Grateful Agents and Agents that Hold a Grudge*
The Role of Affective Behaviors in Sustained Multi-agent Interactions

César F. Pimentel
INESC-ID and Instituto Superior Técnico, UTL
Av. Prof. Dr. Aníbal Cavaco Silva, 2744-016 Porto Salvo, Portugal

Keywords: Affective Computing, Multi-agent Interactions, Emotions, Anger, Gratitude.

Abstract: Interactions among self-interested agents present classical challenges concerning cooperation and competitiveness. Cooperative behavior may be unappealing if unilateral cooperation represents a loss, and adverse behavior may be difficult to avoid if one agent's losses imply another agent's gains. Agents could benefit from mechanisms that promote cooperation and dissuade adverse behaviors. We propose a generic approach, where such mechanisms can emerge from the simulation of affective behaviors that are associated with the human emotions of gratitude and anger. These emotions define implicit contracts about predefined patterns of behavior that agents are capable of following, and recognizing in each other. We use a few examples to illustrate how this approach can help an agent persuade another to cooperate and become an ally, or dissuade it from adopting adverse behavior, as a result of rational decisions.

1 INTRODUCTION

In societies of self-interested agents, individual behaviors are typically a result of two driving forces: cooperation and competitiveness. Examples of these agent societies can be found in domains such as computer supported cooperative work (Baecker, 1993), electronic commerce negotiation (Kephart and Greenwald, 2002), virtual environments (including video games), social simulations (Gilbert and Conte, 1995; Gilbert and Doran, 1994; Moss and Davidson, 2001), among others (Jennings and Wooldridge, 1998).

In such domains, it is often challenging to design agents that behave efficiently: In scenarios where unilateral cooperation represents a loss, and an agent has no guarantee that another agent will also cooperate, cooperating might seem unreachable when it is seen as an irrational course of action; in competitive scenarios where an agent's loss represents another agent's gain, an agent must carefully manage its use of adverse behaviors while trying to avoid the adverse behaviors of other agents.

In the first type of scenarios, agents could benefit from a mechanism that promotes mutual cooperation. In the second type of scenarios, agents could benefit from a mechanism that dissuades others from adopting adverse behaviors. In this paper we argue that specific affective aspects of human behavior can be simulated, to produce such mechanisms, in the context of sustained multi-agent interactions.

In particular, we suggest that if an agent behaves in consistence with the emotion of *gratitude* and another agent recognizes this type of behavior, mutual cooperation may start occurring, as a result of rational decisions, in situations where it would otherwise not occur. In addition, we suggest that if an agent behaves in consistence with the emotion of *anger*, "holding a grudge" against another agent that recognizes this type of behavior, future adverse behaviors may be prevented, as a result of rational decisions. This, in fact, corresponds to reproducing the roles that these two emotions play on human social interactions.

There are other ways of addressing the challenges concerning cooperation and competitiveness, such as using program equilibrium (Tennenholtz, 2004). In many environments, however, such approaches cannot be used because program strategies are not available, mediators do not exist, or communication is not possible. Our approach can be used in such environments and, because it is inspired in human behavior, it is expected to simultaneously be effective and more accurately simulate human behavior.

*This work was supported by national funds through FCT - Fundação para a Ciência e a Tecnologia, under project PEst-OE/EEI/LA0021/2011.

F. Pimentel C..
Grateful Agents and Agents that Hold a Grudge - The Role of Affective Behaviors in Sustained Multi-agent Interactions.
DOI: 10.5220/0004334003240329
In *Proceedings of the 5th International Conference on Agents and Artificial Intelligence* (ICAART-2013), pages 324-329
ISBN: 978-989-8565-38-9

2 COOPERATION AND COMPETITIVENESS

A common problem in multi-agent systems is how to get agents to cooperate when each agent is self-interested and has no guarantee that others will co-operate as well. A classical example from game theory, that illustrates this problem, is known as the *Prisoner's Dilema* (Tucker, 1950; Straffin, 1980).

We recall that in the prisoner's dilemma two agents must separately choose to either cooperate or defect. Each agent obtains best payoffs by defecting, both when the other agent defects and when it cooperates. Defecting is, therefore, the dominant strategy and, consequently, the rational course of action for both agents. However, if the agents would both cooperate, instead of defecting, they would obtain higher payoffs.

In game theory terms, one can say that utility is being wasted because the outcome where both agents defect is not Pareto efficient; there is an outcome with higher payoffs for both of the agents. The problem is that, to achieve this outcome, the agents must apparently act irrationally.

When scenarios include more than two interacting agents, the problem of cooperation becomes even more interesting. In addition to the question of whether or not to cooperate, there is also the question of whom to cooperate with. Teaming up and forming implicit alliances is often the way to obtain the highest payoffs (for the allied agents, obviously). The problem, again, is that an agent is never sure that the other will cooperate and act as an ally.

Besides cooperation, another important issue in multi-agent interactions is competitiveness. In particular, we are interested in addressing the problem of making an agent successful in scenarios where other agents may act adversely, to increase their own payoffs. The question we address here is: How can an agent dissuade other agents from acting in ways that decrease that agent's payoff? In interactions between two agents, being cooperative or being non-aggressive may be seen as the same concept, but that is not the case when we consider interactions with more than two agents.

3 GRATITUDE AND ANGER

We start this section with a personal story, to help illustrate the motivation for the ideas in this article. A group of five friends, occasionally, used to play a board game of military strategy. In this game, each player starts out with control over certain territories, and may attack the territories of other players to gain control over them. It is important to note that there is often more than one way for a player to win the game, so one usually has flexibility in choosing one's goal territories.

The first time that the game was played, one of the players, David, displayed a behavior typically known as "holding a grudge"; when he was attacked by another player he retaliated and continued attacking that player, repeatedly, until the end of the game, even in situations where such attacks did not seem to be advantageous, for David. In the end, both David and his initial attacker failed to win the game, but this was not the end of the story.

In subsequent, games David continued displaying behaviors consistent with the personality of someone that easily "holds a grudge" against whoever attacks him. The other players quickly realized this and (despite finding amusing to show anger in the context of a game) started being biased against attacking David, because being the target of his constant attacks makes the task of winning the game a difficult one. Players did not entirely stop attacking David, but they started taking his personality into account when choosing which opponent to attack. This gave David a great advantage in the game, because he was rarely attacked.

David's emotional behavior might be seen as irrational, in the sense that he typically fails to win the game when he decides to continuously attack a specific player until the end of that game. However, if we consider the long term effects of his behavior over a sequence of games, we conclude that it is a rational behavior, because it gives him an advantage over the other players, increasing his chances of wining games.

"Holding a grudge" is an affective behavior associated with the emotion of *anger*. In the appraisal theory of emotions known as *OCC* (Ortony et al., 1988), anger is the emotion that results from an event that has undesired consequences for oneself and was caused by the actions of another agent. Conversely, when an event caused by the actions of another agent has desired consequences for oneself, the resulting emotion, according to OCC, is *gratitude*.

Just like anger was useful to David in a competitive situation, so can gratitude be useful to promote cooperation. For instance, consider a situation where helping someone makes you slightly worse off but greatly improves the other person's situation, and vice versa. Helping would be irrational, if it had no future implications. But if the person that is helped behaves grateful and returns the favor, then mutual cooperation can be initiated and both people will benefit. In the strategy board game, this could lead to the appearance of unofficial alliances among players.

It is important to notice that for anger and gratitude to have the effects described above, it is necessary that the affective nature of one's behavior is recognized by others and attributed to one's personality. David would not have been successful if the other players would not have realized that he is acting out of anger and that his personality is of someone that easily holds a grudge.

4 GRATEFUL AGENTS AND AGENTS THAT HOLD A GRUDGE

From the point of view of *Affective Computing* (Picard, 1997), human emotional phenomena are a promising source of inspiration to address problems in *Artificial Intelligence*, because such phenomena are generally helpful to humans (Damásio, 1994; Frijda et al., 2000). In this paper we propose that, just as humans benefit from acting on anger or gratitude, artificial agents may also benefit from adopting such behaviors, in the context of sustained multi-agent interactions. In particular, we propose that behaving grateful may promote mutual cooperation while "holding a grudge" may dissuade other agents from acting adversely.

To achieve the discussed effects, we propose that agent architectures contemplate three key abilities:

Affective Behavior. The ability to decide and act in congruence with an emotion. In particular, when an agent's situation is improved by the actions of another agent, it should act in congruence with gratitude, by returning the favor (possibly more than once); when an agent's situation is harmed by the actions of another agent, it should act in congruence with anger, by retaliating (possibly more than once).

Identification of Affective Behavior. The ability to identify that another agent's actions are the result of an affective state, and associate, with that agent's personality, the predisposition for such affective behaviors. This ability is within the scope of a theory of mind (Baron-Cohen, 1995), i.e. the ability to attribute mental states to oneself and others. In particular a simplistic theory of mind is required, to identify the gratitude and anger-congruent behaviors described in the previous item.

Reasoning about Personality. The ability to make decisions that account for the other agents' personalities. Based on the aspects of personality that were identified according to the previous item, the agent should be able to extrapolate and predict the behaviors of other agents under certain conditions. These predictions should be taken into account when deciding the best course of action.

In the next sections we present examples of using this approach, to illustrate its potential benefits. In Section 4.1 we use two scenarios that focus on gratitude behaviors, one with two agents and another with three agents. Analogously, in Section 4.2 we use two scenarios that focus on anger behaviors, one with two agents and another with three agents. All scenarios are presented using the common terminology of game theory.

4.1 Examples with Gratitude

Consider the situation where each of two agents must choose between the actions of helping the other agent and running away. Helping requires a small effort (decreasing the helper's utility by 1), but greatly improves the other agent's situation (increasing it's utility by 2). Running, on the other hand, doesn't affect any of the agents' utility. The resulting payoff matrix, for agents i and j is presented in Table 1.

Table 1: Payoff matrix for the gratitude 2-agents example. Utilities are shown in the order $(u_j,\ u_i)$, where u_x is the utility for agent x.

	i runs	i helps
j runs	$\mathbf{0},\mathbf{0}$	$2,-1$
j helps	$-1,2$	$1,1$

Notice that running is a *dominant strategy* (i.e., a strategy that guarantees the highest payoff for the agent, no matter what the other agent chooses to do). Consequently, there is one Nash equilibrium, marked in bold, where both agents decide to run. The problem corresponds to the well known prisoner's dilemma, because running seems to be the rational option but if both agents would cooperate (help), they would both get higher payoffs.

Now we consider several hypothetical interactions, in the above scenario, organized in sets of 10 rounds. For these interactions, we consider that agent i is prone to act on gratitude, when helped, and agent j has the ability to identify this in i's personality and to make decisions accordingly (see the proposed abilities in Section 4). Table 2 illustrates these hypothetical interactions.

In the first set of interactions both agents start by running, but once j decides to help, i returns the favor twice. j realizes that i probably acted on gratitude, and predicts that it may repeat that behavior when

Table 2: Sets of 10-rounds interactions of the gratitude 2-agents example. R denotes running and H denotes helping.

	Set #1									
i	R	R	R	H	H	R	R	H	H	H
j	R	R	H	R	R	R	H	H	H	H
	Set #2 and subsequent sets									
i	R	H	H	H	H	H	H	H	H	H
j	H	H	H	H	H	H	H	H	H	H

helped. Having this assumption, the rational behavior is to help because mutual help has a higher payoff (1) than the one that is obtained when both agents run (0). As a result, both agents help each other until the end of the set, and may continue doing so in the second and following sets, because j takes i's personality into account, when making decisions.

Now let us consider a military scenario that involves three agents. Each agent must choose to either build, improving the value of its territory, or support one of the other agents, by sending troops that help that agent in defense and attack situations. Supporting improves the supported agent's situation with no apparent gain to the supporting agent; but when two agents support each other, they perform a successful attack on the remaining agent's territory, greatly improving the situation of the attacking agents at the expense of the attacked agent. The actions affect the agents' utilities in the following ways:

- When an agent builds, its utility is increased by 1.
- When an agent x supports another agent y, and y does not support x, the utility of y is increased by 1.
- When two agents support each other, their utilities are increased by 2 and the remaining agent's utility is decreased by 2.

The resulting payoff matrix, for agents i, j and k is presented in Table 3.

Table 3: Payoff matrix for the gratitude 3-agents example. Utilities are shown in the order (u_k, u_j, u_i), where u_x is the utility for agent x. B denotes the action of building and S_x denotes the action of supporting agent x.

		$j\ B$	$j\ S_k$	$j\ S_i$
$i\ B$	$k\ B$	$\mathbf{1,1,1}$	$2,0,1$	$1,0,2$
	$k\ S_i$	$0,1,2$	$1,0,2$	$0,0,3$
	$k\ S_j$	$0,2,1$	$\mathbf{2,2,-1}$	$0,1,2$
$i\ S_j$	$k\ B$	$1,2,0$	$2,1,0$	$\mathbf{-1,2,2}$
	$k\ S_i$	$0,2,1$	$1,1,1$	$-2,2,3$
	$k\ S_j$	$0,3,0$	$2,3,-2$	$-2,3,2$
$i\ S_k$	$k\ B$	$2,1,0$	$3,0,0$	$2,0,1$
	$k\ S_i$	$\mathbf{2,-1,2}$	$3,-2,2$	$2,-2,3$
	$k\ S_j$	$1,2,0$	$3,2,-2$	$1,1,1$

There are four Nash equilibria, marked in bold, one where all agents build and three where two of the agents support each other. From these equilibrium states, mutual cooperation seems to be the best option for the two agents that support each other. The problem is that there are no alliances defined, so an agent does not know whom to expect mutual cooperation from.

Now we consider several hypothetical interactions, in the above scenario, organized in sets of 10 rounds. For these interactions, we consider that agent i is prone to act on gratitude, when supported, and agent j has the ability to identify this in i's personality and to make decisions accordingly. Table 4 illustrates these hypothetical interactions.

Table 4: Sets of 10-rounds interactions of the gratitude 3-agents example. B denotes the action of building and S_x denotes the action of supporting agent x.

	Set #1									
i	B	B	B	B	S_j	S_j	S_j	S_j	S_j	S_j
j	B	B	B	S_i	B	B	S_i	S_i	S_i	S_i
k	B	B	B	B	B	B	B	B	B	B
	Set #2 and subsequent sets									
i	B	S_j	S_j	S_j	S_j	S_j	S_j	S_j	S_j	S_j
j	S_i	S_i	S_i	S_i	S_i	S_i	S_i	S_i	S_i	S_i
k	B	B	B	B	B	B	B	B	B	B

In the first set of interactions all agents start by building, but once j decides to support i, i starts returning the favor. j realizes that i is probably acting on gratitude, and predicts that it will maintain this behavior continuously if j also does. Having this assumption, j's rational choice is to support i, because mutual cooperation has a higher payoff (2) than the one that is obtained when the agents build (1). As a result, i and j establish an implicit alliance, at the expense of k. This alliance extends to the end of the set, and may continue to the second and following sets, because j takes i's personality into account, when making decisions.

4.2 Examples with Anger

Consider a military scenario where each of two agents must choose between the actions of building and conquering. Building slightly improves the value of one's territory (increasing the one's utility by 1). Conquering consists of stealing a part of the opponent's territory (increasing one's utility by 2 and decreasing the opponent's utility by 2). The resulting payoff matrix, for agents i and j is presented in Table 5.

Notice that conquering is a dominant strategy. Consequently, there is one Nash equilibrium, marked in bold, where both agents decide to conquer. Just as was the case with the gratitude 2-agents example,

Table 5: Payoff matrix for the anger 2-agents example. Utilities are shown in the order (u_j, u_i), where u_x is the utility for agent x.

	i conquers	i builds
j conquers	$\mathbf{0,0}$	$2,-1$
j builds	$-1,2$	$1,1$

in Section 4.1, this problem corresponds to the prisoner's dilemma (in fact, the payoff matrices of the two scenarios are exactly the same). Conquering seems to be the rational option but if both agents would cooperate (build), they would both get higher payoffs.

Now we consider several hypothetical interactions, in the above scenario, organized in sets of 10 rounds. For these interactions, we consider that agent i is prone to act on anger, "holding a grudge" against whoever conquers from its territory, and agent j has the ability to identify this in i's personality and to make decisions accordingly. Table 6 illustrates these hypothetical interactions.

Table 6: Sets of 10-rounds interactions of the anger 2-agents example. B denotes building and C denotes conquering.

Set #1

i	B	B	B	B	C	C	C	C	C	C
j	B	B	B	C	B	B	C	B	B	C

Set #2 and subsequent sets

i	B	B	B	B	B	B	B	B	B	B
j	B	B	B	B	B	B	B	B	B	B

In the first set of interactions both agents start by building, but once j decides to conquer, i retaliates and attacks j (i.e. conquers) until the end of the set. j realizes that i probably acted on anger, and predicts that it might do so again in the next sets, if j decides to conquer again. Having this assumption, j knows that aiming for a payoff of 2 (when j conquers and i builds) is a lost cause because, at most, it could only be obtained once in a set. So, as long as i keeps building, j might as well aim for mutual cooperation (when both agents builds) that has a higher payoff (1) than the one that is obtained when both agents conquer (0). As a result, both agents cooperate (build) during the second and following sets, because j takes i's personality into account, when making decisions.

Now let us consider another military scenario, this time involving three agents, where each must choose to conquer land from one of its two adversaries. Agents are positioned in a circular fashion, where k is at the left-hand side of j, j is at the left-hand side of i, and i is at the left-hand side of k. Conquering improves the situation of an agent, even more so if the agent is conquering from its left-hand side opponent (the utility is increased by 2 when conquering from the right-hand side opponent, and by 3 when conquering from the left-hand side opponent). This action harms the situation of the conquered agent (decreasing its utility by 2). The resulting payoff matrix, for agents i, j and k is presented in Table 7.

Table 7: Payoff matrix for the anger 3-agents example. Utilities are shown in the order (u_k, u_j, u_i), where u_x is the utility for agent x. C_x denotes the action of conquering from agent x.

		$j\ C_k$	$j\ C_i$
$i\ C_j$	$k\ C_i$	$\mathbf{1,1,1}$	$3,0,-1$
	$k\ C_j$	$0,-1,3$	$2,-2,1$
$i\ C_k$	$k\ C_i$	$-1,3,0$	$1,2,-2$
	$k\ C_j$	$-2,1,2$	$0,0,0$

Notice that conquering from one's left-hand side opponent is a dominant strategy. Consequently, there is one Nash equilibrium, marked in bold, where i conquers from j, j conquers from k, and k conquers from i. An agent can hope to get a higher payoff, of 2 or 3, if it is not attacked (i.e. conquered) by any of the other agents. The problem is that in order to achieve this, its right-hand side opponent must be dissuaded from following its dominant strategy.

Now we consider several hypothetical interactions, in the above scenario, organized in sets of 10 rounds. For these interactions, we consider that agent i is prone to act on anger, "holding a grudge" against whoever conquers from its territory, and agent k has the ability to identify this in i's personality and to make decisions accordingly. Table 8 illustrates these hypothetical interactions.

Table 8: Sets of 10-rounds interactions of the anger 3-agents example. C_x denotes the action of conquering from agent x.

Set #1

i	C_j	C_k	C_k	C_k	C_k	C_k	C_k	C_k	C_k	C_k
j	C_k	C_k	C_k	C_k	C_k	C_k	C_k	C_k	C_k	C_k
k	C_i	C_i	C_j	C_j	C_j	C_i	C_i	C_i	C_i	C_i

Set #2 and subsequent sets

i	C_j	C_j	C_j	C_j	C_j	C_j	C_j	C_j	C_j	C_j
j	C_k	C_k	C_k	C_k	C_k	C_k	C_k	C_k	C_k	C_k
k	C_j	C_j	C_j	C_j	C_j	C_j	C_j	C_j	C_j	C_j

In the first set of interactions each agent starts by conquering from its left-hand side opponent (the dominant strategy). But since k conquers from i, right in the first round, i retaliates and attacks k (i.e. conquers from k) until the end of the set. k realizes that i probably acted on anger, and predicts that it might do so again in the next sets, if k decides to conquer from i again. Having this assumption, k is better off conquering from j than from i, because the extra utility that is obtained in conquering from i (1) does not compensate for the loss of utility suffered when i retaliates

(2). As a result, both *i* and *k* attack *j*during the second and following sets, because *k* takes *i*'s personality into account, when making decisions.

We highlight the fact that *i*'s behavior in the first set might seem irrational, because conquering from *k* produces less utility than conquering from *j* (the dominant strategy). But, in the long run, *i* gained a clear advantage, because it did not suffer any attacks during the next sets, obtaining always the payoff of 3.

5 CONCLUSIONS

We revisit classical challenges concerning cooperation and competitiveness, in interactions that involve self-interested agents. We propose an approach, to these challenges, inspired in human affective behaviors, attempting to reproduce the beneficial roles that the emotions of gratitude and anger play in human social interactions. In our approach we propose that agent architectures contemplate simplistic approaches to a) producing affective behavior, b) recognizing affective behavior, and c) reasoning about personality.

We used this approach in four examples, to show how a) acting on gratitude can promote cooperation and help form alliances among agents, and b) acting on anger can also promote cooperation, as well as dissuade other agents from having adverse behaviors toward the agent in question. After recognizing the emotional behaviors, agents decided to cooperate or to avoid adverse behavior, not as artificial decisions designed to simulate human behavior, but as the rational decisions that aimed at maximizing the overall present and future payoffs. These decisions take into account the personality of other agents, to help predict their future behaviors in specific situations.

There are numerous ways to approach cooperation and competitiveness (see, e.g., (Binmore, 1994; Binmore, 1998)). Our approach uses emotions to define implicit contracts for predefined patterns of behavior (e.g., "if you help me, i will also help you", or "if you attack me, i will also attack you"). Other patterns of behavior could be used instead of these, but we suggest following these particular patterns because they occur in human beings and, therefore, a) this approach may potentially inherit the already proven benefits that such patterns convey to human social interactions, and b) this approach may more accurately simulate the behaviors of humans, which is important for purposes such as achieving believability (e.g. in the context of synthetic characters or video games), achieving more accuraccy in social simulations, and establishing more successful interactions with human agents.

REFERENCES

Baecker, R. M., editor (1993). *Readings in Groupware and Computer-Supported Cooperative Work: Software to Facilitate Human-Human Collaboration*. Morgan Kaufmann Publishers, San Mateo, CA.

Baron-Cohen, S. (1995). *Mindblindness: An Essay on Autism and Theory of Mind*. MIT Press, Cambridge, MA.

Binmore, K. (1994). *Game Theory and the Social Contract Volume 1: Playing Fair*. MIT Press, Cambridge, MA.

Binmore, K. (1998). *Game Theory and the Social Contract Volume 2: Just Playing*. MIT Press, Cambridge, MA.

Damásio, A. R. (1994). *Descartes' Error: Emotion, Reason and the Human Brain*. Grosset/Putnam Press, New York, N.Y.

Frijda, N. H., Manstead, A. S. R., and Bem, S., editors (2000). *Emotions and Beliefs - How Feelings Influence Thoughts*. Cambridge University Press, Cambridge, UK.

Gilbert, N. and Conte, R., editors (1995). *Artificial Societies: The Computer Simulation of Social Life*. UCL Press, London.

Gilbert, N. and Doran, J., editors (1994). *Simulating Societies: The Computer Simulation of Social Life*. UCL Press, London.

Jennings, N. R. and Wooldridge, M., editors (1998). *Agent Technology: Foundations, Applications, and Markets*. Springer-Verlag, Berlin, Germany.

Kephart, J. O. and Greenwald, A. R. (2002). Shopbot economics. *Autonomous Agents and Multi-Agent Systems*, 5(3):255–287.

Moss, S. and Davidson, P., editors (2001). *Multi-Agent-Based Simulation*, volume 1979 of *LNAI*. Springer Verlag, Berlin, Germany.

Ortony, A., Clore, G. L., and Collins, A. (1988). *The Cognitive Structure of Emotions*. Cambridge University Press, New York, N.Y.

Picard, R. (1997). *Affective Computing*. MIT Press, Cambridge, MA.

Straffin, P. D. (1980). The prisoner's dilemma. *UMAP Journal*, 1:101–103. Reprinted in: Rasmusen, E., editor (2001). *Readings In Games and Information*. Wiley-Blackwell, Oxford, UK.

Tennenholtz, M. (2004). Program equilibrium. *Games and Economic Behavior*, 49(2):363–373.

Tucker, A. W. (1950). A two-person dilemma. Unpublished notes, Stanford University. Reprinted in: Rasmusen, E., editor (2001). *Readings In Games and Information*. Wiley-Blackwell, Oxford, UK.