

The More Radical, the Better: Investigating the Utility of Aggression in the Competition among Different Agent Kinds

Matthias Scheutz Paul Schermerhorn
Artificial Intelligence and Robotics Laboratory
Department of Computer Science and Engineering
University of Notre Dame
Notre Dame, IN 46556
{mscheutz,pscherm1}@nd.edu

Abstract

Aggression is wide-spread in nature and seems to serve, among others, an important role in the interspecies competition for resources. In this paper, we argue that displaying aggression as a means to signal action tendencies (in particular, the probability to continue an encounter) is beneficial for social groups and show that discriminating between “own” and “other” is more beneficial than treating “other” the same as “own”. In particular, we demonstrate that aggression plays a crucial role in strategies applied to “other”. To test the theoretical discussion, we define a seven basic agent types which give rise to 42 different *discriminating agents*, i.e., agents with different strategies for “own” and “other”. In extensive simulation studies we show that discriminating agents, which assume an aggressive attitude towards others, while playing a strategy that distributes resources fairly among “own”, are ultimately the most successful ones. We discuss the implications of these results for natural and artificial agents and conclude with a brief outlook on further studies.

1. Introduction

Competition for resources is wide-spread in nature and has been studied extensively by ethologists. Typically, conflicts or contests among individuals are considered games (in a game-theoretic sense, e.g., (Maynard Smith, 1974, Maynard-Smith, 1982), etc.), where all players want to maximize their payoff and the winner is determined after a number of rounds (e.g., after one round in games like the “Prisoner’s Dilemma”, the “Dove-Hawk Game”, or the “Game of Chicken”). Some of these games can be repeated many times (e.g., the “Iterated Prisoner’s Dilemma”

(Axelrod and Hamilton, 1981) , or the “Sequential Assessment Game” (Enquist and Leimar, 1983)) in order to study the change in strategies based on previous outcomes, while others might end terminally after one round (e.g., when both players accept the dare in a “Game of Chicken”). Given the payoff matrices for these games (i.e., the payoff for winning and losing), different strategies for playing the game can be studied. In particular, it is often possible to determine whether games have one or more *Nash equilibria* (or *evolutionary stable strategies*), which determine the best response to a best reply.

Many real life incarnations of these above games are special iterated versions, called *stopping games* ((Shmaya et al., 2003, Touzi and Vieille, 2002) where a (possibly varying) cost is associated with playing the game. Different from standard game-theoretic approaches, players typically have to pay in every round (e.g., the basic bet in poker games or the cost of fighting for a period of time). Furthermore, the payoff is not determined after each round (as is typically the case in repeated versions of the above mentioned games), but only when one player quits, which often happens when one player lacks the resources to continue the game. In the competition for food in the animal kingdom, for example, it may mean that the animal runs out of energy and dies. Consequently, it may not be possible for an individual to play what would have otherwise been an optimal strategy.

To account for the influence of the context, in which competitions for resources are embedded in nature, we view the “competitions games” as part of the larger “survival game”, in which agents have to survive for long enough to be able to procreate.¹ Specifically, we study and compare different strategies for resolving conflicts among individuals in the context of a simple discrimi-

¹Elsewhere we have called the games played by individuals in the context of larger games “embedded games” (Scheutz and Schermerhorn, 2004).

nation between “own” (i.e., members of one’s own kind) and “other” (i.e., members of another kind). The role of aggression as a means to signal action tendencies in order to resolve conflicts better and faster will be of particular interest, given that aggressive behavior is widespread in nature and that ethologists and anthropologists have repeatedly emphasized the utility of aggression (e.g., see (Lorenz, 1963, Lorenz, 1977, Brain, 1979, Eibl-Eibesfeld, 1974)).

The paper is organized as follows: we start with a description of a very general conflict scenario that has many instances in natural and artificial domains and consider different strategies for resolving conflicts. We argue that displaying aggression as a means to signal action tendencies (in particular, the probability to continue an encounter) is beneficial for social groups. We then argue that discriminating between “own” and “other” is more beneficial than treating “other” the same as “own” and that aggression plays a crucial role in strategies applied to “other”. To test the theoretical discussion, we define a seven basic agent types which give rise to 42 different *discriminating agents*, i.e., agents with different strategies for “own” and “other”. In extensive simulation studies we show that discriminating agents, which assume an aggressive attitude towards others, while playing a strategy that distributes resources fairly among “own”, are ultimately the most successful ones. We discuss the implications of these results for natural and artificial agents and conclude with a brief outlook on further studies.

2. Conflicts, Aggression, and Survival

Conflicts arise in nature because resources such as food, mates, territory, etc. are often scarce. Typically, there are more individuals interested in a kind of resource than is available in their environment, so a conflict ensues in which individuals compete for the resources they desire or need. The strategies used in such a competition can take various forms, ranging from brute force fighting (disregarding any possible signals of opponents about their behavioral disposition) in which the stronger individual is likely to win, to careful negotiations in which a decision is reached based on comparing advantages and disadvantages of the conflict situation with respect to their expected utility. It can also be carried out at different levels of social organization, from conflicts among two individuals, to conflicts among small groups of individuals (e.g., among families), to conflict between nations. In its most general form, the conflict will end in one of three ways: (1) some agents win, the others retreat, (2) nobody gets the resource (everybody loses), or (3) the resource can be shared (everybody gets a part, but not the full resource). In natural settings, individuals need to survive long enough to be able to have offspring, to which they can pass on their genes. This, in

turn, will also guarantee the survival of their kind (i.e., the group to which they belong, e.g., as determined by their genetic features).

In the context of survival, aggression levels, which depend on various external, but also on internal, individual factors (e.g., (Barlow et al., 1986)), seem to play an important role. And while aggression may serve several regulatory purposes for organisms (e.g., (Brain, 1979, Lorenz, 1977, Albert et al., 1993)), its role in social conflicts that relate to *interspecies competition* (as opposed to *intraspecies competition*) for resources is of particular importance for the present investigation.

When animals fight for resources (food, mates, territory, etc.), the contests typically involve various displays of aggression or prowess (e.g., (Adamo and Hanlon, 1996, DeCourcy and Jenssen, 1994, Hofmann and Schildberger, 2001)). These expressions of aggression, whether facial expressions, gestures, or other demonstrations of strength or determination, can be construed as signals that communicate the probability with which an animal will continue to fight (if put in a contest): roughly speaking, the stronger the display, the higher the likelihood that the animal will keep fighting after a given time. Several questions arise: (1) is displaying this information beneficial, (2) what strategy should animals play, (3) is there an optimal strategy, and (4) is cheating (i.e., displaying wrong action tendencies) beneficial (at the individual and/or the group level)?

Previous work with agents that display their action tendencies—whether to continue an encounter or whether to abort it—has shown that taking other agents’ displayed action tendencies into account leads to better group outcomes (Scheutz and Schermerhorn, 2004). For example, if it is obvious that an opponent is very likely to continue to fight over the resource (i.e., that it has a high action tendency to fight) and ultimately win the encounter, then it is not in an agent’s best interest to enter the fight when it is less likely to continue to fight (i.e., it has a lower action tendency to fight) and win the encounter, thereby wasting resources fighting while gaining no benefit. Retreating immediately may also be costly, but compared to the cost of prolonged fighting, it is in the agent’s best interest to retreat. Furthermore, it is in the more aggressive agent’s best interest for its opponent to leave early, because prolonged fights reduce the net benefit of the resource being contested.

3. Conflict resolution in the competition for resources

In this study, we consider a particular kind of embedded two-player stopping game: the competition of individuals for food in the larger context of foraging for food in order for them to survive long enough to be able to

procreate. Each agent has an internal energy store, from which energy is subtracted for movements, fighting, and even just existing. In other words, there is a cost associated with each action which is measured in terms of energy, and ingesting food is the only way to replenish the internal energy store. During each agent’s search for food, there may be occasions in which two or more agents are trying to obtain the same food source. These encounters then lead to the embedded stopping games, where agents can decide to quit and leave the food for the other agent(s), or to fight for food. Costs (e.g., for fighting) are assessed at every step of the game. When participants run out of resources, they are removed from the (larger) game, and hence cannot have offspring. In the evolutionary context, this means that their genes are not passed on.

There is ample evidence that in nature the differences between opponents with respect to their capacity to inflict harm or impose cost on each other matters in making decision about whether to fight for a resource or whether to retreat (and leave the resource to the opponent) (e.g., see (Maynard-Smith, 1982, Parker, 1974)). In addition, the value of the resources plays a crucial role (a factor we will ignore in the experiments to simplify the agent model and the theoretical analysis). Hence, a very general model will need to include some measure of the kind of harm that could be done and the likelihood that it will be done as well as the measure for the value of the resources over which the conflict is carried out.

3.1 The “embedded” competition game

We will use the following abbreviations for constants that determine the benefits and costs involved in the competition for a resource: B_W denotes the benefit of winning an encounter (i.e., the expected utility of the resource), B_L denotes the benefit of losing an encounter (i.e., the expected utility of obtaining other resources), C_F denotes the cost of fighting to the individual, and C_R denotes the cost of retreating. All costs are measured in terms of energy. In general, we can assume that the utility of winning outweighs the utility of losing, i.e., $B_W + C_F > B_L + C_R$.

We model an agent’s decision making process as a random variable that assumes values in {FIGHT, RETREAT}. We denote the probability (in $[0, 1]$) that an agent will fight by P_S and that it will retreat by $1 - P_S$ (for opponents we will use the P_A and $1 - P_A$, respectively).

Each competition can then be described by a sequence of interactions (or “rounds”) which continue as long as both player decide to fight and end as soon as one player decides to retreat. The utility function $U(1)$ determining the expected outcome for a one round interaction is then given by:

$$U(1) = P_S \cdot P_A \cdot C_F + P_S \cdot (1 - P_A) \cdot (B_W + C_F) + (1 - P_S) \cdot P_A \cdot (B_L + C_R) + (1 - P_S) \cdot (1 - P_A) \cdot (B_L + C_R)$$

The above equation can be further transformed and simplified into

$$U(1) = P_S \cdot ((1 - P_A) \cdot B_W + C_F) + (1 - P_S) \cdot (B_L + C_R)$$

In general, the utility for an $(n + 1)$ -round game can be determined in terms of the utility of a 1-round game and the cost of the “fight-fight” outcome, denoted by C_F (note that for an $n + 1$ round-game, there must have been n preceding rounds with “fight-fight” outcomes):

$$U(n + 1) = n \cdot C_F + U(1)$$

The utility function $U(n)$ is strictly decreasing. Consequently, it is clear that long games only incur costs and that, therefore, the best game (other things being equal) is a one-round game. $U(1)$ can be maximized by player S depending on player A ’s choice in the following way: suppose P_S is fixed, then player A should play $P_A = 1$ if $(1 - P_S) \cdot B_W + C_F > B_L + C_R$.² Conversely, if $(1 - P_S) \cdot B_W + C_F < B_L + C_R$, then player A should play $P_A = 0$. For $(1 - P_S) \cdot B_W + C_F = B_L + C_R$ player A will be indifferent between $P_A = 0$ and $P_A = 1$ (any value in $[0, 1]$ will pay $B_L + C_R$). Similarly, player S will be indifferent if player A plays P_A such that $(1 - P_S) \cdot B_W + C_F = B_L + C_R$. Therefore, the pair (P_A, P_S) such that $(1 - P_S) \cdot B_W + C_F = (1 - P_A) \cdot B_W + C_F = B_L + C_R$ is a Nash equilibrium.³

It may be tempting for player S to play $P_S = 1$ as this promises the highest payoff. However, in the absence of information about P_A this can lead to prolonged games, e.g., if both S and A decide to play “fight” (i.e., $P_S = P_A = 1$), where the “payoff” per round is the only the cost C_F with no benefits to either contestant (and reduced expected utility for both players in the next round).

If information about the likely behavior of the other participant is available, however, then better decisions can be reached that will ensure short and for some strategies one-round games. Contestants can, for example, indicate their behavioral disposition in the upcoming conflict, i.e., their respective apriori probability to choose “fight” or “retreat” in the encounter. In nature, one such indicator displayed by animals in terms of facial expressions and gestures (and used in this investigation) is the individual’s level of aggression, which is causally linked to intensity and duration of a fight. Based on

²In general, if an expression α (e.g., $(1 - P_S) \cdot B_W + C_F$) is greater than an expression β (e.g., $B_L + C_R$), then $x \cdot \alpha + (1 - x) \cdot \beta$ has a maximum for $x = 1$ for all $x \in [0, 1]$.

³Whether there are other equilibria will depend on the particular choice of the involved parameters.

these apriori probabilities, contestants can gauge their opponents’ willingness to fight and figure it into their decisions.

For example, an agent might reduce its propensity to fight if it notices that the opponents’ aggression level, and thus likelihood to engage in long encounters, is higher than the agent’s own aggression level. Such an adjustment pays off, because the likelihood of long encounters is reduced and the more likely outcome—that the agent with the higher aggression level will win—is reinforced.

A one-shot rule to end games after the first round based on such encounters is what we called the “rational rule” (Schermerhorn and Scheutz, 2003, Scheutz and Schermerhorn, 2004): agent S plays 0 if $P_S \leq P_A$ and 1 otherwise. This rule is based on the assumption that contestants do not know the actual value of B_W, B_L, C_F , and C_R , hence they cannot compute whether $(1 - P_A) \cdot B_W + C_F > B_L + C_R$. Furthermore, the rule is intended to guarantee one-round games, which are not guaranteed if any of the Nash equilibria are played: suppose that both agents indicate probabilities p such that $(1 - p) \cdot B_W + C_F$, then they both should play “fight”, which will lock them into a prolonged game with only costs. The rational rule, on the other hand, prevents this by forcing the contestant with the lower probability to “give up” right away and still get at least $B_L + C_R$ (instead of the winner’s $(1 - p) \cdot B_W + C_F$).

Of course, this rule not is *fair* in that repeated encounters between the same two individuals will lead to the same outcome, i.e., the same individual will win over and over again. In a group that means that the top individual (with the highest level of aggression) will have $n \cdot ((1 - p) \cdot B_W + C_F)$ payoff after n encounters and the bottom one (with the lowest level of aggression) will have $n \cdot (B_L + C_R)$, the others being in-between depending on their encounters and aggression levels. From an evolutionary perspective, this will lead to an “arms race” of aggression levels (as the more aggressive is bound to win), which eventually will lead to agents that always fight and only “give up” when they die (Schermerhorn and Scheutz, 2003). These extended conflicts are very costly, and as a consequence a population of such agents will dwindle or become extinct altogether.

A fair way to distribute resources equally among members in a group is to take turns (e.g., (Neill, 2003)). In a two-agent environment, players would alternate between getting $B_W + C_F$ and $B_L + C_R$, thus getting the average payoff $(B_W + C_F + B_L + C_R)/2$ every turn. For environments with more agents, turn-taking is not as straightforward as it is possible that both agents in a conflict took their turn at losing in the previous encounter, so both will expect to win in the current encounter. There are a number of ways to implement a fair turn-taking rule

for multi-agent environments. An agent could, for example, keep track of the other agents it interacted with and figure out whose turn it is to win next based on the interaction history. Such an approach, however, is not only computationally expensive, but also requires the ability to track other agents and their encounters over time (which, for example, presumes mechanisms for reidentify agents). We have proposed a simple tallying mechanism that keeps track of the number of encounters the agent won or lost by increasing or decreasing its aggression base aggression level. This adaptation process leads to higher aggression levels if the agent has not been able to obtain the desired resources and to lower aggression levels if the desired resource has been obtained. Formally, the rule can be defined as follows:

Definition [The Turn-Taking Rule] Let r be the basic aggression level of an agent A and let m be the current aggression level ($r = m$ if the rule has never been applied). Then the turn-taking rule $TT(m)^+$ is defined (for losses) as follows: if $m \geq r$, then $TT^+(m) = m + (1 - m)/2$; if $m \leq r/2$, then $TT^+(m) = 2 * m$; else $TT^+(m) = r + (2m - r)(1 - r)/2r$ (this maps values in the interval $(r/2, r)$ into $(r, (1 - r)/2)$). Similarly, $TT(m)^-$ is defined (for wins) as follows: if $m \geq r + (1 - r)/2$, then $TT^-(m) = m - (1 - m)$; if $m \leq r$, then $TT^-(m) = m/2$; else $TT^-(m) = r/2 + r(m - r)/(1 - r)$ (this maps $(r, (1 - r)/2)$ into $(r/2, r)$).

This rule effectively keeps track of how often an individual was able to obtain a resource by increasing or decreasing the level of aggression relative to the rest level and the current level of aggression in a way that the computation is straightforward and inexpensive (in fact, the rule can easily be implemented by a few recurrent perceptrons).

We have shown (Schermerhorn and Scheutz, 2003) that this rule, if combined with the “rational rule” described above, is *fair* in a clearly specified sense: the absolute difference between the number of wins and losses of agents in a group in conflicts is bound by a constant c for all agents: $|wins - losses| < c$, where c depends on the size of the group. This means that at the worst an individual member of the group will have c more losses than wins.⁴

4. Strategies and Agent Models

The above discussion suggests that signaling action tendencies through displays of aggression is beneficial for the survivability of individuals in an embedded game. In particular, making “rational decisions” while taking

⁴It is important to note that this is a theoretical limit. Our previous simulations show that in practice, this number is almost never reached (especially not for large populations), hence the group practically is even more balanced than one would expect from the upper bound c .

turns using the above suggested mechanism to adapt aggression levels has proved superior to all other strategies in previous simulation studies with agents that do not discriminate between “own” or “other”. The question to be answered here is whether this is also true of agents that are capable of discrimination and if not, what the best strategy would be for such agents. We will in the following first define seven different kinds of agents and then report the performance ordering of new agents created from combinations of two different out of the seven strategies. These agents apply one strategy to agents from their own kind and one to agents from other kinds, respectively.

4.1 7 Non-Discriminating Agent Types

We distinguish two different classes of agents, call them *social* and *asocial* agents, depending on whether or not they use perceptions of opponents’ aggression in their decision to fight or retreat. All asocial agents decide whether they will enter and stay in a conflict situation based only on their own aggression level in a probabilistic fashion: the more aggressive they are, the more likely they are to stay. In the experiments reported here the mapping between aggression level and action tendency is the identity function, but other functions are possible. Social agents also decide their actions in the same probabilistic fashion but include perceptions of other agents’ aggression levels, possibly to a point where they reduce their own decisions to non-probabilistic certain actions (i.e., either “fight” or “retreat”).

While we assume that all agents display their aggression for the following experiments, it is worth mentioning that the display for asocial agents has no added benefit as they do not use perceived aggression levels in their decision making. Hence, the strategies for asocial agents would work exactly the same for non-signaling agents.⁵

In the class of asocial agents, we distinguish three kinds: the (prototypical) *asocial* agents, which play their aggression level (i.e., a value in $(0, 1)$), *timid agents*, which always play 0, and *aggressive agents*, which always play 1.⁶

In the class of social agents, we distinguish two kinds: the (prototypical) *social* agents, which increase or lower their action tendencies based on both their own and the others’ aggression levels, and the *rational* agents, which implement the “rational rule” from above. Prototypical

⁵It is also possible to augment the strategy of social agents so as to work with non-signaling agents by simply making assumptions about their level of aggression. E.g., in the case of the rational strategy described above an agent could start by ascribing no aggression to the other agent and then systematically increasing the ascribed aggression level such that it will leave a conflict in time before the accumulated cost of fighting outweighs the potential benefit of winning.

⁶Note that we use “asocial agent” to denote a class of agents as well as an agent kind for ease of terminology as the intended interpretation will always be clear from the context.

social agents use the following rule: if $P_A > P_S$, then they play $P_S - (P_A/P_S) \cdot (P_A - P_S)$, if $P_A < P_S$, they play $P_S + (1 - P_S/P_A) \cdot (P_S - P_A)$, and if both levels are equal they simply play P_S .⁷

Finally, we consider a third category, that of *adaptive agents*, which implement the “turn-taking rule” and thus change their aggression levels based on past encounters (contrary to the above described social and asocial agents, which modify only their behavioral disposition based on their perceptions, but keep their aggression levels constant throughout their lifetime). In particular, we consider (prototypical) *asocial adaptive* agents and *rational adaptive* agents.

4.2 42 Discriminating Agents Types

Based on the above non-discriminating agent types, we can define 42 different discriminating agent types by endowing agents with the ability to distinguish perceptually between “own” and “other” and applying different strategies depending on the membership of the opponent.⁸ We will denote these agents by the name of the strategy used against others followed by the name used against their own (e.g., an “aggressive asocial agent” is aggressive towards others and asocial toward its own without any modification of aggression, although the action tendencies towards other is to always “fight”).

Of the 42 discriminating agents we will focus on the six asocial discriminating agents, i.e., *the timid*, *the aggressive*, *the social*, *the the rational*, *the the adaptive asocial*, and *the the adaptive rational* asocial agents. The choice for the class of asocial agents is based on the previous studies that show that asocial agents are worse than social agents (both when placed in direct competition with them and when among their own). The question is whether discriminating among “own” and “other” can boost the performance of this agent kind above the performance of the others. Given that we compare 6 discriminating agents to 7 non-discriminating ones, we will consider 42 pairings of discriminating and non-discriminating agents altogether.

5. Simulation Experiments

All simulations described in this section were carried out in the simulated world agent grid experimentation system SWAGES under development in our lab.⁹

⁷For a more detailed discussion of this rule, see Scheutz and Schermerhorn, forthcoming. We are currently exploring the relationship between a similar rule used for the analysis of dominance interactions by (Hemelrijk, 1996).

⁸We exclude those 7 discriminating kinds that apply the same strategy to their own kinds and to other kinds, which effectively amounts to no discrimination and thus reduces them to the respective non-discriminating agents.

⁹SWAGES is an agent-based artificial life simulation built on top of the SIMAGENT toolkit, both of which are freely available at <http://www.nd.edu/~airolab/swages> and

For the simulation reported here, we configured the simulation environment for a “one-resource foraging task”. Instances of the resource, call them “food”, contain 800 energy units, pop up at random locations within a predetermined rectangular 1800 x 1800 food area—agents have a size of 9 in comparison—at a predefined frequency (1 resource per cycle) and stay indefinitely until consumed by an agent.

Agents are in constant need of energy and, hence, constantly engaged in foraging for resources. They have a visual system that computes directional force vectors that are scaled by the square of the distance to each resource within their sensory range (of 300 distance units). The vectors are summed and the agent moves in direction of the resultant vector at a speed of at most 4 units per cycle (for details of the architecture, see (Scheutz, 2001, Scheutz and Schermerhorn, 2002)). The energy expenditure for movement is the square of their speed. In addition, each agent consumes one unit of energy per cycle for processing.

Agents also have touch sensors, which they use to detect whether they can ingest a resource or whether they are about to collide with another agent, which means that they enter a competition (typically, this happens when food is close by, but does not have to involve a food source). In the former case, an ingestion mechanism will be triggered to ingest the food item (which takes one cycle), while in the latter a self-preservation collision detection mechanisms will be triggered unless suppressed by the agent. This mechanism is coupled with a reflex that will the agent from crashing into other agents (which would be fatal) by moving them away from the potential collision site at a speed of 7. This reflex is also triggered when agents lose a conflict, hence its duration determines C_R , which on average is 330 energy units (simply because of movement at a higher than usual speed). $C_F = 50$, in contrast, is charged to agents that suppress this mechanism and fight for the resource.

Agents also have an “energy alarm” for self-preservation, which automatically limits their overall speed to 1 if their energy level drops below $Energy_{crit} = 400$ (the speed will remain at 1 until the energy level is raised above the critical level again).

After a certain age α (measured in terms of simulation cycles), agents reach maturity and can procreate asexually, if their energy levels are above the minimum necessary for procreation (set to 2200). The energy necessary for creating the offspring (2000) is subtracted from the parent, and a new agent will appear in the vicinity of the parent in the subsequent simulation cycle. Parents pass all the parameters of all of their strategies to their

<http://www.cs.bham.ac.uk/research/poplog/newkit.tar.gz>. It consists of a continuous, potentially unlimited two-dimensional surface populated with various kinds of spatially extended objects, in particular, different kinds of agents and resources they need as well as various kinds of static and moving obstacles.




Figure 1: The average number of survivors in the experiment sets comparing timid-asocial discriminating agents to all seven non-discriminating agent kinds.




Figure 2: The average number of survivors in the experiment sets comparing social-asocial discriminating agents to all seven non-discriminating agent kinds.

offspring without modification (e.g., mutation).

5.1 Experiments and Results

We conducted 42 sets of experiments, each consisting of 40 different runs for 10000 cycles each with different random initial conditions (the same 40 different initial conditions were used in all 42 sets to guarantee a fair comparison). In all runs, 25 discriminating and 25 non-discriminating agents were placed at random locations within the 1800 x 1800 food area in the environment together with 50 randomly placed food items. The initial energy of all agents was set to 2000 and their initial aggression levels were distributed following a Gaussian distribution with spread 0.125 around 0.5.

Previous results from simulations of multi-species environments (Scheutz and Schermerhorn, 2003, 2004) with non-discriminating agents show that it is better to be “timid” than to be “aggressive”, and better to be “social” than “asocial”, and even better to be “rational” or to take turns on top of being rational. In short, merely aggressive agents showed the worst performance,




Figure 3: The average number of survivors in the experiment sets comparing aggressive-asocial discriminating agents to all seven non-discriminating agent kinds.




Figure 4: The average number of survivors in the experiment sets comparing rational-asocial discriminating agents to all seven non-discriminating agent kinds.

and all asocial agents performed worse than all social or adaptive agents.

Adding discrimination, however, changes the picture and we see that applying two standards (in this case, strategies for interactions), one for “own” and a different one for “other”, pays off, even if the two employed strategies used in isolation are not very good at all. Figures 1 through 6 show the average number of survivors at the end of the simulation run for each discriminating agent kind compared to those of all non-discriminating kinds (the error bars depict the 95% confidence intervals).

The results show that asocial agents with the highest level of aggression against others (i.e., an “always fight” strategy) outperform all other combinations of asocial discriminating agents in addition to outperforming all non-discriminating agents (see Figure 3 compared to the others). In fact, there is no other combination that always leads to the extinction of all non-discriminating agent kinds within only a few generations.

Discriminating timid asocial agents only perform well against other timid and aggressive agents, where they

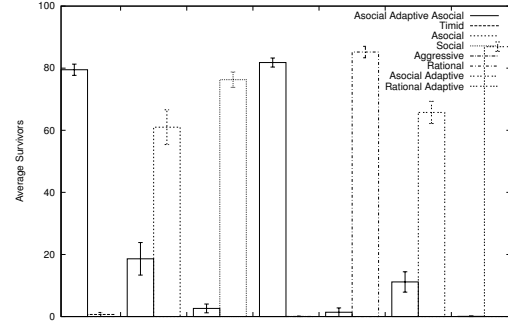


Figure 5: The average number of survivors in the experiment sets comparing asocial adaptive-asocial discriminating agents to all seven non-discriminating agent kinds.

can reap the benefits from their intragroup conflicts. They lose against all other kinds, similar to the timid non-discriminating agents.

Social asocial discriminating agents perform better than timid and aggressive, on a par with prototypical asocial and social agents, but worse than non-discriminating rational agents or adaptive agents (i.e., the benefit of adaptation is still higher than that of discrimination). Note that they manage to survive against the non-discriminating rational adaptive agents, which is possible because they will occasionally win encounters against them (when their adaptation has lowered their aggression level to a point where multiple losses in a row are possible, which significantly increases the likelihood for the social strategy to win).

The discriminating rational asocial agents show an expected high performance against asocial and prototypical social agents given that they play the best social strategy against others (the second best performance after the discriminating aggressive asocial agents). Compared to non-discriminating rational agents, however, their intraspecies asocial conflicts reduce their overall performance.

Finally, the discriminating adaptive agents suffer from the same interspecies conflicts that eventually reduce their fitness. Adaptation caused by conflicts with others also seems to work against them (compared to the discriminating rational asocial agents, for example). Especially, if the majority of opponents are “others”, discriminating adaptive asocial agents will have raised their aggression to very high levels, which in turn are likely to lead to harmful “fight-fight” encounters in intragroup competitions.

5.2 Analysis

The results of the simulation studies may seem somewhat surprising at first glance. However, we believe that they should not come completely unexpected at

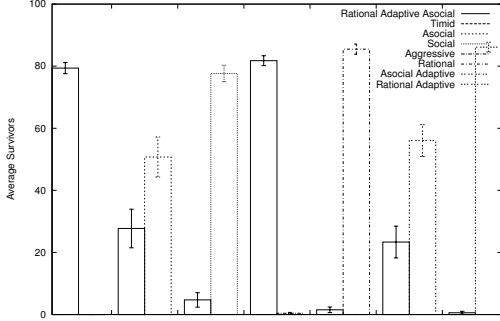


Figure 6: The average number of survivors in the experiment sets comparing rational adaptive-asocial discriminating agents to all seven non-discriminating agent kinds.

second thought, if the goals of discriminating vs. non-discriminating (i.e., intraspecies vs. interspecies) competitions are taken to be different. Roughly, the goal in non-discriminating competitions is to maximize the individual’s payoff, while in discriminating competitions, as we would like to suggest, it is to minimize the relative payoff of all competing members in other groups compared to competing members in the own group. In other words, a strategy for intergroup competition is good if it inflicts more damage on members of the other group than on members of the own group, while a strategy for intragroup competitions is good if it gives more benefit to the individual than to other contestants. It is important to note that these two conditions are not symmetric. The discriminating rational asocial agents, for example, attempt to maximize their own utility while playing against others, rather than harming the others as in the case of the discriminating aggressive asocial agents. Consequently, despite their generally high performance against any other non-discriminating agent types (see Figure 4), they do not even come close to the performance of the discriminating aggressive asocial agents. We will illustrate this asymmetry with a brief example.

Consider an environment with two non-discriminating agents of one kind (call it A) and two discriminating agents of another kind (call it B). There are 6 possible mixed two-way interactions of these four agents, i.e., where at least one A -agent interacts with at least one B agent: $a1 - b1$, $a2 - b2$, $a1 - b2$, $a2 - b1$, where the remaining two agents do not interact, and $a1 - b1/a2 - b2$, and $a1 - b2/a2 - b1$. Suppose the B agents are aggressive asocial agents with aggression level 0.3 and 0.5, respectively, and the A -agents are rational agents with aggression level 0.6 and 0.75, respectively. The total expected utility for A summed over the six encounters is $6 \cdot (C_R + B_L)$, while it is $6 \cdot (C_F + B_W)$ for the B agents. Now we consider the intraspecies competition, where the A -agents get a combined utility $C_F + B_W + C_R + B_L$,

while the B -agents will at most get that much, given that with probability 0.15 their interaction is prolonged. If we consider only one-round encounters for simplicity sake, then B -agents get $0.35(C_F + B_W) + 0.15 \cdot (C_F + B_W) + 0.35 \cdot (C_R + B_L) + 0.15 \cdot C_F$. Combining these two utilities we can see that the discriminating agents do much better on average than the non-discriminating agents given that the discriminating agents play the harshest possible strategy against “others”. We can also see that it is important for the discriminating agents to be confrontational, as playing a non-confrontational strategy like “always retreat” would not have worked (rather it would have resulted in higher average payoffs for the non-discriminating agents, as exemplified in the results shown in Figure 1).

6. Discussion

One way to understand why the difference between discriminating and non-discriminating agents gives rise to such a significant increase in performance is to appreciate that discriminating agents have the option to choose how to distribute resources to non-group members, whereas non-discriminating agents do not have this option (a change in strategy affects non-members and members alike). The simulation studies confirm our analysis that once this choice becomes an option, then the best strategy is not one that favors others, but rather represses them. It is almost ironic that discriminating aggressive agents do better when other agents are around, as they gain more from others’ losses (because of their aggression) than they gain from interactions among themselves. Yet, their strategy is so “successful” that they quickly cause the extinction of non-discriminating agents (in all the above simulation the total number of competitors of discriminating aggressive asocial agents never exceeded 200, while their numbers were in the thousands).

Given that the worst strategy, being aggressive and being inclined to fight all the time, for non-discriminating agents turns out to be the best for discriminating ones (as far as “others” are concerned), one would expect aggression to evolve if discrimination is possible, and one would expect discrimination to evolve, if (perceptual) categorization is possible (as discrimination, the way we modeled it here, is based on perceptual difference, which requires a categorization mechanism that is capable of distinguishing between “own” and “other”).

The outcomes of these investigations may, thus, help answer the question why there are aggression mechanisms in biological organisms in the first place, when it seems that aggression is not beneficial in the non-discriminating case. One answer given by Lorenz, for example, is that intra-species aggression serves to force animals to distribute better over a given territory (as it prevents animals from sticking together, thus having to compete for scarce resources in a small area, while

resources are plentiful elsewhere (Lorenz, 1963)).

A different answer based on our results, which would, of course, require further investigation, might be this: suppose species developed an aggression mechanism because it was beneficial to be aggressive towards others in the competition for resources, thus furthering (or maybe even ensuring) the survival of their kind. Once such a mechanism is in place, there are two possibilities for its activation: either through innate categories of “other kinds” or through an adaptation phase, where the difference between own and other is learned and an association between other and the aggression mechanism is formed. In that case, the association between own and other is not fixed as in the innate case, but amenable to change and thus subject to modifications during an individual’s lifetime. While this adds flexibility and thus higher adaptivity to individuals (as they may be able to survive in environments with other competitors than their current ones as well), it also opens up the possibility of forming associations between “own” and the aggression mechanism. It is then possible, that in the absence of other species, the discrimination is applied to the own kind and a new division between other and own is formed such that the aggression mechanism is associated with new category of “other”, which really is a subgroup of “own”.

Of course, this consideration assumes that there is a (possibly innate) mechanism that is disposed to divide individuals in the environment into own and other. The categorization could be based on many factors ranging from difference in physical appearance to mental attitudes. The important point here is that there is great evolutionary utility to having such a mechanism that will elicit aggressive behavior towards others, i.e., a mechanism that causes the individual to look out for differences among other individuals, categorize other individuals in terms of these differences into “own” and “other”, and elicit aggressive behavior towards individuals that are categorized not “own”, but “other”. In that case, the problem would not be with the aggression mechanism *per se*, but rather with the lack of target individuals against which the aggression is to be employed.¹⁰

Note that there are good evolutionary reasons to make this mechanism innate, i.e., to prevent individuals from failing to develop such a categorization, as the failure to develop such a discrimination might not serve the group as a whole, and consequently (future) individuals. Specifically, it may give rise to “free-riders” benefit from their conspecifics’ aggression towards others (e.g., an example in the context of a war between nations might be soldiers that do not put their lives at risk in combat.¹¹

¹⁰It seems that this consideration is consistent with what it commonly believed about the utility or role of sports in modern societies.

¹¹It seems that main reason why people fight is to not let their “buddies” down, see (Little, 1964).

Finally, we would like to mention that the above proposed agent architectures fits within the “frustration-aggression theory of “aggression” (e.g., see (Berkowitz, 1989)). Specifically, the turn-taking mechanism can be taken to feed into the frustration level, so that whenever a resource is not obtained frustration grows according to the turn-taking rule, which then feeds into aggression, which is subsequently displayed. In this setup, agents get more aggressive as their frustration increases (which increase because they do not obtain their resources based on the decision mechanism built into the turn-taking rule). Eventually they will get the desired resource and aggression levels will subsequently drop (as frustration levels drop caused by goal accomplishments).

7. Conclusion

In this paper, we have addressed the question whether aggression is a useful control mechanism in the competition of agents for resources. Specifically, we demonstrated that aggression, while not beneficial by itself, is highly beneficial in combination with discrimination between groups (i.e., between agents from one’s own group and all other agents). Extensive simulation experiments showed that other combinations of strategies for the seven basic agent types considered in the paper do not fare as well. This should be an encouraging result for those who would like to ground a theory of aggression in evolutionary pressures. Although we have only been able to sketch a possible account, it should be clear how our conjecture could, at least in principle, be tested in further, much more elaborate simulation experiments with agents that have the respective architectural capacities (such as categorization, learning mechanisms, etc.).

The above results also suggest another interesting direction for future research, which is feasible in the short term and which we intend to pursue. While in principle discriminating agents derive their increase in performance from the harm they inflict on others, it is not clear to what extent this actually beneficial for them and how, for example, the above investigated aggressive asocial agents would fare against other discriminating agents (that also played an aggressive strategy against “others”). In particular, it would be interesting to determine the relative productive success of the discriminating agents that had more encounters with agents from another kind (i.e., which mostly played the aggressive strategy) compared to their conspecifics, which mostly interacted with their own kind. This should shed light on the dynamics of the interactions between agents of different kinds and help answer questions about why some biological species seem to be intrinsically more aggressive than others.

Finally, we would like to point out that the results are likely to have implications beyond the biological domain, as the general game-theoretic principles do not

only apply to the competition for resources that are immediately relevant and a constitutive part of the embedded survival game. Rather, they are likely to apply to many social situations and groups of different social structure. If this is correct, then more detailed studies of the kinds of embedded games we have considered in this paper might contribute to our understanding of the implications of aggression to the social dynamics as they manifest themselves in national and international conflicts.

References

- Adamo, S. A. and Hanlon, R. T. (1996). Do cuttlefish (cephalopoda) signal their intentions to conspecifics during agonistic encounters? *Animal Behavior*, (52):73–81.
- Albert, D., Walsh, M., and Jonik, R. (1993). Aggression in humans: what is its biological foundations? *Neurosci Behav Rev*, (17):405–425.
- Axelrod, R. and Hamilton, W. D. (1981). The evolution of cooperation. *Science*, 211:1390–1396.
- Barlow, G., Rogers, W., and Fraley, N. (1986). Do midas cichlids win through prowess or daring? it depends. *Behav Ecol Sociobiol*, (19):1–8.
- Berkowitz, L. (1989). Frustration-aggression hypothesis: examination and reformulation. *Psychological Bulletin*, (106):59–73.
- Brain, P. (1979). Classical ethology and human aggression. In Brain, P. and Benton, D., (Eds.), *The biology of aggression*, pages 603–623.
- DeCourcy, K. R. and Jenssen, T. A. (1994). Structure and use of male territorial headbob signals by the lizard *Anolis carolinensis*. *Animal Behavior*, (47):251–262.
- Eibl-Eibesfeld, I. (1974). Aggression in the ikobushmen. *Aggression*, (52):1–20.
- Enquist, M. and Leimar, O. (1983). Evolution of fighting behavior: decision rules and assessment of relative strength. *Journal of Theoretical Biology*, (102):387–410.
- Hemelrijk, C. K. (1996). Dominance interactions, spatial dynamics, and emergent reciprocity in a virtual world. In *Proceedings of the Fourth International Conference on the Simulation of Adaptive Behavior*, pages 545–552.
- Hofmann, H. A. and Schildberger, K. (2001). Assessment of strength and willingness to fight during aggressive encounters in crickets. *Animal Behavior*, (62):337–348.
- Little, R. (1964). Buddy relations and combat performance. In Janowitz, M., (Ed.), *The new military*.
- Lorenz, K. (1963). Das sogenannte bse. *Naturgeschichte der Aggression*.
- Lorenz, K. (1977). Aggressivitt: arterhaltende eigenschaft oder pathologische erscheinung? *Aggression und Toleranz*.
- Maynard Smith, J. (1974). The theory of games and the evolution of animal conflicts. *Journal of Theoretical Biology*, pages 209–221.
- Maynard-Smith, J. (1982). *Evolution and the Theory of Games*. Cambridge University Press, Cambridge.
- Neill, D. (2003). Cooperation and coordination in the turn-taking dilemma. In *14th International Conference on Game Theory at Stony Brook*.
- Parker, G. (1974). Assessment strategy and the evolution of fighting behaviour. *Journal of Theoretical Biology*, pages 223–243.
- Schermerhorn, P. and Scheutz, M. (2003). Implicit cooperation in conflict resolution for simple agents. In *Agent 2003*.
- Scheutz, M. (2001). The evolution of simple affective states in multi-agent environments. In Cañamero, D., (Ed.), *Proceedings of AAAI Fall Symposium*, pages 123–128, Falmouth, MA. AAAI Press.
- Scheutz, M. and Schermerhorn, P. (2002). Steps towards a theory of possible trajectories from reactive to deliberative control systems. In Standish, R., (Ed.), *Proceedings of the 8th Conference of Artificial Life*. MIT Press.
- Scheutz, M. and Schermerhorn, P. (2004). The role of signaling action tendencies in conflict resolution. *Journal of Artificial Societies and Social Simulation*, 1(7).
- Shmaya, E., Solan, E., and Vieille, N. (2003). An application of ramsey theorem to stopping games. *Games and Economic Behavior*, 2(42):300–306.
- Touzi, N. and Vieille, N. (2002). Continuous-time dynkin games with mixed strategies. *SIAM J. Control Optim*, 4(41):1073–1088.