws are required to specify a number of constraints like directions of motion of robots in odd and even rows and handle boundary conditions like reaching rightmost or leftmost column. The sophisticated traffic law that they propose involves more constraints. These constraints are derived using a trial and error procedure.

The potential field based representation that I use here does not require the trial and error procedure.

Theorem 1. In absence of local minima, potential fields can be used to define social laws to model robots and seek goals in a colony of robots in polynomial time.

Proof - This task needs that robots should avoid each other and reach their goal. This naturally leads to two social laws - "avoid other robots" and "progress towards goal". Choosing potential fields to encode these laws, we can define potential functions for goals that exert attractive force on the robots and each robot itself will have a potential function to exert repulsive force on its neighbors. At each step each robot sums the forces exerted by its goal and other robots it senses and computes the next direction and magnitude of its motion as per $x_{t+1} = f(x_t, \Sigma\Phi_{repulsive} + \Phi_{goal})$, $y_{t+1} = f(y_t, \Sigma\Phi_{repulsive} + \Phi_{goal})$.

Let us say that there are N robots in the colony and the average number of steps required for a robot to reach its goal is S. At each step let us assume that each robot senses on an average, pN other robots where $0 \leq p \leq 1 - \frac{1}{N}$. Then each robot will have to compute pN repulsive forces and an attractive force. Hence the number of steps required for goal seeking using potential fields in a robot colony is $(N\ S\ (p\ N + 1))$, polynomial time. In worst case $p = 1 - \frac{1}{N}$, still the complexity is polynomial. Let us say that at each step, each robot is concerned about collisions with other robots. This requires them to predict the magnitude and direction of movements of other robots. Then each robot can imagine itself to be present in other robots' current locations and determine the distance to their goals and distances to robots sensed by them from that location. Hence at time t, each robot can determine the path traced by every other robot during the time interval $[t, t+1]$ using the laws. Hence the result.□

The two laws could be superimposed to achieve the desired result because of the representation chosen.

4 BDI Paradigm

[Georgeff & Lansky 87] develop a procedural reasoning system based on the attitudes of belief, desire and intentions. The beliefs are facts about the world and desires are the goals to be realized. An intention is a commitment to an act whilst in a certain mental state. At any given instant, the actions being considered by PRS depend not only on its current desires or goals but also on its beliefs and previously formed intentions. The system can modify its beliefs, desires and intentions. [Jennings 95] develops a model of cooperation based on joint intentions that specifies pre-conditions to be attained before collaboration can commence and prescribes how to behave both when joint activity is progressing satisfactorily and when it has trouble. Joint intentions are a joint commitment to perform a collective action while in a certain shared mental state. In the BDI paradigm, providing an agent A with model of agent B means letting A know the beliefs and desires of B. Then A can use beliefs and desires of B and knowledge of the current situation to infer B's intentions and predict its actions. In the work [Nagao 93], a plan recognition module determines a speaker's intention by constructing his belief model and dynamically adjusting and expanding the model as the conversation progresses. Social agents using the BDI paradigm consider other agents' (including humans) beliefs and intentions and behave cooperatively.

If all agents had unlimited processing power and complete model of the beliefs, desires and intentions of other agents (ideal case), it would be possible for each agent to justify current actions of other agents and generate a trajectory of their future actions. Bandwidth limitations make it hard for agents to continuously receive communication about the status of all other agents and it is difficult for agents to perform local processing, reason about activities of other agents and use those inferences to modify local processing. It is not possible for agents to update other agents' beliefs in real time, particularly for a colony of large size. The model of multiagent conversational state in [Nagao & Takeuchi 94] allow each agent to have knowledge of belief space of every other agent, leading to *N (N - 1)* interactions. [Nagao & Takeuchi 94] identify three types of communication mismatch - (1) Illocutionary act mismatch (A asks B, "Do you know what happened today" with the intention of knowing it from B but B says that he does know what happened and terminates the conversation. A in this case does not get the required information.) (2) Belief inconsistency (A asks B the same question and B assumes that A knows the answer.) (3) Plan inconsistency (A asks B the same question as in 1 and B assumes that A has a plan of informing about the event). I claim that each type of mismatch can be avoided by social laws like "do not make queries that state intentions ambiguously", "do not assume anything that is not explicitly stated in the query" etc. The law-based approach may result in longer queries but it will save the effort going into debugging the mismatches. Extending BDI paradigm to complex collaborative activities requires modeling of communication, joint goal and joint beliefs. The interfaces among these need to be carefully established. A set of social laws provides a single homogeneous medium to model agents. BDI based agents negotiate in case of failure to predict other agents' actions or deadlock. One advantage of social laws over negotiation is - laws can be designed for any size of some domains e.g. in the case of traffic, the number of vehicles or robots is an indicator of the domain size. The effort involved in negotiation (e.g. the time taken to reach an agreement) is however proportional to the number of agents involved in the process.

It is necessary to design social systems in such a way that the interaction among agents of the system and their interaction with the environment converges towards the desired performance. The convergence properties are easier to test if social laws are used to model agents (as illustrated in theorem 1).

It is shown here that using social laws allows us to predict long term effects of agent interactions that are not apparent if BDI paradigm is used. [Gordon 95] mentions that the state of an individual ant in a colony depends only on the sum of the weighted interactions between individuals. The theorem below states provable properties of a similar class of laws.

Let there be a $M \times M \times M$ *3-D* grid made up of cubes of side h. Let there be an agent at each point in the grid (x_1, x_2, x_3) *s.t.* $x_1, x_2, x_3 \in \{0, h, 2h, ..., \frac{M}{h}h\}$. Let f be a function of x_1, x_2, x_3, t where the triple (x_1, x_2, x_3) represents a point in the *3-D* world and t denotes time. Let $\beta_1 = \nabla^2 f, \beta_2 = \nabla^3 f, \beta_3 = \nabla^4 f, ..., \beta_{n-1} = \nabla^n f$. Let $\alpha_1, \alpha_2, \alpha_3, ..., \alpha_{n-1}$ be real constants such that not all of them are zero.

Theorem 2.

(i) The solution of any linear combination $\alpha_1\beta_1 + \alpha_2\beta_2 + \alpha_3\beta_3 + ... + \alpha_{n-1}\beta_{n-1} = 0$ corresponds to a social law for updating f that will result in an infinite deadlock, where the deadlock is reached due to $|\frac{\partial f}{\partial t}| \leq \epsilon$, $\epsilon \geq 0$ is a constant, where $\frac{\partial f}{\partial t}$ is given by

$$\sum_{i=1}^{n-1} \alpha_i \nabla^{i+1} f$$

(ii) This class of laws lead to deadlock irrespective of the radius of communication.
(iii) Changing the distribution of f locally or globally does not resolve the deadlock.

Some agents (especially learning agents) have the potential to form their own goals and intentions, to initiate actions on their own. Such autonomous agents may be used as consultants. Such agents may be employed for secretarial work. Such agents might respond to e-mails, reveal organizational details and even release personal records. Agents with more autonomy might modify records, make undesirable commitments and intentionally supply incorrect information. The conformance of behavior of these agents to their owner's expectations is more important than their intelligence. It is not clear how a pure BDI model can control such agents. Such agents need a model of expectations of their human user. These expectations can be encoded as laws to be obeyed by the agents.

[Kautz *et al* 94] describe software agents called "visitorbots" that handle activities involved in scheduling a visitor to their laboratory but raise the issues of reliability and predictability of such agents. Social laws can be used to control and model these agents. These agents can have laws like "do not delete a mail without permission of the human user", "do not forward any mail to other agents without permission of the human user". These laws in fact supply an agent with at least a partial model of other agents e.g. if all software agents are to obey the law of not forwarding a mail to other agents without permission of their human user, agents can infer that other agents will not forward any mails to them without such a permission. Such a model also allows the agents to infer that they should send a request to the concerned human user to allow his agent to forward his mail to them.

5 Other aspects

- **Meta-knowledge** Laws can be used as a medium to represent meta-knowledge. Their primary use can be in conflict resolution. [Asimov 42] proposes three laws of robotics - (1) A robot may not injure a human being, or, through inaction, allow a human being to come to harm. (2) A robot must obey orders given to it by human beings except where such orders would conflict with the first law. (3) A robot must protect its own existence as long as such protection does not conflict with the first or second law. Let us say that a robot has a behavior to extinguish fire. This behavior may be triggered when fire is created by A for the purpose of cooking food. However A certainly does not want the robot to extinguish this fire. The expectation of A can be modeled by the robot in the form of the law "do not extinguish fire created for cooking". Let us say that a softbot is asked to delete accounts of students to free up some space at the end of semester. But what about those students who have got "incomplete" grade? This erroneous behavior of the softbot can be eliminated by controlling it by the law "do not delete an account of a student who got an incomplete grade."

- **Algorithm** Webster's dictionary defines a law as a rule binding on a community, a system of such rules or an invariable sequence of events in nature. [Shoham & Tennenholtz 95] define a law as a set of constraints. [Briggs & Cook 95] define a flexible law as a set of laws giving various options to an agent. The set of options and the set of constraints can be represented as a computational procedure analogous to an algorithm. The reverse also holds - it may be possible to show that an algorithm if used by an agent does restrict its actions imposing one or more constraints, behaving like a law. This leads to the issues - What criteria should be fulfilled by a representation to qualify as a law? An attempt is made to answer this question here. Let us say that an agent can execute m actions in absence of a particular representation R, where R is any representation which when possessed by the agent can affect its decision making. Let m' be the number of actions that the agent can execute when it is supplied with R. Then R can be a law if $\frac{m'}{m}$ is less than 1 and is sufficiently small. Hence the power to *shrink the space of executable actions* is the characteristic of a law. This is a necessary criterion but not sufficient. The space of actions executable by an agent can be reduced in many other ways like reducing its degrees of freedom or modifying its environment such that fewer actions are executable. R can be a law if the computation required to execute actions recommended by R is bounded and can be carried out based on local information. If R is a law and a', a are the actions recommended by R and other representations of an agent (excluding R) respectively, then a' will receive priority over a. R cannot be a law if actions recommended by it are all that an agent can execute and these actions are executed at a very low frequency *e.g.* let us say that robots in a colony have a behavior to pick up cans and there are conflicts because of multiple robots trying to pick up same can. Let us assume that to resolve this conflict, the robots are programmed to obey the law "do not pick up a can if some other robot is approaching it". This can lead to deadlock where no robot picks the can because of the assumption that some other robot will pick it up. If a robot from this colony is isolated and programmed to pick up such cans to resolve the deadlock, then its representation will be more of a behavior or exception handler rather than a law.

• **When to use?** Here an attempt is made to identify the conditions under which social laws can be used. If agents in a colony have intentions that need similar control, this control can be provided by a social law. If there are n softbots such that i *th* softbot can delete upto a_i files for the purpose of cleaning up directories, each softbot can be controlled by the law "do not delete more than $\frac{\Sigma a_i}{n}$ files". If a user decides that no more than g files should be deleted, one can modify the above law to get "do not delete more than $\frac{g}{n}$ files". This law eliminates the need for agent modeling. In the absence of this social law, each agent would be required to keep a track of files deleted by other agents and accordingly modify its decision to ensure that no more than g files were deleted. However agents using this law will not delete any files if the number of files to be deleted is less than the number of agents (as in that case the ratio is $\frac{g}{n}$ is less than 1) proving the law to be ineffective. In that case one can use models of other agents. Using social laws to predict actions of other agents essentially involves perceiving the current situation, finding out what laws of other agents are applicable in that situation and using that information to predict their actions. The next step is to see if these actions of other agents are co-operative, harmful or neutral. If an agent A_1 is trying to push a table and sees agent A_2 nearby and A_1 knows that A_2 obeys the laws "do not try to push objects weighing more than W kg" and "avoid obstacles", then A_1 can conclude that A_2 will either be co-operative or neutral (to arrive at an exact conclusion, A_1 needs to know weight of the table it is pushing). There may be situations in which no laws are applicable, in particular when agents are doing tasks that are not controlled by laws. The utility of other agent modeling paradigms like BDI becomes apparent here. A more efficient approach would be to combine the BDI paradigm and social laws. This approach to agent modeling is best described by figure 1.

6 Conclusion

A key problem in characterizing the effects of a given agent action on the world involves how to specify which aspects of the world do not change, the *frame problem.* Social laws can be restrictive enough to alleviate the frame problem (by reducing the reasoning an agent has to do about the

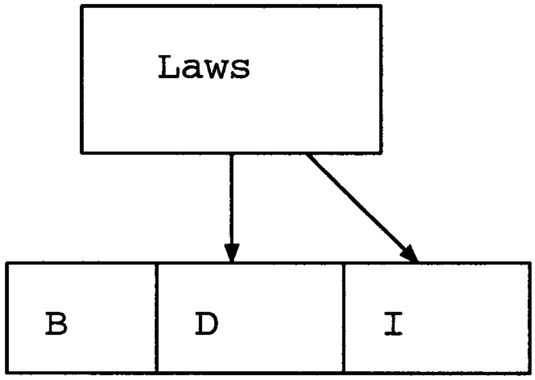

Figure 1: *Integrated approach*

actions of other agents). The ease with which social laws allow us to control the scope of reasoning is absent in BDI paradigm. I do not claim that laws can always completely replace the BDI paradigm though in simple domains like robot navigation it is possible. BDI paradigm however will be benefitted by allowing incorporation of social laws, especially for reducing negotiation, avoiding revision of beliefs and intentions etc.

The criteria defined in section 2 are essential for comparing the utility of social laws in modeling agents. As shown in section 3, using social laws expressed in proper representations allows us to test convergence properties. Most theoretical progress in DAI continues to be in the domain of homogeneous agents. An important question is - how can heterogeneous agents built by different developers using different techniques collaborate? I believe that bootstrapping a system of such agents with a set of social laws will be a useful solution.

References

[Asimov 42] Asimov I, Runaround, Astounding Science Fiction, 1942.

[Briggs & Cook 95] Will Briggs and Diane Cook, Flexible social laws, Proc. of International Joint Conference on Artificial Intelligence, 1995, 688-693.

[Durfee 92] Edmund H. Durfee, What your computer really needs to know, you learned in kindergarten, Proc. of tenth National Conference on Artificial Intelligence, 1992, 858-864.

[Georgeff & Lansky 87] Michael P. Georgeff, Amy L. Lansky, Reactive reasoning and planning, Proc. of AAAI-87, 677-682.

[Gordon 95] Deborah M. Gordon, The development of organization in an ant colony, American Scientist, Jan-Feb. 1995, 50-57.

[Hill & Johnson 94] Randall W. Hill, Jr. and W. Lewis Johnson, Situated plan attribution for intelligent tutoring, Proc. of AAAI-94, 499-505.

[Jennings 95] N. R. Jennings, Controlling cooperative problem solving in industrial multiagent systems using joint intentions, Artificial Intelligence 75, 1995, 195-240.

[Kautz *et al* 94] Henry Kautz, Bart Selman, Michael Coen, Steven Ketchpel and Chris Ramming, An experiment in the design of software agents , Proc. of AAAI-94, 438-443

[Nagao 93] Nagao K, Abduction and dynamic preference in plan-based dialogue understanding, Proc. of IJCAI-93, 1186-1192.

[Nagao & Takeuchi 94] Katashi Nagao and Akikazu Takeuchi, Social interaction: Multimodal conversation with social agents, Proc. of AAAI-94, 22-28.

[Shoham & Tennenholtz 95] Yoav Shoham and Moshe Tennenholtz, On social laws for artificial agent societies: off line design, Artificial Intelligence 73, 1995, 231-252.

[Shoham & Tennenholtz 92] Yoav Shoham and Moshe Tennenholtz, On the synthesis of useful social laws for artificial agent societies, Proc. of Tenth National Conference on Artificial Intelligence, 1992, 276-281.

Webster's New Dictionary, Promotional Sales Books, 1994.

[Tambe *et al* 95] Milind Tambe, W. Lewis Johnson, Randolph M. Jones, Frank Koss, John E. Laird, Paul S. Rosenbloom and Karl Schwamb, Intelligent agents for interactive simulation environments, AI Magazine, Spring 1995, 15-39.