

Socially Intelligent Reasoning for Autonomous Agents

Lisa M. J. Hogg and Nicholas R. Jennings, *Member, IEEE*

Abstract—Socially intelligent agents are autonomous problem solvers that have to achieve their objectives by interacting with other similarly autonomous entities. A major concern, therefore, is with the design of the decision-making mechanism that such agents employ in order to determine which actions to take to achieve their goals. An attractive and much sought after property of this mechanism is that it produces decisions that are *rational* from the perspective of the individual agent. However, such agents are also inherently social. Moreover, individual and social concerns often conflict, leading to the possibility of inefficient performance of the individual and the system. To address these problems we propose a framework for making socially acceptable decisions, based on social welfare functions, that combines social and individual perspectives in a unified and flexible manner. The framework is realized in an exemplar computational setting and an empirical analysis is made of the relative performance of *varyingly sociable decision-making* functions in a range of environments. This analysis is then used to design an agent that adapts its decision-making to reflect the resource constraints that it faces at any given time. A further round of empirical evaluation shows how adding such a metalevel mechanism enhances the performance of the agent by directing reasoning to adopt different strategies in different contexts. Finally, the possibility and efficacy of making the metalevel mechanism adaptive, so that experience of past encounters can be factored into the decision-making, is demonstrated.

Index Terms—Intelligent agents, social reasoning.

I. INTRODUCTION

SOcially intelligent agents are autonomous problem solvers that have to achieve their objectives by interacting with other similarly autonomous entities (be they other artificial agents or humans). When designing and building such agents, a major concern is, therefore, with the decision-making apparatus that they should use. Traditionally, designers have sought to make their agents rational so that they can “do the right thing” [1]. To this end, a major strand of research has adopted an economic viewpoint and looked at self-interested agents [2] that consider what action to take solely in terms of its worth to themselves. However, this is only part of the story. When an agent is situated in a social context, its actions can often have nonlocal effects. For example, the actions of different agents can conflict or result in duplication of action. This can lead to undesirable results and inefficient utilization of common resources. This may have implications for both the performance of the individual and of others. Consequently,

the benefits, to both the individual and the overall system, of a more social perspective on decision-making are beginning to be realized [3], [4] and notions of social rationality [5], [6] are emerging. The notion of socially acceptable decisions has long been of interest in human societies and in particular in the field of socio-economics [7]. However, to date, there has been comparatively little cross fertilization with agent-based computing research (except notably [8]). To help rectify this situation, this research seeks to examine the link with human style social reasoning and use its insights to explore the conflict between individual and global concerns in designing and building socially intelligent artificial agents.

In addition to balancing individual and social concerns, socially intelligent agents typically need to operate in a resource-bounded manner. They do not have unlimited time or computational resources. Moreover, such bounded rationality should be responsive to fluctuations in the amount of resources available. Hence, agents should be able to modify how they make decisions based on their current context. In our case, this means agents should be able to dynamically vary their balance between individual and social considerations depending on the amount of resources available in the system. Moreover, because computing the social effects of action choices consumes resources, agents need to be able to vary the effort they expend on this task. Thus, when resources are plentiful an agent may wish to expend a significant amount of effort computing the social implications of an important choice. However, when resources become scarce, the same agent may choose to adopt a computationally cheaper approach to the same decision.

This paper investigates the feasibility and efficacy of *socially rational decision-making*. We define a decision-making framework based on work found in socio-economics (in particular social welfare functions) that explicitly characterizes how agents can determine which action to perform in terms of a balance between individual and social concerns. By being explicit about the constituent components, the framework provides the flexibility to enable agents to dynamically tune their operation in order to be as rational as possible in the prevailing circumstances. This framework is implemented in an exemplar social setting and the ensuing empirical evaluation highlights the effectiveness of various decision-making strategies in different problem solving environments. These results are then used to design a metalevel controller that adapts the agent’s social strategy to the resource constraints that it is faced with at any given moment in time. This mechanism enables the agent to vary its degree of sociality at *run-time* according to its perception of the environment and the other agents. The benefits of the metalevel controller are shown

Manuscript received December 18, 2000; revised April 4, 2001.

The authors are with the Department of Electronics and Computer Science, University of Southampton, Southampton, U.K. (e-mail: lmjh@ecs.soton.ac.uk; nrj@ecs.soton.ac.uk).

Publisher Item Identifier S 1083-4427(01)07726-8.

through empirical evaluation. Finally, the metalevel controller is made adaptive, using a Q -Learning model, so that the agent can use its experience of previous encounters to adapt its social behavior to best fit with that of the other agents. Again, the benefits of this mechanism are highlighted empirically.

This work makes a number of important contributions to the state-of-the-art in socially intelligent agents. First, it describes a marriage of socio-economic and agent-based techniques to demonstrate how social reasoning can be effectively employed by an agent situated in a multiple agent system. Second, the effectiveness of a number of social reasoning strategies is evaluated as a dependent of the problem solving environment and the types of other agents it contains. Third, the design, implementation, and evaluation of a metalevel controller is given that enables the agent to vary its social problem solving disposition according to its prevailing resource constraints. Finally, the means and the benefits of making the metalevel controller adaptive are demonstrated. When taken together, this work can be seen as bringing together the main constituent components for designing and building socially intelligent agents.

The remainder of the paper is structured as follows. Section II details the socially rational decision-making framework and introduces the multi-agent platform used in our empirical evaluation. Section III describes the experiments we performed to assess our hypotheses about socially rational decision-making. Section IV expands the basic design to include a metalevel module that allows the agent to deal with the problems of bounded rationality and builds on the results of the experiments in Section III. Section V investigates the addition of a learning component to the metacontrol level. Related work is discussed in Section VI, followed by the conclusions and future work in Section VII.

II. INDIVIDUAL AND SOCIAL NEEDS

To date, the dominant decision-making philosophy in agent design has been to equate rationality with the notion of an individual maximizing a self-biased utility function. Thus, an agent's motivation is the maximization of benefits with regards to its own goals. However, in a multi-agent setting, for the reasons outlined above, a more social perspective on decision-making is often desirable. Traditionally, this has been achieved by making the overall system the primary unit of concern. This has the consequence of subordinating an agent's autonomy to the needs of the system. For this reason, we believe such top-down approaches fail to exploit the full potential of the agent-oriented approach; therefore, we propose an alternative means of achieving the same end. Thus, we wish to build agents from the micro to macro level, but still retain the benefits of a more social perspective. To this end, our approach is to incorporate an element of social consideration into each agent's individual decision-making function.

One means of achieving good system performance from the micro level is to incorporate all the necessary social information into a single, amorphous utility function. This is the method that would be followed by advocates of traditional decision theoretic approaches. However, such an approach conceals important details of how (and why) the agent actually reasons. Such details

are not only important for the analysis of agent behavior, but also provide a vital tool to designers when building complex systems. Therefore, we advocate an approach that provides detailed guidance as to how social agents may be constructed.

Making the theory of social decision-making finer grained in this manner is also essential for progress on the issue of bounded social rationality. Here, parallels can be drawn between conceptions of metareasoning [9] and the idea of controlling the amount of social reasoning that should be performed by contracting and expanding the set of acquaintances the agent considers in its reasoning.

A. Social Decision-Making Framework

In order to ascertain the social impact of an action, an agent needs to be able to determine the value that a state (as a result of an action) has for other individuals and possibly for the whole society. To do this, the agent needs to empathize with what others value (i.e., know how others value states and be able to make interpersonal comparisons). In this case, the social decision framework developed here builds upon and extends the idea of social rationality proposed by Jennings and Campos [5] and is based on Harsanyi's social welfare function [10]. Social welfare functions were first introduced by sociologists and they deal with choice by a group of individuals in a society. The decision maker can either be a group making a joint decision or an individual making a choice that has global consequences. The general theory of social welfare is formalized as follows. A set of agents $A = \{A_1, \dots, A_n\}$ must produce a collective decision over a set of alternative social situations $\{a_1, \dots, a_m\}$. Each individual has a preference ordering of the alternatives (this could be a simple ordinal ranking or a cardinal utility function). The group preference ordering, or social choice function, is some function W , such that W represents the preferences of the group. In Harsanyi's formulation of social choice, each individual's preferences are represented by a von Neumann–Morgenstern cardinal utility function u that obeys the standard axioms of Bayesian rationality [11]. In particular, W is defined to be the weighted summation of the individual utilities

$$W_i(a_k) = \sum_{j \in A} \lambda_j u_j \quad (1)$$

where

- i individual agent;
- λ_j weight given to agent j 's utility function in the overall equation;
- a_k action under consideration.

To be socially rational, an individual maximizes W over the different alternatives. This function represents how an individual agent may judge states of the world from a moral or social perspective by taking into consideration the benefit to others of its course of action and weighing it against its own benefits. Harsanyi [10] has shown that his social welfare function satisfies the following important postulates. First, the utility functions of all the individuals satisfy the axioms of rational behavior under risk. Second, the social preference function of the agent also satisfies these axioms. Finally, if all individuals are indifferent between any two alternatives, from the viewpoint of their

personal preferences, then an individual i will be indifferent between them from the standpoint of its social preference function. This provides an invaluable link between conceptions of individual and social rationality.

Our work adopts the basic mechanism of social choice defined in (1) (although see [12] for variants) as the means of ensuring that individual agents make socially responsible decisions. Using this formulation of social welfare, an individual has two separate preference orderings: its individual preferences (indicating an agent's own preference structure) and its social preferences (representing what an agent believes is best for the society in which it is situated). The task of analyzing how the agent decides what action to perform can be formulated by differentiating how the agent makes decisions in terms of what it expects to get out of the action itself and what effect the action has on others. The effect that an action has on other agents can be determined by examining their utility functions. This is a feasible method for agent systems consisting of tens of agents, which is our primary target. However, when the number grows to hundreds or even thousands, agents will have large acquaintance sets and the calculation of all the necessary utility functions will be computationally prohibitive in most cases. In such circumstances, the agent will need to control this calculation by considering how resource bounded it currently is (see Section IV for a further discussion of this point). A further potential difficulty with the approach we have advocated stems from the inclusion of the acquaintances' utility functions. In an ideal world, each agent would know the utility function of all the other agents. However, in practice this is infeasible. Thus, in the experiments, each agent uses an approximation of these functions, based on information it extracts from the interactions that it has with others.

In terms of the balance between individual and social needs, (1) can be rewritten as

$$W_i(a_k) = \lambda_i u_i(a_k) + \sum_{j \in \{A-i\}} \lambda_j u_j(a_k). \quad (2)$$

By setting the weighting of a specific utility function to zero, an agent can eliminate the influence of that acquaintance on its decision-making. For example, setting the agent's own weighting λ_i to zero removes any personal utility benefit consideration from its decisions. In this way, the agent can tailor its function to give more weighting to either individual or social concerns. Thus, we can say at a coarse level, the above equation becomes

$$W_i(a_k) = \lambda_i u_i(a_k) + \lambda_{\text{soc}} \sum_{j \in \{A-i\}} u_j(a_k) \quad (3)$$

where λ_{soc} is the weighting given to the social utility part of the function (i.e., how much one cares about the needs of others). To gain a perspective on the feasibility of using such a mechanism in practical systems, the above decision-making function has been implemented in a multi-agent testbed and evaluated as follows.

B. Experimental Domain

The decision-making framework was evaluated by applying it to the Phoenix fire fighting multi-agent simulation [13]. Here,




Fig. 1. Distributed Phoenix environment.

the park is policed by several firebosses, each of which is responsible for a specific geographic region (see Fig. 1).

Firebosses control a number of firefighting resources (bulldozers), that can be exchanged as and when necessary. The goal of the firebosses is to prevent fires claiming too much land. Hence, the success of an agent, and the utility it achieves, is measured as a function of land loss. The system's overall utility can also be measured in this way (i.e., land lost in the whole park). In order to use the framework described in Section II-A, we assume that the agents' utility functions are comparable.¹ A fireboss can estimate, given the current environmental conditions, how a fire is likely to spread and it can, therefore, choose the most effective allocation of its firefighting resources. These projections provide the agent with a means of determining the likely outcomes of different courses of action. Thus, for example, an agent can ascertain whether asking for a loan of resources from another fireboss is likely to be profitable. This choice is affected by the uncertainty of the reply of the other agents, and hence, the decision maker will use a probability distribution based on whether the agent believes that 1) there is a fire in the other agent's part of the park and 2) whether the other agent is likely to give resources if it is possible for it to do so.

In what follows, we assume all agents are motivated toward the goal of reducing the amount of land lost as a result of fires. If an action results in a large land loss, then that action has less utility in comparison with an action that produces lower land loss. Attaining more resources allows the agent to complete fighting the fire more quickly and so is preferred to simply fighting the fire with its current resources. However, asking for extra resources from another fireboss decreases that fireboss's effectiveness (and utility) when it is faced with a fire of its own. Therefore, agents need to be able to estimate the likely utility of a particular resource distribution to their acquaintances in terms of predicted land lost in their region. Agents do this by assessing whether a fire is likely to occur in the acquaintance's area of the park and what effect particular resource distributions will have for the agent in this context (based on the acquaintance's estimated utility function). This is then used in (3) to determine which action is socially preferred.

To illustrate the nature of social decision-making in Phoenix, consider a scenario involving three agents (A , B , and C),² each

¹This is acceptable if we assume that the agents' utility functions are comparable up to an affine transformation [7].

²For reasons of clarity, the example scenario considers just three of the systems' agents.

of which has control over two bulldozers. If a fire breaks out in A 's area, A must decide whether to ask B or C for extra resources or proceed to fight the fire with its current resources. By obtaining an extra bulldozer, A will probably reduce the amount of land it loses and hence increase its utility. However, taking a bulldozer from B or C reduces their firefighting power and hence decreases their expected utility in the event of a fire. In addition, sharing bulldozers involves overheads based on the time it takes to communicate, as well as the delay of waiting for the extra resource(s) to arrive. Furthermore, A is not certain that its request will be granted, hence time may be wasted if its request is refused. Against this background, A 's decision can be formulated in the following manner:

$$W(\text{ask}_C) = \lambda_A u_A(\text{ask}_C) + \lambda_{\text{soc}} \sum_{k \in \{B, C\}} u_k(\text{ask}_C) \quad (4)$$

where

- $W(\text{ask}_C)$ social welfare function of agent A asking agent C for a loan of resources;
- λ_A weighting A gives to its own utility function;
- λ_{soc} weighting given to the social part of the utility equation;
- u_k expected utility function of fireboss k .

The utility functions u_A and u_k capture the domain specific considerations detailed above, such as the likelihood of success and the communication overheads involved in requesting a bulldozer. Each agent uses a discrete probability distribution to represent the likelihood that a resource request is accepted. Initially all agents use the same distribution, but these are updated over time as a result of interaction and resource exchange between the agents (e.g., when resource requests are accepted or rejected). Agent A will compute similar formulations for its other options (ask_B and fight the fire with current resources) and then select the action that maximizes the social welfare.

III. EVALUATING THE EFFECTIVENESS OF SOCIAL DECISION-MAKING

Our aim is to evaluate the effectiveness of different decision-making attitudes under varying levels of resource pressure. To this end, the experimental control variables are the following.

- The utility function weightings λ_i and λ_{soc} from (3). These values can be altered to implement a wide range of decision-making strategies. Here we use the following values:
 - 0:1 *Selfless*: care nothing about their own utility. These agents will not ask for help, but will grant any resource request that they receive (provided they are able to do so);
 - 1:0 *Selfish*: care only about their own utility. They will ask for resources from others, but will not loan out their own resources;
 - 0.5:0.5 *Balanced*: place equal concern on individual and social needs;
 - 0.25:0.75 *Social tendency*: are more social than selfish;
 - 0.75:0.25 *Selfish tendency*: are more selfish than social.

- The number of firefighting resources each agent controls, i.e., one, two, or three bulldozers. In this context, three represents as many as the agent ever needs to fight a single fire).
- How fast each fireboss can make a decision (slow, medium, or fast).

Resource pressure is exerted by 1) altering the number of resources an agent controls and 2) manipulating how quickly an agent can make its decisions. A given simulation run involves specifying when and where fires occur, how large they are initially, and what the environmental conditions are (wind speed etc.). Each run is randomly generated, involving between two and six fires. Statistically significant results are obtained by averaging over 12 runs. There are six firebosses in these experiments; although this number represents a fairly small multi-agent system, it nevertheless provides a reasonable indication of the feasibility of our decision-making mechanism.

Two broad types of experiments were carried out. First, ones where all agents used the same decision-making attitude (Section III-A). This is to ascertain the performance profile of the various attitudes under varying resource constraints. Second, those in which agents have different attitudes (Section III-B). This is to ascertain the robustness of the decision-making attitudes against a range of heterogeneous opponents. In all the experiments, individual utility is measured as a function of land lost and system utility is measured as a function of the land lost in the entire park. The latter is calculated as the aggregation of the land lost by all the firebosses.

A. Homogeneous Agent Societies

These experiments seek to test the following hypotheses about social decision-making.

- 1) The performance of the self-biased attitudes (selfish, selfish tendency) will deteriorate, both at the individual and system level, as the number of resources in the system are reduced.
- 2) The performance of the society-based attitudes (selfless, balanced, and social tendency) will slowly degrade as resources are reduced, but not as much as the self-interested types.
- 3) Balanced agents will achieve the best system level performance since their decisions are based upon the welfare of the whole system.

Fig. 2 shows how the average individual performance, for one sample agent, of the different agent strategies is affected when resource pressure is exerted on the system. Runs using different rates of thinking (the other form of resource pressure) produce similar results but, because of space limitations, are not shown here. It can be seen that self-biased agents, as predicted in hypothesis 1, perform poorly in times of scarce resources. This is because they waste time trying to obtain extra resources when 1) it is perhaps unnecessary and 2) the system is comprised of agents of the same type that will not provide assistance in any case. Furthermore, we can see their performance degrades more rapidly than those with social tendencies as stated in hypothesis 2.

The selfish tendency strategy is better than the purely selfish one because requests for extra resources are only made if they




Fig. 2. Individual agent performance.

are likely to have a noticeable benefit. The surprising result is that agents with a social tendency perform very badly in times of scarce resources. This is due to the fact that in times of scarce resources when resources are more valuable, the value of those extra resources outweighs any social considerations an agent may have and so requests are more likely to be made. In turn, the acquaintances are less likely to loan out their resources as it is too costly for them on an individual basis. This introduces a delay in the fire fighting and so means more land is lost. Balanced agents perform well since all utilities are considered equally and so the costs of asking for resources and loaning them out play an important role in the decision. This means balanced agents ask for and loan out resources, but only if it is clearly beneficial. In times of plentiful resources, the performance of the different types becomes less disparate since agents generally have sufficient resources to minimize the impact of social interchange.

Fig. 3 shows the cumulative land loss of the entire system. Here, agents with social tendencies generally perform well as they explicitly attempt to assess the system-wide implication of their choices. We can also see that balanced agents perform the best (hypothesis 3) as they work toward overall utility maximization. However, selfless agents perform worse than the balanced or social tendency agents because they miss the opportunity of attaining available resources from elsewhere, i.e., balanced/social tendency attitudes. They do, however, perform better than the self-biased strategies as they do not waste time asking for resources unless they really need to, i.e., when they have no resources, and simply get on with the task at hand.

B. Heterogeneous Agent Societies

To investigate the performance of a system comprising of agents using different strategies, the runs described for homogeneous societies were repeated using different percentage mixtures of the various strategies. In particular, different percentages of selfish agents (25%, 50%, and 75%) were introduced into societies of the other decision-making attitudes with the resource pressure kept at a constant level. We are especially interested in the impact of selfish agents since these should have the greatest detrimental effect on the performance of socially rational societies. To this end, we wish to explore the following hypotheses.

- 1) The performance of selfless agents will decrease rapidly as more selfish agents are introduced.




Fig. 3. System performance.

- 2) The performance of balanced agents will be resilient to the introduction of selfish agents.
- 3) Mixing selfish and socially motivated agents in one society may produce system performance that is superior to that of the homogeneous societies of either type.

Fig. 4(a) shows both the average individual performance of selfless and selfish agents in societies in which the percentage of selfish agents is steadily increased and Fig. 4(b) a similar graph but with societies that are a mixture of balanced and selfish agents. The individual performance of both the selfless and the balanced agents suffer as the number of selfish agents is increased. However, the balanced agents are less susceptible to the increase in the number of selfish agents since they have an inbuilt concern for their own utility (hypothesis 5). This means they will not unquestioningly give resources to others if they can profit from retaining them. It can also be seen that the performance of the selfless agents decrease more rapidly than the balanced agents as more selfish agents are introduced (hypothesis 4).

Fig. 5 demonstrates how the mixed societies perform on a system level. The gradual introduction of more selfish agents decreases overall system performance for both mixtures. However, the society consisting of balanced agents shows a more steady decline in performance than the one containing selfless agents. Again this occurs because balanced agents are more concerned for the overall system and not just for individual or altruistic concerns. One point to note is the initial performance improvement of the selfless/selfish society. When there are a small number of selfish agents, and several selfless agents willing to accede to requests, overall performance improves since resources in the system are being distributed more effectively than would be the case if the system consisted solely of selfless agents. This can be related to hypothesis 6 where we expected that system performance would actually improve with some mixtures of agents. As the number of selfish agents increases, however, there are fewer opportunities for these agents to gain resources, so performance again deteriorates.

The above results demonstrate the advantage of considering both individual and social needs when making decisions. They also show the factors that can affect the outcome of the decision. For example, an agent adopting different strategies in different resource bounded environments can produce different performance characteristics. The composition of strategies within the system also has an impact on performance. The results shown above indicate that using some mixtures of different strategy types can produce better results than others. Finally, agents do




Fig. 4. Individual performance in (a) selfish/selfless and (b) selfish/balanced societies.




Fig. 5. System performance of heterogeneous societies.

not *learn* how to best interact with other agents. Rather, they simply follow the same strategy regardless of whether a positive or a negative response was obtained from an agent in a previous encounter.

Given these results, we modified the design of the Phoenix agents to take the above points into consideration. A metacontrol mechanism was designed to sit above the agent's decision-making mechanism to aid it to reason about what is the best strategy to adopt depending on what the resource context is. This metalevel takes information about the number of resources an agent has and the state of the environment to determine what it needs to do to tackle its fires.

IV. METACONTROL OF SOCIAL REASONING

We can identify several ways in which the decision-making mechanism of the agents can be changed to improve their performance. First of all, by adopting a static decision-making strategy, the agent can either miss out on gaining extra help because it is too selfless, or persist in trying to attain resources when it is obvious that it is wasting its time (this can be seen in the selfish strategy case in Figs. 2 and 3). This suggests that there are times when having the ability to dynamically vary one's strategy is useful. Second, agents may waste valuable time measuring the social implications of their actions, when they should be taking action. For example, when the agent is faced with a serious fire that needs immediate action, time can be wasted by calculating the full social implications of all the different alternatives. A more appropriate solution is to minimize the calculation of social welfare to only include

relevant or important agents. Finally, agents currently do not learn from their interactions with others, so they do not take into account whether asking a particular agent for resources is more profitable than asking others. The sending of such requests and the waiting for replies also results in time being wasted when the outcome of the request is failure to secure extra resources. Agents can thus save time by pinpointing the agents with whom they can interact successfully as well as determining what strategy may be appropriate given the responses they have received from others in the system.

Based on these observations, the design of the agent was modified to include a metalevel that controls the amount of social reasoning that is performed by the agent. Metareasoning has been used in a variety of systems to overcome problems of bounded rationality [9]. The general idea is that a computational level (metalevel) sits above the basic decision-making mechanism of the agent and controls how much computation should be devoted to deciding what to do. In our case, fires occur at different times and with varying ferocity. Fireboss agents may therefore need to make quick decisions as to how to tackle a fire when fires are high risk and spread quickly. However, for less serious fires, they may be able to plan firefighting in a more detailed and reasoned way.

A. Metalevel Architecture

From the previous set of experiments, the following summarizes the points that were noted to affect the performance of the fireboss agents.

- How resource constrained the agents are in terms of
 - 1) how many resources they have (zero, one, two, or three bulldozers);
 - 2) how many fires there are and what is the ferocity of each (this is related to the number of bulldozers that are available).

Together these represent how resource-bounded the agent is. Performance improves the more resources that an agent has at its disposal (Figs. 2 and 3).

- The constituency of the types of agents in the system. This is because the performance of the agents is related to the possibility of loaning resources from other agents (Figs. 4 and 5). Attempting to loan resources from a selfish agent

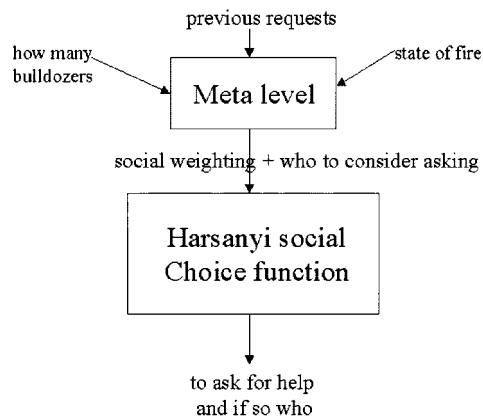


Fig. 6. Metalevel design.

is not likely to be successful and results in a waste of time asking.

These points contribute to the metalevel design (see Fig. 6). The figure shows how a metalevel was designed to sit above the decision-making mechanism of the agent controlling the strategy adopted when making decisions. Information about the state of the environment such as number of bulldozers and state of the fire is used to choose the social weighting, the peers to consider asking, and peers to include in the calculation of social utility. This information is used to change the focus of the decision-making strategy by adapting the amount of social reasoning that is undertaken (changing the peers considered) as well as modifying the weights used in the equation. In more detail, the metalevel uses the following information:

- the number of resources an agent has at its disposal;
- the environmental conditions (this provides information about the state of the environment, such as wind speed and direction);
- classification of fires (this gives the agent a measure of how serious the fire is); this measure is based on the initial size of the fire and how fast it is predicted to grow; fires are classified to be of low risk, medium risk, or high risk);
- the previous requests for this fire (this ensures that the agent does not ask the same firebosses again).

To test whether this design offers an improvement on the basic strategies compared in the first set of experiments, the following empirical evaluation was performed.

B. Empirical Evaluation

To test the effectiveness of this extra reasoning level, the following hypothesis was proposed.

- 1) Adding a metalevel component that helps the agent direct its reasoning will improve the performance of the individual agent over corresponding agents that do not possess such a component. This improvement will be apparent at both the individual level and the system level.

To test this hypothesis, experiments were carried out to compare the performance of agents who adapt their decision-making depending on how resource constrained they were to other strategies. In order to test how effective the metacontrol strategy is, it was compared to two strategies from the previous set of experiments: balanced (social) and selfish. These were




Fig. 7. Metalevel: Individual agent performance.




Fig. 8. Metalevel: Overall system results.

chosen in order to perform a comparison with the best and worst strategies previously tested.

Fig. 7 shows the results of the experiments to compare selfish, social, and metalevel strategies. The graph shows the land lost for an individual agent averaged over a set of firefighting scenarios and over different levels of resource availability. It can be seen, as hypothesized, that implementing a control level above that of the basic reasoning level does indeed produce an improvement in the performance of the individual agent. This is true over the scenarios when the fireboss agents have one, two, or three bulldozers at their disposal. Not only this, but the improvement in performance is quite marked; in some contexts, it represents an almost 50% improvement. This is due to the fact that the agent is adapting its strategy when faced with different contexts, which it is not doing in the other two strategies. This is useful as the agent may not always be able to attain resources from others since they are fighting their own fires, or have adopted a selfish attitude due to the high probability of another fire occurring. In addition, calculating the social utility over only a subset of the possible number of agents affected reduces the amount of time that the agent spends reasoning, so the agent gets on and fights the fire quicker. This subset, also directs the agent to reason about more profitable interactions with others, such as those who would be more inclined to lend it resources. Again this reduces the amount of calculation that needs to be done on the social welfare as there are fewer action alternatives to consider.

Fig. 8 shows the variation on system performance of the selfish, social, and metacontrol strategies. Here, the system as a whole also performed better when the agents are changing their strategies depending on what resource context they find themselves in. All metasocial agents attempted to adapt their strategy based on 1) the strategies of others and 2) how resource

constrained they were. This meant that there was a much more overall efficient utilization of resources (especially time) in the system.

The above results show that controlling what reasoning is done has a marked effect on the performance of the agents. By taking into account what the state of the environment and availability of resources are, the agent can tailor its reasoning to adapt to different situations. Thus, when resources are scarce, the agent can adopt a strategy that reduces the amount of computation that needs to be done and indicates what agents may be more likely to lend resources when requested to. When resources are plentiful, it can include all agents in its calculation of social welfare, ensuring that the full implications of its actions are considered. The next section considers the affect of further enhancing the metalevel control of the agent by adding a learning element to its design.

V. ADAPTIVE METALEVEL CONTROL OF SOCIAL REASONING

The previous section considered the improvements that could be made to the Phoenix agent design. One of the weaknesses of the original system is that the agent adopts a static strategy when making decisions and that it would be appropriate for the agent to adapt its reasoning to deal with different resource bounded contexts. The experiments described in Section IV dealt with this issue. A second weakness was the fact that agents do not learn from their previous interactions. Thus, agents repeatedly ask for a loan of resources from an agent that is simply not willing to accede to the request. To overcome this, it is clear that learning is needed to allow the agent to further improve its reasoning by utilizing information about the success or failure of previous interactions. The agent can learn what sort of responses it is likely to get from others by using the replies it has received from previous resource requests. This information can then be used to determine what agents are likely to give the agent extra resources and those that will not and to help the agent compile a list of agents that it may ask for help and receive a favorable response.

In this work, the method of learning chosen is reinforcement based, and in particular Q -learning [14]. This was adopted because of its natural use of feedback information from acting within the environment from the fact that learning can be done as the agents interact together and that no explicit model of the environment is necessary.

The basic idea is that the agent acts within its world to achieve its goals by taking actions $a \in A$ which allows it to move within the state space S . It receives feedback r_n from its various actions when it traverses from one state s_n to another s_{n+1} (see Fig. 9). The goal is for the agent to learn a control policy $\pi : S \rightarrow A$ that chooses an action a_i in state s_i that maximizes the accumulative reward. The accumulative reward can be defined as the sum of the rewards from the current state to the goal state.

In Q -learning, an agent keeps a table of action state pairs, $\langle a, s \rangle$ that tell it the value of taking action a in state s . In Phoenix, fireboss agents can learn who are the best firebosses to ask for a loan of firefighting resources. Each fireboss will have a Q -value for asking all other firebosses in the system. Every time an agent requests resources from another fireboss,




Fig. 9. Basic Q -learning model.

the agent can update its Q -value representing the value of asking based on whether the agent says yes or no. This information can then be used by the agent to help decide who to ask for resources. Q -values are updated by the following:

$$Q_k(a) = Q_{k-1}(a) + \alpha(r_{k+1} - Q_{k-1}(a)) \quad (5)$$

where

- Q_k Q -value of the k th action trial;
- r reward given after the action has completed;
- α learning rate of the agent.

If a low learning rate is used, the agent is slow to react to changes in the environment, whereas a high learning rate means that the agent reacts quickly to changes.

A simple means of choosing what action is right to perform, would be to simply choose the action that has the best Q -value. In doing this, the agent may commit itself to choosing actions which initially have high Q -values, perhaps failing to explore other actions which may in the long run have higher Q -values and so be more profitable. Therefore, for action choice, it is common to use a probabilistic approach [15].

$$P(a|s_t) = \frac{\exp^{Q(s_t, a)/\tau}}{\sum \exp^{Q(s_t, a)/\tau}} \quad (6)$$

where $P(a)$ is the probability of choosing action a in state s_t and τ is a temperature used to produce different degrees of exploration. A higher value of τ means that the agent chooses actions with more equal probability, so is more inclined to explore different action alternatives. A low value of τ produces behavior which sees the agent choosing the more highly valued actions and exposes the agent to the possibility of being stuck in local maxima.

A. Experiments

Experimentation with the Q -learning was split into two parts. The first part of this phase of experimentation was to test how agents using Q -learning in conjunction with the metalevel control compared to agents who simply used the metalevel control. Hence, in the first set of experiments the learning rate and exploration was kept constant to simply see how the performance compares. Each fireboss agent keeps a table of Q -values for each of its peers. Each time it makes a request to a certain peer, it updates the Q -value of that peer using the amount of land that




Fig. 10. Comparison of metalevel and adaptive strategies: Individual results.

would be saved as a result of the request as a basis for calculating the reward.

In the experiments, the land lost was recorded for various levels of resource availability for both metacontrolled and learning agents. We propose the following hypothesis.

- 1) Providing the agents with the basic ability to learn from previous interactions provides a means of improving decision-making.

The second part of these experiments investigates the effect of changing the learning rate and the exploration constant on performance. Different values of the learning rate α and temperature τ were used. Three different learning rates were used: low, medium, and high. In addition to this, different values of τ were used to indicate a low, a medium, or a high degree of exploration. Here, the following hypotheses were adopted.

- 2) Increasing the rate of learning should improve performance since agents learn the beneficial actions more quickly.
- 3) Increasing the degree of exploration will improve performance in resource constrained situations, though there will be a point at which greater exploration will degrade performance in some situations.

In order to provide a good indication of the strengths of the learning mechanism, different compositions of agents were used in which there was a percentage of selfish agents in the system. These different mixes included

- 1) a system comprising completely of adaptive agents;
- 2) a system consisting of 50% of adaptive agents;
- 3) a system where there was a single adaptive agent.

Due to the fact that selfish agents never provide assistance, placing adaptive agents in this setting provides a means of really testing the effectiveness of learning what agents are best to request resources from. All of the results of the adaptive experiments are given below.

B. Results

Fig. 10 compares the performance of agents using a metalevel strategy and ones, which in addition, adopt the Q -learning strategy. Here, adaptive agents perform even better than the metalevel ones. This is because these agents not only adapt to varying levels of resource pressure, but also learn from previous experience to know who are the firebosses that are more likely to provide assistance. They make the assumption that agents that have been helpful in the past will be helpful in the future. The adaptive agents can thus finetune their metacontrol




Fig. 11. Comparison of metalevel and adaptive strategies: System results.




Fig. 12. Increasing α : Individual performance.






Fig. 13. Increasing α : System performance.

strategy, becoming more prudent with regards to who to ask for loans as well as who to include in their decision-making.



Fig. 11 shows the results at the system level. Again, the adaptive metacontrolled layer shows an improvement over the metacontrolled layer over all levels of resource availability. This is due to the fact that all agents learn that particular agents are better to request resources from than others.

Figs. 12 and 13 show the effect of increasing the learning rate on the performance of the individual and the system. As can be seen, agents using a higher learning rate perform better as proposed in hypothesis 2. This is because agents learn more quickly what other agents are likely to give them resources and so they can identify them more quickly and attain extra resources from them in future fires. In the case where resources are scarce (number of bulldozers is equal to one), there is a sharper increase in performance than in the other two resource cases. In such situations, it is more important for agents to be able to identify to whom they should make a request for resources, as they need to be able to get on and fight the fire quickly.

Fig. 14. Increasing τ : Individual performance.Fig. 15. Increasing τ : System performance.

The performance of the adaptive agents using different temperature values are shown in Figs. 14 and 15. As the temperature parameter is increased, the agents are more inclined to explore alternative agents to ask for resources rather than simply the ones that have been helpful in the past. Here we see that in this environment, the more exploration that is performed the worse the performance of the agents becomes. This is especially true when resources are scarce. There is, however, a slight improvement in performance when agents engage in a medium degree of exploration. This can be explained by the fact that agents can benefit from asking others but may waste time in trying out different agents, some of which may be further away from the fire. Being further away from the fire means that bulldozers will take longer to travel to the fire, in which time, the fire has expanded.

The final graphs show how the performance varies over the different values of α and τ within different compositions of systems. Fig. 16 shows how the increase of α affects the performance of the agents in different mixes of agent system. The system containing all adaptive agents performs the best out of the three mixes since all agents improve their decision-making by adopting the metacontrolled learning strategy. In the case where there is only one adaptive agent, performance is much poorer than in the other two cases. This is because the adaptive agent does not have any opportunity to improve its performance as there are no means of doing so as the system consists of selfish agents. In addition, the improvement in performance is slight as α is increased. Again, learning may teach the agent that it is better not to ask the other agents for resources, though this means that the agent relies on dealing with the fire with its

Fig. 16. Increasing α : Individual performance in different mixes.Fig. 17. Increasing τ : Individual performance in different mixes.

own resources rather than gaining an advantage from obtaining a loan of resources from others. In all cases, as the learning rate is increased, the performance of the agents improved. This is because as in Figs. 12 and 13, as the learning rate increases, agents learn more quickly what other agents are good and bad sources of gaining extra resources.

Fig. 17 shows how performance is affected by changing the temperature parameter. In the case of only one adaptive agent, increasing the degree of exploration has little effect on performance as there is little value to be gained from asking any other firebosses. When there is a 50–50 mix, increasing τ initially degrades performance then slightly improves it. This is because the agent needs to engage in a certain level of exploration in order to find the agents that are more willing to assist it. Below this level, the agent misses out on the opportunity of gaining extra resources as it is less likely to find agents that are willing to help. Above this level, the agent has more chance of being successful in finding an agent that will lend it resources since it is more likely to try a wider variety of different agents. In the all-adaptive case, increasing τ slowly degrades performance, as the agent engages in more exploration. Again, as above, increasing the degree of exploration results in the agent considering others that may be less suitable, due to their distance away from the fire.

The above set of experiments have shown that adding a meta-level control component to the agent architecture has distinct advantages over allowing the agent to follow a static decision-making strategy over different resource-bounded contexts. This shows that in order to maintain a high degree of performance, the agent needs to tailor its decision-making to correspond to the amount of resources that are available to it. By adapting its decision-making strategy the agent can choose the best course

of action applicable to the circumstances. In addition, giving the agent the ability to learn from previous interactions can provide invaluable feedback to the metalevel. The information learned can be used to fine tune the metacontrol mechanism, providing information that will enable the agent to pinpoint the agents with which it is most beneficial to interact, as well as what strategy to adopt when making decisions.

VI. RELATED WORK

Rationality has been widely debated and studied in the field of agent research [1]. Decision theory [16] has emerged as the dominant descriptive and normative theory of rational decision-making. The fundamental principle of decision theory is the maximization of the agent's utility function under certain axioms of uncertainty and utility [17]. Game theory is also concerned with the rational behavior between two or more interacting individuals [18]. Each agent has a payoff or utility function that they attempt to maximize based on the information they have about the strategy of the other individual(s). This payoff function represents the preferences of the individual, though it can be based on altruistic motives in the case where more global/social concerns are the dominant philosophy. There are, however, a number of problems with game theory with regards to the social aspects of decision-making. One is the inability to deal adequately with some social notions such as cooperation [19]. In fact, without the introduction of some binding force ensuring cooperation, the theory can produce suboptimal results, as shown by the prisoners' dilemma example. Furthermore, although both game and decision theory provide simple and attractive formalisms of individual action choice, they have been criticized on the grounds that they reveal nothing about the motivations of the agents making the decisions [19]. For example, both disciplines can produce socially acceptable results if the utility functions used incorporate some social information, but these theories provide no answers as to *how* this can be done or even *why* this should be done. This, in turn, is of little use when attempting to understand, describe, and ultimately build socially intelligent agents. Thus, we adopt some of the fundamental principles of these theories but expand these ideas to explore our ideas of social reasoning.

A consistent theme in the work of Castelfranchi [3], [20] is the concept that sociality is derived from the individual mind and social action. Social rationality, and in particular an agent's social power, is described via manipulation of dependence relationships between agents. Agents may interfere, influence, and adopt goals of their acquaintances as a result of the manipulation of these relationships. Such notions can then form the basis of a variety of social actions. Although underlining the need to explore and emphasize the social aspects of an agent's make-up, this line of work addresses the philosophical rather than practical questions of how this should be achieved. Building on this, Cesta *et al.* [21] explore the practicalities of social decision-making by experimenting with a variety of social attitudes. Their work mainly covers simple, rather rigid, agent systems and concentrates on how the introduction of exploiters into a society effects system performance. Their results are consistent with our findings regarding the introduction of what we have

called selfish agents. When exploiters (selfish agents) are introduced into their system, the performance of the system decreases, a result which is magnified as resource pressure is exerted. Although they look at some effects of resource bounds, this is not the main thrust of the work. Also, there is no discussion of how individual autonomy is balanced with social concerns in such contexts.

In [22], Brainov defined a range of social decision-making strategies that differ in their attitudes toward other agents.

- 1) Altruistic agents consider other agents in their decision-making;
- 2) Self-interested agents never consider other agents when making decisions;
- 3) Envious agents consider others, but in a negative sense.

In [23], Brainov extends this work by comparing the use of these different attitudes in multi-agent planning and negotiation. Our different social attitudes are consistent with his basic definitions, but are grounded in a particular utility configuration: that of Harsanyi's welfare function. This provides a means of moving the theory into practice and allows us to begin our investigations into resource bounded social agents.

Using the shared plans intention-reconciliation (SPIRE) agent framework, Glass and Grosz investigate how a social commitment incentive scheme, which they call the Brownie point model, affects agent performance over time [24]. An agent makes a decision based on a weighted combination of the actual value of doing the task and the brownie points it is rewarded. They manipulate this weighting to produce agents that are more group committed by giving a higher weighting to the brownie points part of the function. Their results show that agents striking a balance between group commitments and monetary gains perform better than ones who have a high level of group commitment. They also look at how environmental factors influence the performance of agents under this model, but admit that further analysis and empirical investigation is needed. Like the social rationality work presented here, they experiment with various social strategies, but differ by examining the effect on performance of how much time the agent is committed to group tasks.

Jennings and Campos [5] define a social equivalent of Newell's conceptualization of individual agent rationality that they term the principle of social rationality. Social rationality is defined as the action choice of an individual based on global concerns. To add substance to this definition, Kalenka and Jennings [25] describe several social attitudes that can be ascribed to agents under this principle. Their work provides a framework for defining the different social attitudes that an agent may possess, including helpfulness and cooperativity. However, the missing element in their work is the practical consideration of resource bounds on the performance of social agents. Their framework also restricts the level of analysis that can be performed with regards to an agent's different relationships in the society. For instance, there is no mechanism to employ when the agent finds itself as a member of multiple groups or coalitions.

More socially minded decision-making attitudes have been investigated in the socio-economic literature under the umbrella

of social welfare functions (also collective choice rules or preference aggregation rules) [7]. Here, the main emphasis is on how a group of human agents can collectively make decisions. The decision maker can either be several agents making a joint decision or an individual making a decision that has global consequences. These functions have been shown to have the advantage of Pareto optimality, but have the disadvantage that equity is not preserved in the group, i.e., the decision is not fair to everyone, for example in the case of the distribution of wealth. There are also concerns as to how the utility functions are derived and how they should be combined in an overall function to reflect group choice. These issues are also important when we consider software agents, and at present, there are no comprehensive solutions to these problems. However, we do believe that practical assumptions can be made about the origin and structure of the utility functions used by agents, as we have demonstrated in this work, and that with further experimentation into these issues, useful insights can be found.

VII. CONCLUSIONS AND FUTURE WORK

This paper has outlined the case for a more socially aware approach to decision-making in a multiple agent context and how this should be tempered to deal with problems of resource boundedness. A novel agent decision-making framework, incorporating insights from work on social welfare functions, has been devised to tackle the problem of decision-making by socially intelligent agents. This framework provides a means of describing and analyzing how an individual agent may approach the task of making socially acceptable decisions in a social system. More importantly, perhaps, is the empirical demonstration of the effectiveness of various socially aware decision functions in a range of problem solving scenarios. Our results indicate that decision attitudes based on social concerns perform better in resource-bounded contexts than the more traditional, self-interested attitudes. In particular, our results for balanced agents demonstrate the importance of considering both the individual and system consequences of decision-making. Furthermore, this work investigated the effect of having several different decision-making attitudes in the same system. Here again, we highlighted the importance and effectiveness of basing decisions on both individual and social concerns by demonstrating the robustness of balanced agents in the face of exploitation by selfish agents. These experiments also demonstrate the importance of social decision-making to the performance of the individual and the system, of the mixture of strategies used by the participating agents, and how this must be considered in the individual's decision-making function.

This prompted an addition to the decision-making machinery of the agent that took the above points into consideration. A metalevel was implemented to help the agent determine what strategy was best to follow in a certain context. The results show that the addition of such a metalevel controller produces an improvement in both individual and system performance. This controller enables the agent to be aware of what current resource pressures it must address and then to adapt its reasoning to handle this. In Phoenix, this means the agent modifies how

it considers others in the system when making decisions. This is in terms of whether 1) it should consider asking others for help or 2) whether it should loan resources to other agents. Resource pressures dictate that the agent needs to address the degree of social reasoning it should undertake. For example, two extremes would be whether it should eliminate social reasoning altogether, i.e., consider only its own benefit, or take into account everyone, i.e., think about the value all the other agents attribute to a particular course of action. Thus, the metalevel outlined in this paper allows the agent to modify its social reasoning to eliminate agents considered in the social welfare function.

Another important aspect of agent performance is whether it retains information and learns from its interactions. To this end, this paper has also evaluated adding a learning component to the agent's reasoning mechanism. Results show that by giving agents the ability to learn from previous interactions a further increase in performance is produced. In addition to this, modifying the learning equation's parameters can affect the agents' performance.

In terms of extending our work, we need to further investigate how socially intelligent agents can dynamically build relationships with one another and then use this knowledge to learn how to operate more efficiently. An example of how this could be achieved was given in this paper in the form of learning who it was worth asking for resources. Another important point to consider is how an agent decides what peers it includes in its social welfare function. This is especially useful when the agent is faced with heavy time pressures since it need only perform calculations for the acquaintances that it deems important, but then the problem is to determine which peers are important or relevant to the decision. We are also interested in how agents manage activities between different subgroups or coalitions that they might be a part of. At the present, only the balance between individual and system concerns have been investigated. We would like to explore in a more detailed way, how an agent balances its own needs and the needs of the various groups of which it is a part.

Finally, we believe that our socially rational agents are ideally suited to participating in hybrid systems in which there is a mixture of humans and artificial agents working together (e.g., in computer supported cooperative work or group decision-making applications). In such systems, the artificial agent needs to be able to act both to achieve individual objectives and cooperate with humans in order to complement their problem solving activities.

ACKNOWLEDGMENT

This paper is a significantly revised and extended version of [26], which was presented at the Sixth International Workshop on Agents Theories, Architectures, and Languages.

REFERENCES

- [1] S. Russell, "Rationality and intelligence," *Artif. Intell.*, vol. 94, no. 1, pp. 55–57, 1997.
- [2] J. Doyle, "Rationality and its role in reasoning," *Computat. Intell.*, vol. 8, no. 3, pp. 376–409, 1992.
- [3] C. Castelfranchi, "Social power a point missed in multi-agent, DAI and HCI," in *Decentralised AI*, Y. Demazeau and J. P. Muller, Eds. Amsterdam, The Netherlands: Elsevier, 1990, pp. 49–62.

- [4] B. Grosz, "Collaborative systems," *Artif. Intell. Mag.*, vol. 17, pp. 67–85, 1996.
- [5] N. R. Jennings and J. Campos, "Toward a social level characterization of socially responsible agents," *Proc. IEE Softw. Eng.*, vol. 144, no. 1, pp. 11–25, 1997.
- [6] L. Hogg and N. R. Jennings, "Socially rational agents," in *Proc. AAAI Fall Symp. Socially Intell. Agents*, 1997, pp. 61–63.
- [7] K. P. Corfman and S. Gupta, "Mathematical models of group choice and negotiations," in *Handbook in OR and MS*, J. Eliahberg and G. L. Lilien, Eds. Amsterdam, The Netherlands: Elsevier, 1993, pp. 83–142.
- [8] R. Axelrod, *The Complexity of Cooperation*. Princeton, NJ: Princeton Univ. Press, 1997.
- [9] S. Russell and E. Wefald, *Do the Right Thing*. Cambridge, MA: MIT Press, 1991.
- [10] J. C. Harsanyi, *Rational Behavior and Bargaining Equilibrium in Games and Social Situations*. Cambridge, U.K.: Cambridge Univ. Press, 1977.
- [11] L. J. Savage, *The Foundations of Statistics*. New York: Wiley, 1954.
- [12] J. F. Nash, "The bargaining problem," *Econometrica*, vol. 28, pp. 155–162, 1950.
- [13] P. R. Cohen *et al.*, "Trial by fire: Understanding the design requirements for agents in complex environments," *Artif. Intell. Mag.*, vol. 10, no. 3, pp. 32–48, 1989.
- [14] C. Watkins and P. Dayan, "Q learning," *Machine Learning*, vol. 8, pp. 279–292, 1992.
- [15] T. M. Mitchell, *Machine Learning*. New York: McGraw-Hill, 1997.
- [16] J. von Neumann and O. Morgenstern, *Theory of Games and Economic Behavior*. Princeton, NJ: Princeton Univ. Press, 1944.
- [17] S. Russell and P. Norvig, *Artificial Intelligence: A Modern Approach*. Englewood Cliffs, NJ: Prentice-Hall, 1995.
- [18] T. C. Schelling, *The Strategy of Conflict*. Cambridge, MA: Harvard Univ. Press, 1963.
- [19] C. Castelfranchi, "Limits of economic and strategic rationality for agents and MA systems," *J. Robot. Autonom. Syst.*, vol. 24, pp. 127–139, 1998.
- [20] —, "Modeling social action for AI agents," *Artif. Intell.*, vol. 103, pp. 157–182, 1998.
- [21] A. Cesta, M. Micheli, and P. Rizzo, "Effects of different interaction attitudes on multi-agent system performance," in *Proc. MAAAW*, W. Van de Welde and J. Peram, Eds., 1996, pp. 128–138.
- [22] S. Brainov, "The role and the impact of preferences on multiagent interaction," in *Intelligent Agents VI*, N. Jennings and Y. Lesperance, Eds. New York: Springer-Verlag, 1999, pp. 349–363.
- [23] —, "Altruistic cooperation between self-interested agents," in *Proc. 12th Euro. Conf. Artif. Intell.*, W. Whalster, Ed., 1996.
- [24] A. Glass and B. Grosz, "Socially conscious decision making," in *Proc. BISFAI (Bar-Ilan Symp. Foundations Artif. Intell.)*, 1999.
- [25] S. Kalenka and N. R. Jennings, "Socially responsible decision making by autonomous agents," in *Cognition, Agency and Rationality*, K. Korta, E. Sosa, and X. Arrazola, Eds. Boston, MA: Kluwer, 1999, pp. 135–149.
- [26] L. Hogg and N. R. Jennings, "Variable sociability in agent based decision making," in *Proc. Sixth Int. Workshop Agent Theories, Architectures, Languages*, Y. Lesperance and N. R. Jennings, Eds., 1999, pp. 276–289.
- [27] G. Boella, R. Damiano, and L. Lesmo, "Cooperating to the groups utility," in *Intelligent Agents VI*, N. Jennings and Y. Lesperance, Eds. New York: Springer-Verlag, 1999.
- [28] A. K. Sen, *Choice, Welfare and Measurement*. Cambridge, MA: MIT Press, 1982.



Lisa M. J. Hogg received the undergraduate joint honors degree in mathematics and computer science from the University of York, York, U.K., the M.S. degree in applied artificial intelligence from the University of Aberdeen, Aberdeen, U.K., and the Ph.D. degree from Queen Mary and Westfield College, London, U.K., in 1994, 1995, and 2001, respectively.

She is a Research Assistant with the Department of Electronics and Computer Science, University of Southampton, Southampton, U.K. Her research inter-

ests include social reasoning in multi-agent systems and teamwork.



Nicholas R. Jennings (M'92) received the B.Sc. degree in computer science from Exeter University, Exeter, U.K. in 1988 and the Ph.D. degree from the University of London, London, U.K. in 1992.

He is a Professor in the Department of Electronics and Computer Science, University of Southampton, Southampton, U.K., where he carries out basic and applied research in both the theoretical and the practical aspects of agent-based computing. He has published approximately 135 articles on various facets of agent-based computing, holds two patents with two

more pending, has written one monograph, and co-edited five books.

Dr. Jennings was the recipient of the Computers and Thought Award in 1999 for his contributions to practical agent architectures and applications of multi-agent systems and the IEE Achievement Medal in 2000 for his work on agent-based computing.