Leadership in movement-aligning groups

MSc Research Project
(Specialisation Autonomous Systems)

Daniel A. P. Reid
s1083589

April 4, 2005

Internal supervisor:
Prof. dr. L. Schomaker
External supervisor:
Dr. C.K. Hemelrijk
Faculty of Biology, department of Theoretical Biology
Rijksuniversiteit Groningen

Artificial Intelligence
Rijksuniversiteit Groningen

Contents

Abstract

Many models have shown that schooling behaviour can arise through self-organisation, using individual-based models. In this study, I adapt a model in which agents engage in dominance interactions with winner-loser effects, resulting in a dominance hierarchy. I add schooling and goal-directed behaviour to the model. The effects of different gradients of hierarchy on group structure are examined, in schools both containing agents with a spatial target(*goal-directed*) and schools without goal-directed members. The gradient of hierarchy is found to have no influence on group structure of movement-aligning groups in this model. Goal-directed agents are found to strongly move towards the front of the group over time. This process is affected by both gradient of hierarchy and whether the goal-directed agents are dominant or subordinate: if the gradient is steep, dominant goal-directed agents reach the front of the group sooner than in groups with a shallow gradient; the opposite holds when subordinate agents are goal-directed.

1 Introduction

1.1 Self-organisation and emergence

In the fields of both theoretical biology and artificial intelligence, self-organisation is an important topic. It can be defined in short as order and complexity arising from the interactions between a system's components, without there being a central controller or organiser. The process whereby this order arises is self-organisation, also called *order for free* [10]. Although it is a somewhat vague concept that lends itself easily to long discussions about exact definitions, most people instinctively have an idea of what is meant. The most obvious and most complicated example of a self-organising system is life itself. Life creates a multitude of patterns, and vastly increases global complexity and order. It arises, however, from very simple interactions at a molecular level

One of the most elegant definitions of emergence I have come across is one given by Luc Steels [22]: *"Emergent functionality means that a function is not achieved directly by a component or a hierarchical system of components, but indirectly by the interaction of more primitive components among themselves and with the world."* This nicely emphasises the main point about emergent phenomena, which is that they are observable at the macroscopic level, even though they are generated by micro-level elements. A common misconception about emergence is that the macro-level behaviour must not be immediately obvious from the micro-level. This is not, in my view, a sensible criterium, since it means a behaviour would then be classified as emergent dependent on the analytical properties of the observer instead of the behaviour itself. A nice example of emergence is schooling behaviour in simulations of fish. The macro-level, emergent behaviour, the school, arises from very simple micro-level rules based on the agents' immediate neighbours.

Emergence and self-organisation are closely-related concepts. In general, self-organisation is a more specific term than emergence – using the definition given above, emergent functionality is *any* functionality at macro-level that arises through micro-level interactions. This can be best illustrated by the following example: An explosion is a macro-level emergent function of the behaviour of the molecules at the micro-level, yet it is not a self-organising system, since it *decreases* the system's organisation.

1.2 Self-organisation and A.I.

When the field of artificial intelligence first came into being, the focus was on making large, complex systems, that would eventually be just as intelligent, if not more, than humans. This approach is also known as Classic A.I. and involves elaborate internal representations of the world, large deduction systems and so forth. However, another approach has arisen since then that focuses on aggregate systems consisting of many simple components instead of a single complex one. This alternative paradigm is known as behaviour-based A.I. and focuses heavily on self-organisational, distributed approaches to problems. Examples include projects using genetic algorithms, cellular automata, agent-based models and some recent research using groups of relatively simple robotic agents [15, 24]. These approaches have the advantage of being generally robust in that they operate well under noisy conditions and that they give insight into how new phenomena and functionality may emerge [23].

1.3 Self-organisation and biology

There are many aspects of biology that can be modeled using self-organisation, ranging from single cells to human populations. The following are interesting examples that show the use of theoretical modeling in exploring animal behaviour.

A three-dimensional cellular automaton was used to model the complete life cycle of the slime mold (*Dictyostelium discoideum*) from its behaviour as single cells to the eventual aggregation into moving slugs and then fruiting bodies [14]. This is interesting in that a single model, using only local rules just as the amoeba presumably do, accounts for many seemingly disparate behaviours. All the amoebae exude a chemical called cyclic AMP, and have a refractory period so that wave patterns are generated as in an excitable medium. Combined with chemotaxis towards this chemical, the cells accomplish aggregation, slug movement, morphogenesis and fruiting body formation.

Another model was used to investigate the trail-laying and trail-following behaviour of foraging army ants [21]. The model was very similar to the one used by Deneubourg et al. [4] – a two-dimensional grid, with agents limited in their senses to the squares directly to their front. The net energy gain of the colony was used as a fitness measure, and a small set of parameters was varied between simulations to maximise fitness. The various species of army ant vary in their diet, ranging from large but rare to common but small food sources. Changing the food distribution in the model produced foraging patterns that resembled those found in real army ants with similar diets.

In both these cases a mechanism is implemented in a model that yields results that closely resemble a behaviour found in reality. While it is by no means certain that reality's underpinning mechanisms are the same, it is useful to demonstrate that behaviour can arise without complex individuals or intentionality. The models can be used to make predictions about behaviour in reality which can then be tested.

1.4 Self-organisation and movement alignment

One of the behaviours in biology that can be modeled well using a self-organisational approach is swarming or flocking behaviour, where a group of entities moves *en masse* in a single direction. This behaviour is generally believed to arise from a set of simple rules for each agent, where they align their movement direction with their close neighbours, avoid the ones that get too close and approach others if they are far away. In this study this behaviour is called *movement alignment* so as to distinguish it from coordinating groups where movement is not aligned. The distinction is easily illustrated using the two general terms for groups of fish: A shoal is a group of fish that stays together yet does not move as a whole; a school is a group of fish that is moving in a single direction. A shoal is still coordinating in that its members try to stay near others of their kind.

Shortly defined, therefore, a group is movement-aligning if its members are generally aligned in a single direction and the group is thus moving not much slower than the individuals' speeds. Good examples of this type of coordination behaviour are fish schools, bird flocks or primate progressions.

Only relatively recently has this behaviour been viewed as arising from simple, local rules [5]. Instead, the prevailing view was that there must be a few, experienced group members that somehow led the group and determined its behaviour and direction. Using computer simulations to mirror and examine this kind of grouping behaviour can yield valuable insights in the mechanics of schooling on an abstract level, and can inspire more biological research to test the results that are found in the computer simulations. It is this feedback loop between real-life research and simulations that makes this kind of research worthwhile.

Several models have been made that display movement-alignment behaviour that mirrors real-life flocking and schooling. One of the earliest was the BOIDS model by Reynolds [18], who built a fairly simple model based on agents' averaging of their responses to the group members around them, with the responses being split into the now standard repulsion, attraction and alignment behaviours. Though the model was not intended to be an explanation of animal behaviour, but was instead meant to create visually pleasing flocks, it has served as the inspiration for many more scientific and biologically grounded models. Since then, several models have been made, each studying a slightly different aspect of movement alignment [9, 13, 17]. Most movement alignment models are very limited in the *kind* of interactions the agents have – the only proximity-triggered interactions agents have is simple avoidance, and factors like body shape and size have only recently entered the picture. [9, 13].

One of the types of interaction missing in current implementations is the dominance interaction, even though many of the schooling or flocking species modelled display dominance interactions and hierarchies. Examples for fish include yellowtail [20], sticklebacks[1] and scalloped hammerhead shark [11]. Since dominance interactions have an important influence on group distribution in stationary groups [7], it would be interesting to look at this effect on moving, movement-aligning groups, and this is one of the aims of this study.

1.5 Leadership

Of late, there is much interest in how leadership works in a group context. Bumann and Krause [2] have observed that in several species of fish, group direction is heavily dependent on the direction of the frontmost individual. This can be contrasted with baboon groups, for example, where the frontmost individual is merely an *initiator*, and the direction of the group is ultimately supposed to be determined by a leader who is further towards the group's rear [3]. Krause et al. [12] also showed that in juvenile roach, hungry fish were to be found in the front of the group significantly more than satiated ones. It follows, then, from the importance to group direction of individuals in the group front, that these hungry individuals will also be the ones mainly determining where the group goes. If it is taken as a given that moving towards food is a positive influence on group survival, and *if* these hungry individuals know where food is to be found, they will then lead the group to this food, and thus increase the group's food intake and odds of survival.

Reebs [16] found that even when there was only one entrained individual in a group of 12 fish (golden shiners, *Notemigonus crysoleucas*), in several cases it successfully led the other 11 to a food source they had no previous knowledge of. If there were more entrained group members the success rate rose to 100%. The entrained fish were always the first ones to arrive at the source, hence they were in the front of the group. The study shows two interesting results: a small percentage of group members can have a large effect on group direction, and individuals with a directional bias are found at the front of the group. An interesting question therefore is whether the leading fish *wished to be* at the front of the group in order to lead it in a certain direction, or whether the frontal position it attained was instead a side-effect of its directional preference. This is one of the questions I will be investigating in this study.

In a related study using a computer simulation, Romey [19] also found that small numbers of individuals with different behavior within a group can have a significant effect on the group's behavior. This was also the conclusion of an investigation by Huse et al. [8] in which they examined the effects on group behaviour of individuals with a spatial target. In a simple movement alignment model, they studied how often the rest of the group followed a group of individuals with a directional preference.

So, in short, the aim of this study is two-fold: to study firstly the effects of dominance interactions and a dominance hierarchy on group structure in movement-aligning groups, and secondly whether goal-directed individuals move to the front of movement-aligning groups automatically, and whether it matters if these goal-directed individuals are dominants or subordinates.

2 Research Questions

To reach these goals, the following questions must be answered:

Given a movement-aligning group,

i What are the effects of dominance interactions and a dominance hierarchy on group structure, and how does the gradient of the dominance hierarchy affect this?

ii What are the consequences of having source-directed individuals for group structure? More specifically, will the source-directed individuals end up at the front of the group?

iii What is the difference between groups where the goal-directed agents are the dominants or the subordinates?

This study is mainly exploratory in purpose. Some of the assumptions and predictions made for the model are based on biological data, mainly Reebs', Bumann & Krause's and Huse et al.'s [16, 2, 8], and one of the aims of the study is to see how well these biological data can be parallelled using a computer simulation.

In order to answer these questions, a model of interacting agents in a group with a dominance hierarchy is adapted. First movement-alignment behaviour is added and then some group members are given a spatial target. The results on group structure are then investigated.

3 Methods

3.1 Model Basis

To investigate the results of dominance interactions on groups of movement-aligning individuals, I have adapted the DomWorld model, as created by C. K. Hemelrijk [7]. Domworld was chosen because it is relatively easy to adapt so the agents align their movements, since the behavioural rules are already formulated in terms of distances and headings relative to group members. Furthermore, DomWorld comes with a wealth of spatial and statistical analyses.

3.2 DomWorld

DomWorld is an agent-based, spatially explicit model on grouping and dominance interactions. It deals with the emergence of group structure and dominance hierarchies through interactions among agents.

The following is a description of how DomWorld works, adapted from the one given by C. K. Hemelrijk. [7]

The Model

The model is individual-oriented, and consists of a toroid, two-dimensional 'world' and its interacting agents, its visualisation and analytical components. The world is continuous, as opposed to grid-based, therefore agents can move in any direction. Agents have a limited angle and range of vision and thus have asymmetrical front/back embodiment, while they are left/right-symmetrical. The activities of agents are regulated by a timing regime as follows. After each time they act, agents draw a random waiting time from a uniform distribution. The agent with the shortest waiting time is activated first. The decay in waiting time is usually the same for each agent, but if a dominance interaction occurs within a short range of an agent, that agent's waiting time is reduced even more, thus increasing the probability that it will be the next one to act. Agents group and perform dominance interactions according to the sets of rules described below and in Figure 1.

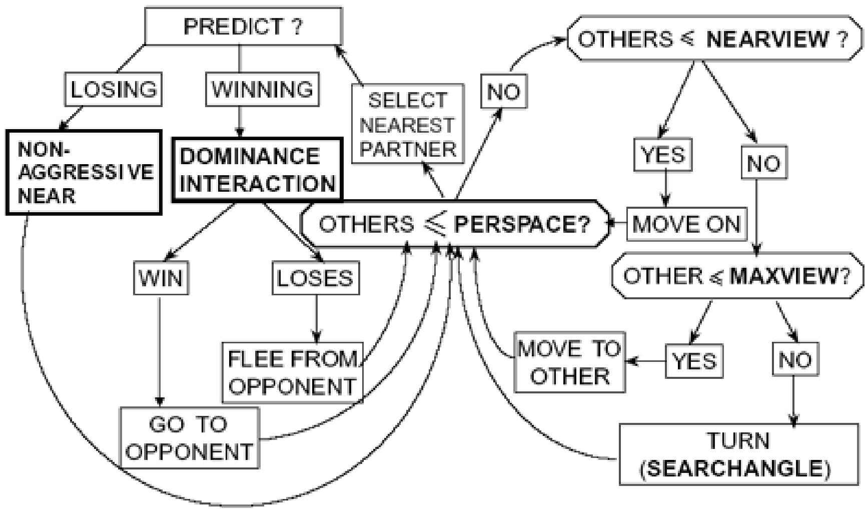

Figure 1: The behavioral rules for DomWorld

Grouping rules

An agent's behaviour depends on the distance to the nearest visible agent as follows:

i If the nearest visible agent is within its 'personal space' (parameter PersSpace), it predicts whether it will win an interaction. If it predicts victory, it performs a dominance interaction with that agent. If it predicts defeat, it remains near without interacting. The agent that wins the interaction moves towards the opponent, while the loser does a full turn and moves away.

ii If the nearest visible agent is within the larger range of NearView, the agent keeps moving in its original direction.

iii If an agent detects its nearest neighbour outside NearView, but within its maximum range of vision (MaxView), it moves towards them.

iv If an agent does not perceive any other agent within MaxView, it searches for group members by turning over an angle (SearchAngle) of 90° randomly to the left or right

Dominance Interactions

Dominance interactions represent competitive interactions over resources that are not specified in DomWorld, but are presumed to include food, mating and spatial location. They are modeled as follows.

i Each entity has a variable dominance value (Dom) which represents its capacity to win a hierarchical interaction.

ii After encountering one another in their PersSpaces, entities 'decide' whether or not to attack using the Risk-Sensitive[6] system. Here, the probability of attacking depends on the potential risk of defeat as follows. On meeting another agent and observing its Dom value, an agent may predict whether it will win or lose on the basis of a 'mental' battle, which follows the rules of a dominance interaction as described below. If the acting agent loses the mental interaction, it will refrain from any action (thus displaying 'non-aggressive' proximity). If it wins the mental battle, it will start a 'real' dominance interaction.

iii If an actual dominance interaction takes place, then entities observe each other's Dom. Subsequently, winning and losing is determined, as follows, by chance and their relative Dom Value.

$$w_i = \left[\begin{array}{ll} 1 & \frac{DOM}{DOM_i + DOM_j} > RND(0,1). \\ 0 & else \end{array} \right.$$

Here, w_i is the outcome of a dominance interaction initiated by agent i (1= winning and 0=losing). In other words, if the relative dominance value of the interacting agents is larger than a random number drawn from a uniform distribution, agent i wins, otherwise it loses. Thus, the probability of winning is proportional with the relative Dom of an agent with the Dom of its partner.

iv Updating the dominance values is done by increasing the dominance value of the winner and decreasing that of the loser, as follows.

$$Dom_i := Dom_i + \left(w_i - \frac{Dom_i}{Dom_i + Dom_j} \right) \times StepDom,$$

$$Dom_j := Dom_j - \left(w_j - \frac{Dom_j}{Dom_j + Dom_i} \right) \times StepDom.$$

The consequence of this system is that it behaves as a damped positive feedback: victory of the higher-ranking agent reinforces their relative

Dom value only slightly, whereas a victory by the lower-ranking agent gives rise to a relatively large change in Dom. To keep Dom values positive, their minimum value was arbitrarily set at 0.01. StepDom is a scaling factor that varies between zero and one and represents the intensity of aggression. In line with the larger rank differences in despotic rather than egalitarian societies, high values imply a large change in Dom value when updating the system and thus indicate that single interactions may strongly influence the future outcome of conflicts. Conversely, low StepDom values represent low impact.

v Winning includes chasing the opponent one unit distance and then randomly turning 45° to the right or left in order to reduce the chance of repeated interactions between the same partners. The loser responds by fleeing under a small random angle over a predefined fleeing distance.

3.3 Changes to operation of DomWorld

Several changes to DomWorld were necessary to get the desired behaviour. They are as follows.

Wait Time

In DomWorld, as has been explained, if a dominance interaction takes place, the time until the next action of all agents within a small radius is also decreased. Thus, agents that are near dominance interactions are more likely to be the ones to act next. This causes a feedback loop whereby agents in the centre of a group will act more often than those at the edges, which means agents at the rear of a movement-aligning group will fall behind and never be able to catch up. To eliminate this problem, the system was changed so that an agent is only eligible to become active once all others have acted. Agents are chosen to act in random order. This new system was chosen because it is the implementation most similar to a truly parallel one, wherein the agents would act simultaneously.

Addition of measurements

The following measurements have been added.

Confusion: The mean squared deviation of each agent's heading from the group's heading. This is a measure of group alignment.

Speed: The distance traveled by the group's centre of gravity in one period, a time-unit in which each agent has acted twenty times.

Average of Nearest Neighbour Distance: The average of the distances of the agents to their nearest neighbour. This is a measure for group density.

Width / Length: The width of the group divided by its length, providing a measure for group shape. Width and length are calculated along the mean group heading, as in Kunz & Hemelrijk's model [13].

Centre Distance: The average distance of each agent to the group's centre of gravity. This is a measure of the surface of the group.

Dominance & Distance to front: The Kendall rank correlation between dominance and distance to the front of the group. This is a measure of the location of the dominant agents in the group.

Frontal location of centre of gravity: The degree to which the centre of gravity is located towards the front of the group.

Together, these measures give a good picture of what happens to the group structure.

Fleeing & chasing direction

An agent that loses a dominance interaction ordinarily flees in the direction away from the victor, with a random error. The victor then moves towards the loser, again with a random error. This has a strongly disruptive effect on the group's alignment, since the fleeing direction may thus be completely opposite to the group direction. To reduce this effect, the fleeing direction is now averaged with the group direction. In the worst case, therefore, the fleeing agent will flee in a direction perpendicular to the group direction.

The chasing behaviour has been completely removed. This is based on the view that dominance interactions in non-aligning groups are about resources, while in movement-aligning groups the main reason for interactions is crowding. This means that in movement-aligning groups the flight of the loser will satisfy the winning individual and no chase will occur.

3.4 Movement alignment

After a predetermined number of periods (currently 260) of interactions under standard DomWorld rules, the agents' behaviour is switched to movement alignment, as outlined below. This interval allows for the emergence of a dominance hierarchy. The switch in behaviour can be interpreted as the group changing from foraging or resting to travel behaviour. During the movement alignment phase, the personal space for each agent is shrunk to half its original size to allow the group to be more cohesive during movement alignment. This can be interpreted as decreasing the tendency of the agents to have dominance interactions.

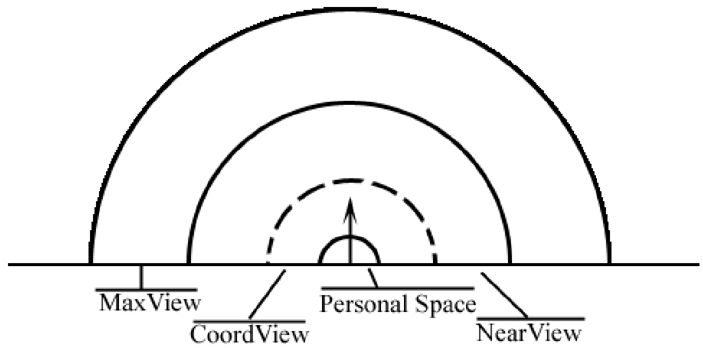

Figure 2: Vision ranges for movement aligning agents. The dashed line indicates the added Coordination View

Once movement alignment starts, the view angle for each agent is increased to 180°. This allows especially the agents near the group front to align their movement more easily, since it decreases the frequency of their searching behaviour. The real-world explanation for this increase in perceptiveness is that individuals are checking on their neighbours more often, and therefore turning their head more, or using auditory cues.

Coordination View

A fourth range is added in between nearview and personal space, called coordination view. Its value is currently set to half that of nearview (See Figure 2).

Movement alignment behaviour

The agents' behaviour depends on the range of the nearest visible entity as follows.

i If the nearest visible agent is in an agent's PersSpace there is no movement alignment. Instead, it predicts whether it will win an interaction. If it predicts victory, it performs a dominance interaction with that agent. If it predicts defeat, it remains near without interacting. The agent that wins the interaction keeps moving straight ahead, while the loser does a full turn and moves away. The direction of flight is averaged with the fleeing agent's alignment direction to make dominance interactions less disturbing to group alignment.

ii If the nearest visible agent is in coordination view, there is movement alignment. Avoidance, alignment and attraction directions are calculated for the visible neighbours and are weighted and averaged. Avoidance is the heading away from the neighbour, alignment is the neighbour's heading and attraction is the heading towards the neighbour. See also Figure 3.

iii If the nearest visible agent is in near view *or* max view, the acting agent moves towards that agent.

Vectors

The model uses vector-based calculations to determine the agents' heading in each time step. The new heading is calculated as follows:

$$\vec{v}_{new} = w_{old} * \vec{v}_{old} + \vec{v}_{mvalign}$$

Where w_{old} is the relative weight of the old direction, compared to the movement alignment vector. The current weight of the old direction is 5, thus limiting the size of the turning angle for each step. The movement

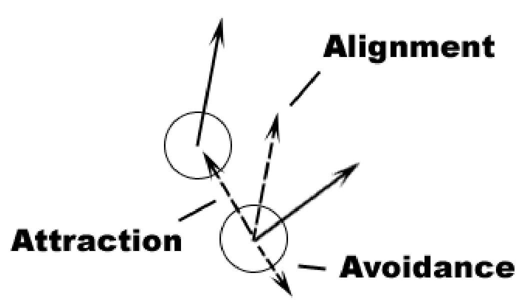

Figure 3: The different behaviour vectors.

alignment vector is determined by weighing and adding three component vectors. These are the standard avoidance, alignment and attraction ones (see Figure 3).

$$\vec{v}_{mvalign} = w_{al} * \vec{v}_{al} + w_{at} * \vec{v}_{at} + w_{av} * \vec{v}_{av}$$

Where w_{al}, w_{at} and w_{av} are the alignment, attraction and avoidance weights, respectively, and $\vec{v}_{al}$, $\vec{v}_{at}$ and $\vec{v}_{av}$ the alignment, attraction and avoidance vectors. The current weights used are $w_{al} = 1$, $w_{at} = 1$ and $w_{av} = 1$ – equal priority for each behaviour, in other words.

The avoidance vector is adding weighted vectors that correspond to the headings *away* from the agents within coordination view. The weights for these vectors are calculated as follows:

$$w_{av} = \frac{1}{dist}$$

The alignment vector is determined by adding weighted vectors that correspond to the directions of the agents within coordination view. The weights for these vectors are calculated as follows:

$$w_{al} = \begin{cases} dist/6 & \text{if } dist \leq 6 \\ 6/dist & \text{if } dist > 6 \end{cases}$$

The attraction vector is determined by adding the weighted vectors pointing towards all other agents within maxview. The weights for these vectors are a low, constant value (currently 0.1).

Looking at the weight functions (see Figure 4), one can see that the different behaviours are dominant in different regions. This parallels the discrete region-based behaviour of Reynold's BOIDS and Huth's fish [18, 9]. Using continuous weights instead of discrete regions results in behaviour that looks somewhat more natural however, with fewer sharp behavioral transitions. In general, only in the most extreme of circumstances, such as an invasion of one's personal space, would I expect real animals to exhibit a sudden switch in priority. Real life, on the whole, is an analog process. The model used in this project reflects this view: it is a hybrid system, using discrete behaviour for the most extreme of the behaviours – seeking others when no one is in view and dominance interactions when someone is too close.

Target source

A target is added to the world, at a certain distance to the centre of gravity of the group. The distance can be varied (though obviously constrained by the size of the 'world', and is currently 400. The target is always placed in the northeasterly direction on the map. Since group orientation is effectively random until movement alignment starts, this set direction does not yield different results from a random direction, and makes visual debugging easier. The source is meant to represent an abstract target, such as a water hole or food source. There is no benefit in being the first to reach the source.

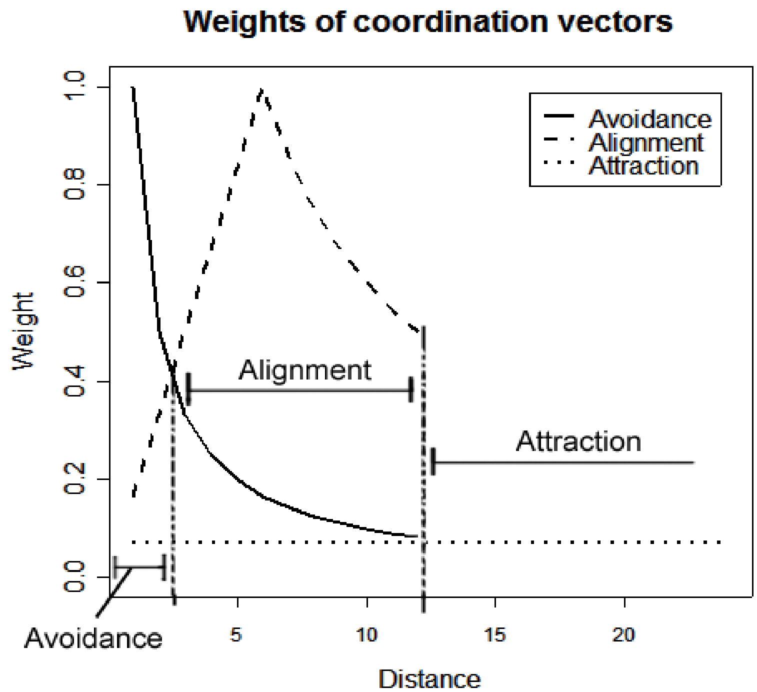

Figure 4: Weights for three responses to neighbours in relation to distance, as used in movement alignment. The dominant influences are indicated for different ranges.

Goal-directed agents

A certain number of agents is given a bias towards the target: After determining their movement direction using the movement alignment rules, they now take the weighted average of an unit vector pointing in the movement direction, and an unit vector pointing towards the target, thus:

$$\vec{v}_{new} = \frac{w_{bias} * \vec{v}_{tgt} + \vec{v}_{old}}{2}$$

Where w_{bias} is the set weight for the directedness.

3.5 Data collection

Groups of 30 agents were used. Groups with both a steep and a shallow gradient of dominance (StepDom 1.0 and 0.1, respectively) were investigated.

Variable name	Used value
Runs per setting	50
GroupSize	30
PersSpace standard DomWorld	4
ViewAngle standard DomWorld	120°
CoordView	12
NearView	24
MaxView	50
High StepDom	1.0
Low StepDom	0.1
StartDom	45
ViewAngle during movement alignment	180°
PersSpace during movement alignment	2
Bias weight	1
OldDir weight	5
Target distance	600
Period duration factor	20

Table 1: Parameter values for data collection

Either the most or least dominant agents were selected to be goal-directed. The number of goal-directed agents was varied from 1 to 28 in steps of 3. The simulation was run 50 times for each setting.

The parameter settings used can be found in Table 1.

Only runs where the group stayed together were used for analysis, because I wanted to examine groups with a consistent group size, among other reasons. The percentage of runs where the group stayed together was tracked and can be found in Figure 8

Movement alignment and goal-directed behaviour are only enabled from period 260 onwards. Runs with goal-directed members reach the goal in at most 32 periods, so the relevant periods are 260-292. The same periods were selected from the runs without goal-directed agents for comparison.

4 Results

4.1 Aggregation behaviour

In standard DomWorld, groups do not align their movement direction. This leads to the following results, which were collected as a control group for the movement-aligning groups.

In non-aligning groups, the group shape is equal in width and length (Figure 5d). This is as expected because there is no constraint on the movement direction for the agents so the direction the group is moving in shifts wildly. This can be seen from the low speed, which implies a high turning rate. This means there is never time for the movement direction to influence group structure. It is also unsurprising, therefore, that the group centre of gravity coincides with the geometrical centre of the group (Figure 5e), or that the average nearest neighbour distance is constant (Figure 5a). Non-aligning groups have a high confusion and low speed (Figure 5b-c).

4.2 Movement alignment

When movement alignment starts, group structure changes as follows[1].

The group becomes denser, as can be seen in the decrease of the average distance to the nearest neighbour (Figure 5a). This increase in density is caused by several factors. Firstly, the agents' attraction to one another is increased. Attraction is always present when movement alignment is on, whereas it only happens in aggregating groups when the only visible agent is at a large distance. Furthermore, due to the smaller personal space, the number of dominance interactions decreases. Since dominance interactions are the main cause of group dispersal, decreasing the number of interactions results directly in an increase in group density.

The centre of gravity shifts towards the group front(Figure 5e). This is similar to the findings of Kunz & Hemelrijk [13], and arises similarly:

Firstly, since speed is constant for all agents, agents that have fallen behind will not be able to catch up easily.

Secondly, because agents in the front of the group have no-one ahead of them, they have a relatively high chance of having visible neighbours to only one side, so that they'll have a tendency to turn to the side, instead of going straight as they would if there was an equal number of visible agents to their left and right. This increases turning rate and hence decreases speed.

Thirdly, agents in the front of the group have a high chance of not having any other agents within their field of view, and thus of engaging in searching behaviour. This will increase their turning rate and thus decrease their speed. This will in turn cause the frontmost agents to be overtaken by one of those directly behind it, and cause crowding in the front, increasing density there.

[1] Since there was little to no difference between despotic and egalitarian groups, results of despotic groups only are shown

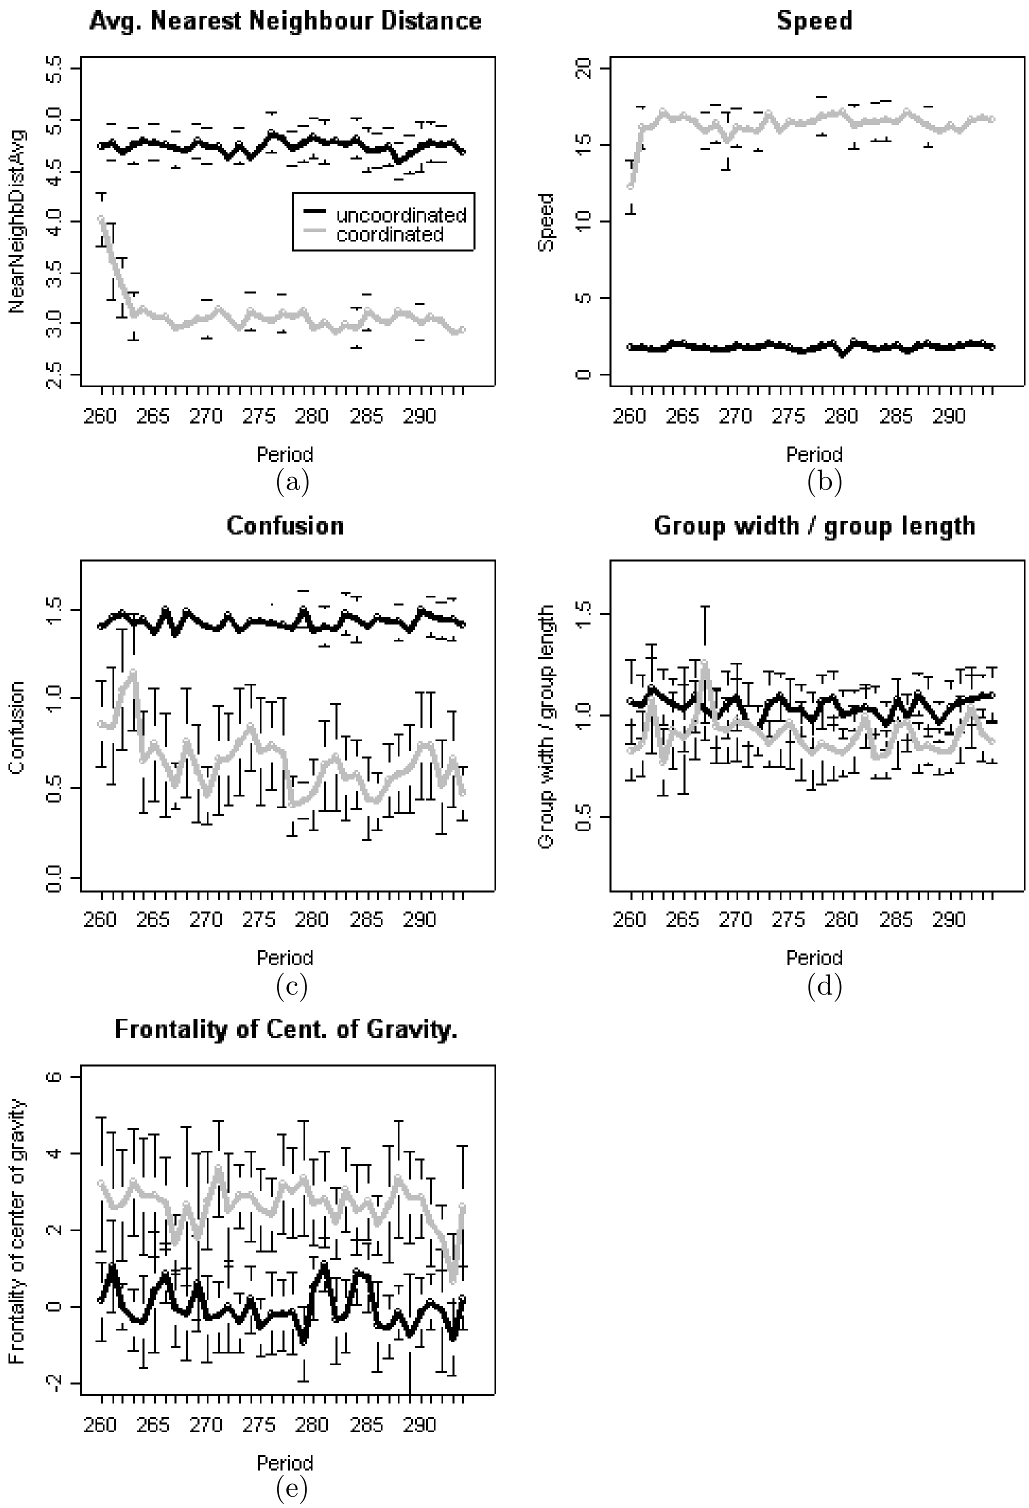

Figure 5: A comparison of movement-aligning and non-aligning groups, for despotic groups. Error bars represent the standard error. Egalitarian groups are similar.

The confusion decreases sharply when movement alignment starts, due to alignment behaviour and the smaller personal space, and the decrease in dominance interactions these entail. The speed of the group increases and turning rate decreases. Since the turning rate and speed are closely correlated, only speed is shown here (Figure 5b–c).

The group shape is of similar length. (Figure 5d).

So far, movement alignment has similar effects in this model as were found in the models by Kunz & Hemelrijk and Reuter & Breckling [13, 17] although the lengthening of the group is weaker here.

No significant correlation between dominance and distance to the group front was found (results not shown).

There was no significant difference between despotic and egalitarian groups during movement alignment for average nearest neighbour distance, speed, confusion, group shape or frontality of centre of gravity.

4.3 Effects of source-directed individuals

Looking at the average nearest neighbour distance, there is little difference whether or not groups include goal-directed individuals. The average distance to the centre of gravity, however, increases slightly (Figure 6a–b). This arises because the group lengthens to almost twice the group width (Figure 6c). The centre of gravity moves slightly (but not significantly) more towards the group front than it does in groups without goal-directed members (Figure 6d).

Because the goal-directed agents' headings are averaged between movement alignment and the direction to the target, their priority for alignment behaviour is weaker than that of other agents. One would expect this to result in higher confusion and lower speed because of an increase in dominance interactions due to weaker avoidance and alignment behaviour. It is at first glance surprising, therefore, that the speed is slightly higher and the turning rate lower when there are goal-directed agents (Figure 6e, f). One of the reasons for this is the omission of runs where the group split up. The runs in which the group stays together until the target is reached will be the ones where confusion is relatively low.

Another explanation is that because the goal-directed agents are all biased towards the same direction, they cause less confusion in the group. This is of course an implicit strengthening of the alignment behaviour within the group of directed individuals. The difference is more pronounced the more goal-directed individuals there are in the group. The difference in confusion, however, does not differ significantly between undirected groups and groups

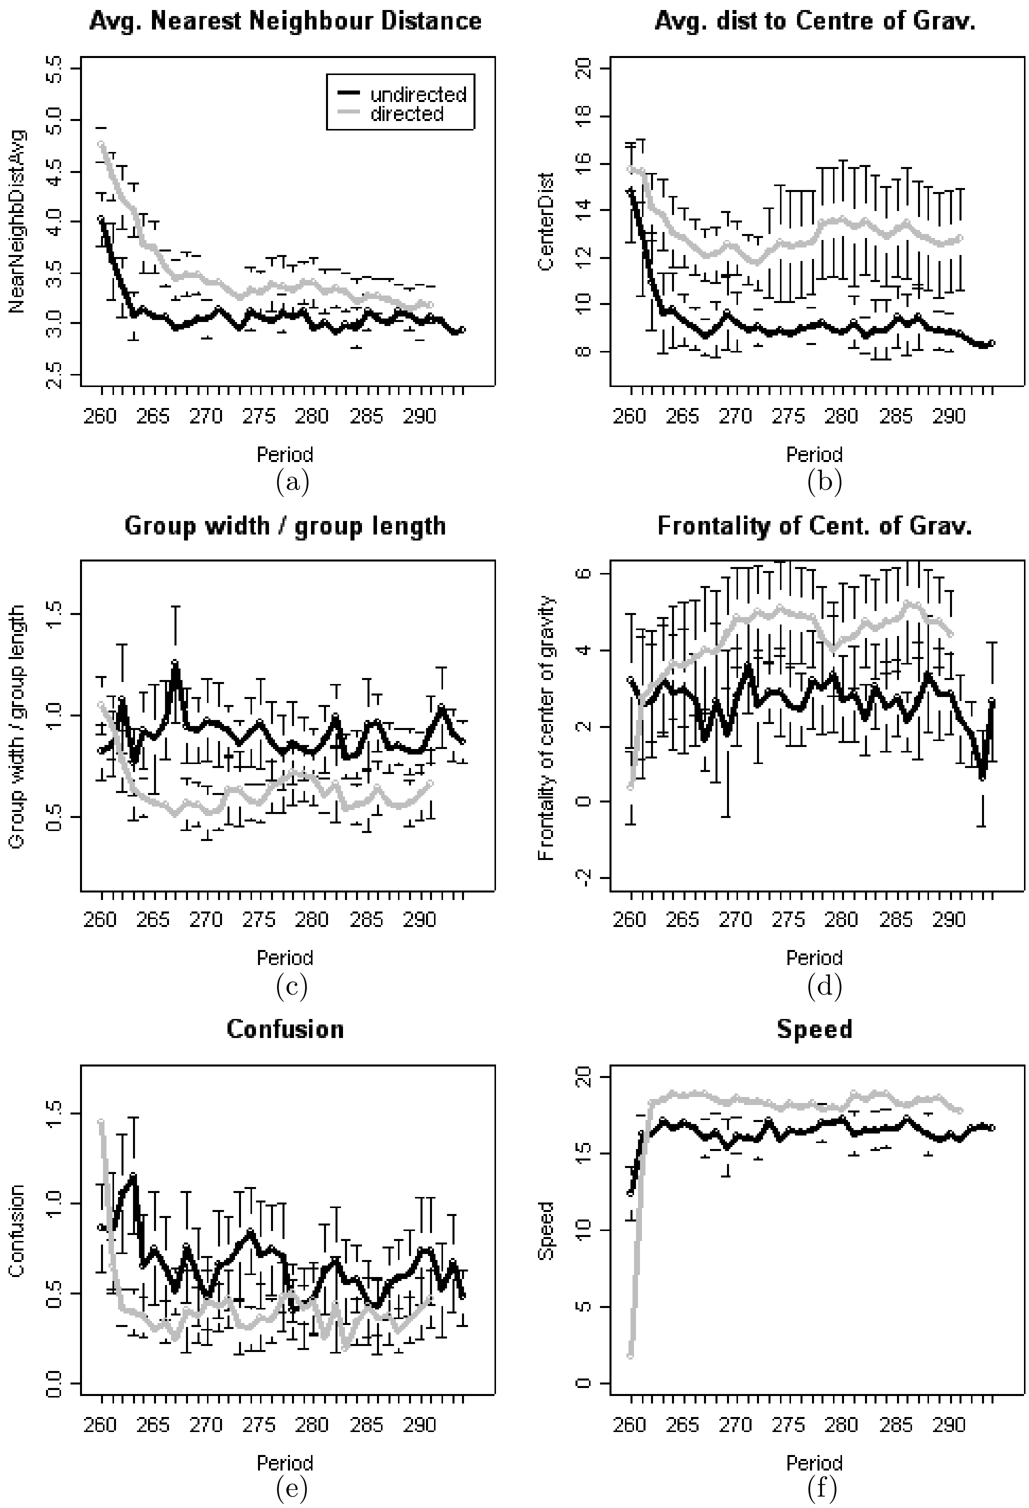

Figure 6: Comparison between movement-aligning groups containing either 10 or no goal-directed individuals. Error bars represent the standard error.

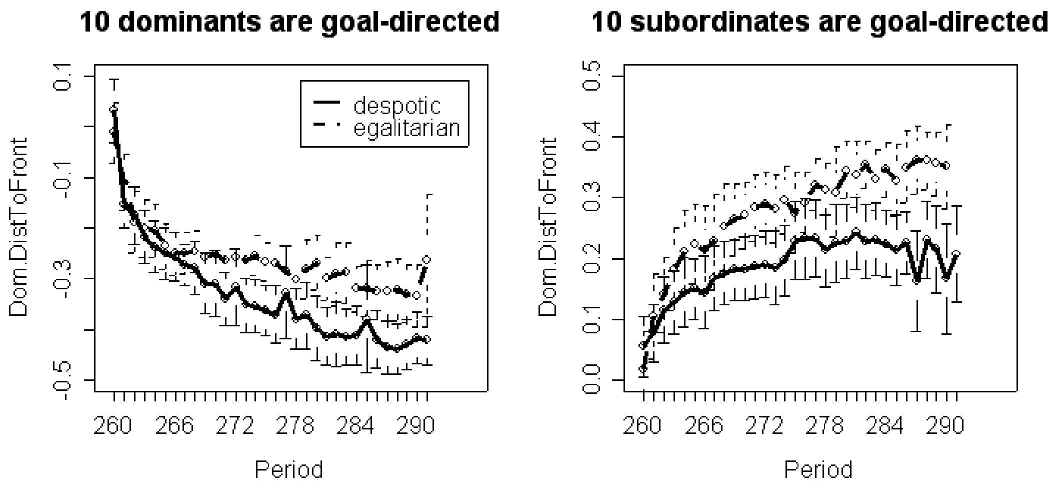

Figure 7: Kendall rank correlation between dominance and distance to front for dominant and subordinate goal-directed group members

containing goal-directed agents (see Figure 6e).

Figure 7 shows the Kendall rank correlation between dominance and distance to the front of the group if goal-directed individuals are dominant or if they are the most subordinate individuals in two types of societies, despotic and egalitarian.

These results lead to several conclusions. Firstly, the directed agents clearly move towards the front of the group over time: if the directed agents are dominant there is a negative correlation between dominance and distance to the group front, and when they are subordinate there is a positive correlation, which means that in the first case the agents at the front of the group are predominately those with high dominance, and vice versa. The correlation gets stronger over time, which shows that the goal-directed agents keep moving forward in the group gradually.

Secondly, the correlation is stronger for dominant directed agents in despotic groups than it is in egalitarian groups. This can be explained as follows. Since losing a dominance interaction results in flight behaviour, which will generally not be towards the group front, and since in successful runs the group heading is generally towards the goal, the loss of a dominance interaction then points the losing agent *away* from the target. This means that a loss impedes an agent's progress both towards the goal, and towards the group front. The difference between dominant and subordinate individuals is much more pronounced in despotic groups than in egalitarian ones, which means that in despotic groups the high-ranked individuals win a much higher percentage of the dominance interactions they engage in than those of equal rank do in an egalitarian group. This in turn explains why goal-directed agents, if they are the most dominant members of the group, automatically reach the front of the group earlier.

Thirdly, the correlation is stronger for subordinate directed agents in egalitarian groups than it is in despotic groups. The explanation is similar to the one above: in egalitarian groups, subordinate agents win many more dominance interactions than they do in despotic groups, and therefore automatically end up at the front sooner.

As can be seen in Figure 8a, there is no significant difference between despotic and egalitarian groups in the percentage of groups that reach the target. This is unexpected, since it has previously been shown that a difference in intensity of aggression leads to differences in group structure and dynamics [7]. Especially the difference in group spread between egalitarian and despotic groups is a factor that one would expect to influence the rate at which groups split up.

Unexpectedly, the success rate of a group is not influenced by whether either dominant or subordinate individuals are the ones to be goal-directed. Since losing a dominance interaction will generally cause an individual to deviate from its previous heading, one would expect the groups with subordinate goal-directed individuals to have lower success rates. This can be explained as follows.

From viewing several runs, there are two main ways for the group to split up. Firstly there are the cases where a group of individuals simply gets turned away from the rest of the group and moves off. This is usually caused by a series of lost dominance interactions that cause the agents to turn away from the group. If the goal-directed agents are spread through the group this causes an increase in dominance interactions, and lead to group breakup. Secondly, since the goal-directed individuals move more efficiently as a group, if the whole front of the group consists of goal-directed agents they move faster than the undirected agents in the group rear, and thus leave them behind. An explanation for the lack of difference in success percentages between the different types of runs is that there is a balance between these two ways for the group to break up: If the goal-directed agents get to the front more easily, group breakups will generally be caused by the goal-directed agents leaving the group behind, and vice versa. Visual

(a)
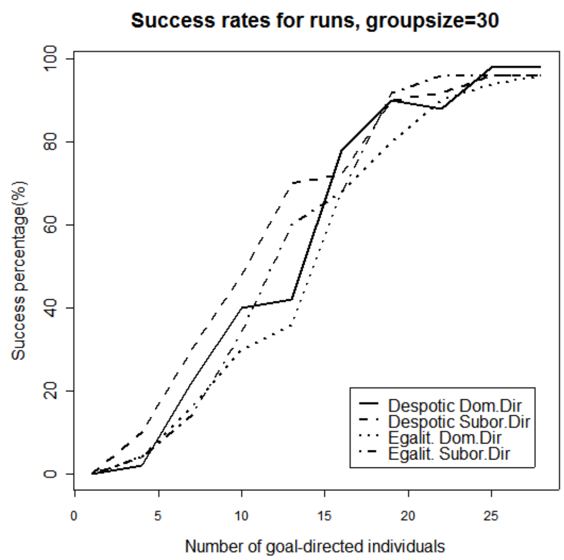

(b)
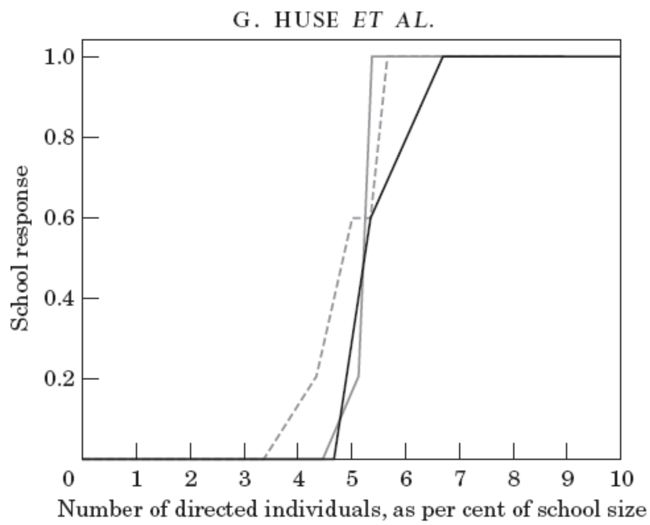

Figure 8: *a*): The percentage of groups in which the target was reached for the various settings. Dom.Dir means the goal-directed individuals are the most dominant ones, Subor.Dir means the goal-directed individuals are the most subordinate ones.
b): From Huse et al. [8]. Fraction of group that followed a number of goal-directed individuals to a target.

examples of these types of breakups can be found in appendix A.

5 Discussion

5.1 Discussion of results

The model shows several emergent phenomena: The group shape, density, speed and confusion all change both with movement alignment and in the presence of goal-directed individuals. Movement-aligning groups behave similarly to those in other models, for example those by Huth and Kunz & Hemelrijk [9, 13].

The model shows that compared to groups without goal-directed agents, groups containing goal-directed individuals have a more oblong group shape, lower turning rate and higher speed, have a centre of gravity that lies more towards the front of the group and a larger average centre distance, while simultaneously the similar nearest neighbour distance remains similar. These changes can all be explained by an increase in movement efficiency due to the goal-directed agents.

The goal-directed agents consistently end up at the front of the group, without this being implemented in their rulesets. This can be seen in the increasing strength of the correlation between distance to the front of the group and dominance. A visual example can be found in Appendix A. Since the agents have no positional preference, another factor must be the cause. I explain this by their lower deviation from the direction to the target, hence their higher speed. This corresponds with the alternative explanation Krause and Bumann [12] give for their finding that hungry fish are consistently found in the front of shoals of juvenile roach: a higher swimming speed and turning rate related to foraging behaviour. The idea that the frontal position of the fish is not intentional but due to their swimming characteristics is further supported by their finding that the hungry fish were to be found at the front of the group increasingly over the course of 30 minutes. An automatic movement of goal-directed individuals to frontal positions is consistent with my results. If fish had an intentional positional preference they probably would manifest this sooner, since real fish can vary their speed and thus can reach the group front very quickly if they desire.

Whether goal-directed agents are high or low in rank does not influence the group's success rate (See Figure 8a). Again, this is probably due to the trade-off between different causes of group breakup mentioned previously.

Though the goal-directed individuals were not given any preference for a position within the group, they ended up at the front. This is because the rest of the group members do not share their priority towards the target, and will therefore deviate from it more, both through random error and local effects of movement alignment. Therefore, they will move less directly towards the target, and their effective speed is lower. This also explains the slight forward shift of the centre of gravity over that seen in regular movement-aligning groups. The stragglers have even more trouble catching up, since the front of the group is moving faster, while the front of the group is denser due to a decrease in dominance interactions. A visual example can be found in appendix A.

My results show a smooth connection between the percentage of goal-directed individuals and success percentage. This differs from the results by Huse et al. [8] (see Figure 8). which show a clear phase transition if 5% directed individuals are goal-directed. This qualitative difference might be due to the dominance interactions in this model, or due to the differences in implementation of movement coordination behaviour. Huse et al. based their model on Reynolds' BOIDS [18]. The most obvious difference is that in the BOIDS model the agents are not limited in their angle of perception.

Another factor is that in Huse et al.'s implementation the group is initially milling in a small ring. This means the group members are already polarised and aligning with their neighbours. In my implementation the group is not aligning their movement direction when goal-directed behaviour sets in. The increased randomness in the initial conditions makes it more difficult to have the group follow the goal-directed agents, which might also account for the much lower success rate in this model.

5.2 Problems and possible improvements

A problem with this version is that there is very little apparent influence of the intensity of aggression on group structure as previously demonstrated by Hemelrijk [7], This is puzzling, since I made no changes to the basic DomWorld behaviour. I have investigated the matter thoroughly and have been unable to find the cause. Nonetheless, the model behaves well enough otherwise that I have carried on regardless.

One of the obvious additions that could be made to the model would be to track subgroups as well and analyse, for example, what percentage of the group reaches the target. This is not an easy task, however, and it is doubtful how much it would add exactly, since many of the measures are only interesting for relatively large subgroups of agents. The degree of confusion of a group is only interesting for groups of 8 agents or above.

An oft-heard criticism from biologists on the kind of abstract research conducted in this study is that it lacks biological plausibility and is either hard or impossible to test. Often these critics are missing the point, which is that we aren't investigating reality so much as we are investigating *behaviour*. Any insights in how group behaviour *might* arise is worthwhile in and of itself, both because it can yield inspiration for additional research to be carried out in reality and because it can, from the Artificial Intelligence point of view, aid in the creation of groups of artificial agents that exhibit useful behaviour. While it is always true that real life is more complicated, it remains important to find simpler ways to explain observed phenomena.

From the perspective of Artificial Intelligence, it is always worthwhile to investigate new ways to create intelligent behaviour, especially those ways that are not as computationally and conceptually complex as the first attempts made in the field.

5.3 Additional Research

There are several avenues of research that might lead to interesting results. Firstly, as an extra control, these results need to be compared to schools without dominance interactions.

Secondly, it would have been interesting to analyse the exact method of group leadership more closely, to see how closely the model parallelled the results by Reebs [16] and Bumann & Krause [2], in other words to note in how many of the successful runs the frontmost agent was a goal-directed one, and to what degree the direction of the frontmost individual influenced group direction.

Thirdly, the group structure and dominance hierarchy were allowed to emerge randomly before movement alignment was activated. It would be interesting to take some pre-determined spatial distributions of agents with pre-set dominance hierarchies and investigate the results. The effects of the gradient of the dominance hierarchy could be investigated by placing all the directed individuals at the rear of the group, and looking at the time interval before they get to the front of the group.

6 Acknowledgements

I would like to thank the following people for helping me in many ways:
Charlotte Hemelrijk for her inspiring supervision.
Hendria van der Haar for basically just being herself.
Caroline Lemmens & Egbert Smit for their support.

References

[1] T. C. M. Bakker, E. Feuth-DeBruijn, and P. Sevenster, *Asymmetrical effects of prior winning and losing on dominance in sticklebacks*, Ethology, 82 (1989), pp. 224–229.

[2] D. Bumann and J. Krause, *Front individuals lead in shoals of three-spined sticklebacks (Gasterosteus aculeatus) and juvenile roach (Rutilus rutilus)*, Behaviour, 125 (1993), pp. 189–197.

[3] R. W. Byrne, *How monkeys find their way: Leadership, coordination and cognitive maps of african baboons*, On the move. How and why animals travel in groups, (2000), pp. 491–518. ed. S. Bonski and P. A. Garber.

[4] J. L. Deneubourg, S. Goss, N. Franks, and J. M. Pasteels, *The blind leading the blind: Modelling chemically mediated army ant raid patterns*, Journal of Insect behaviour, 2 (1998), pp. 719–725.

[5] W. D. Hamilton, *Geometry for the selfish herd*, Journal of Theoretical Biology, 31 (1971), pp. 295–311.

[6] C. K. Hemelrijk, *Risk sensitive and ambiguity reducing dominance interacitons in a virtual laboratory*, From animals to animats: Proceedings of the Fourth International Conference on Simulation of Adaptive Behaviour, (1998), pp. 255–262. ed. R. Pfeifer, B. Blumberg, J.-A. Meyer and S. W. Wilson.

[7] ——, *An individual-orientated model of the emergence of despotic and egalitarian societies*, Proceedings of the Royal Society London B: Biological Sciences, 266 (1999), pp. 361–369.

[8] G. Huse, S. Railsback, and A. Fernö, *Modelling changes in migration pattern of herring: collective behaviour and numerical domination*, Journal of Fish Biology, 60 (2002), pp. 571–582.

[9] A. Huth, *Ein Simulationsmodell zur Erklärung der kooperativen Bewegung von polarisierten Fischschwärmen*, PhD thesis, Department of Biology and Physics, University of Marburg, 1992.

[10] S. A. Kauffman, *At Home in the Universe: the search for laws of complexity*, Oxford University Press, 1995.

[11] A. P. Klimley, *Schooling in the large predator, sphyrna lewini, a species with low risk of predation: a non-egalitarian state*, Ethology, 70 (1985), pp. 297–319.

[12] J. Krause, D. Bumann, and D. Todt, *Relationship between the position preference and nutritional state of individuals in schools of juvenile roach(Rutilus rutilus)*, Behavioral Ecology and Sociobiology, 30 (1992), pp. 177–180.

[13] H. Kunz and C. K. Hemelrijk, *Artificial fish schools: Collective effects of school size, body size, and body form*, Artificial Life, 9 (2003), pp. 237–253.

[14] A. F. M. Marée, *From pattern formation to morphogenesis, multicellular coordination in* Dictyostelium discoideum, PhD thesis, Department of Theoretical Biology and Bioinformatics, University of Utrecht, 2000.

[15] M. J. Mataric, *Designing emergent behaviors: From local interactions to collective intelligence*, Proceedings, From Animals to Animats 2, Second International Conference on Simulation of Adaptive Behavior (SAB-92), (1992), pp. 432–441. ed. J-A. Meyer, H. Roitblat and S. Wilson.

[16] S. G. Reebs, *Can a minority of informed leaders determine the foraging movements of a fish shoal?*, Animal Behaviour, 59 (2000), pp. 403–409.

[17] H. Reuter and B. Breckling, *Selforganisation of fish schools: an object-oriented model*, Ecological Modelling, 75/76 (1994), pp. 147–159.

[18] C. W. Reynolds, *Flocks, herds, and schools: A distributed behavioral model, in computer graphics*, Computer Graphics, 21(4) (1987), pp. 25–36.

[19] W. L. Romey, *Individual differences make a difference in the trajectories of simulated schools of fish*, Ecological Modelling, 92 (1996), pp. 65–77.

[20] Y. Sakakura and K. Tsukamoto, *Ontogeny of aggressive behaviour in schools of yellowtail, Seriola Quinqueradiata*, Environmental Biology of Fishes, 56 (1999), pp. 231–242.

[21] R. V. Sole, E. Bonabeau, J. Delgado, P. Fernadndex, and J. Marin, *Pattern formation and optimization in army ant raids*, Artificial Life, 6 (2000), pp. 219–226.

[22] L. Steels, *Towards a theory of emergent functionality*, From Animals to Animats: Proceedings of the First International Conference on Simulation of Adaptive Behavior (SAB90), (1991), pp. 451–461. ed. J.-A. Meyer and S. W. Wilson.

[23] ——, *The artificial roots of artificial intelligence*, Journal of Artificial Life, 1 (1994), pp. 89–125.

[24] ——, *Emergent functionality in robotic agents through on-line evolution*, Artificial life IV: proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems, (1994), pp. 8–14. ed. Rodney Brooks and Pattie Maes.

A Screenshots

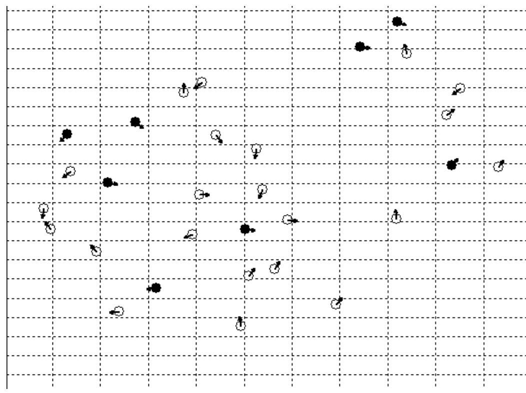

Figure 9: A group using standard DomWorld rules. Dark circles represent the most dominant agents.

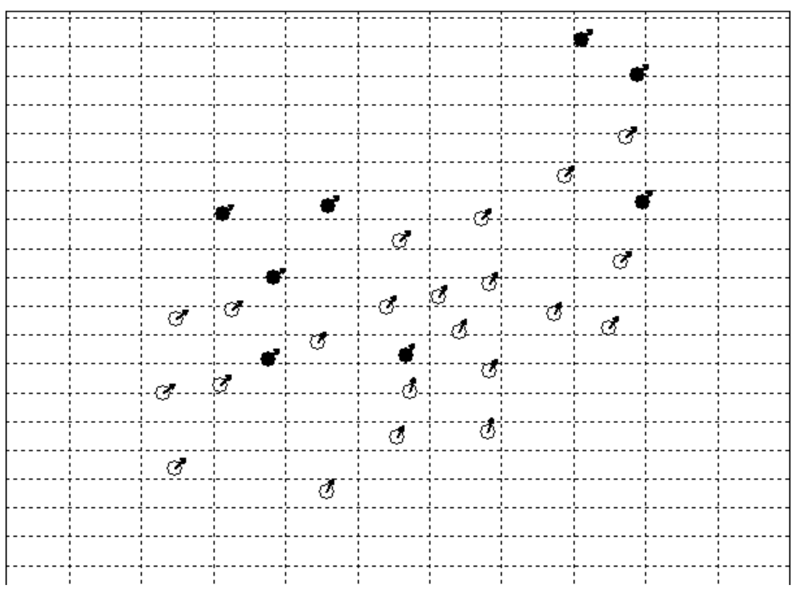

Figure 10: A group using movement-alignment rules shortly after movement alignment has switched on. Dark circles represent the most dominant agents.

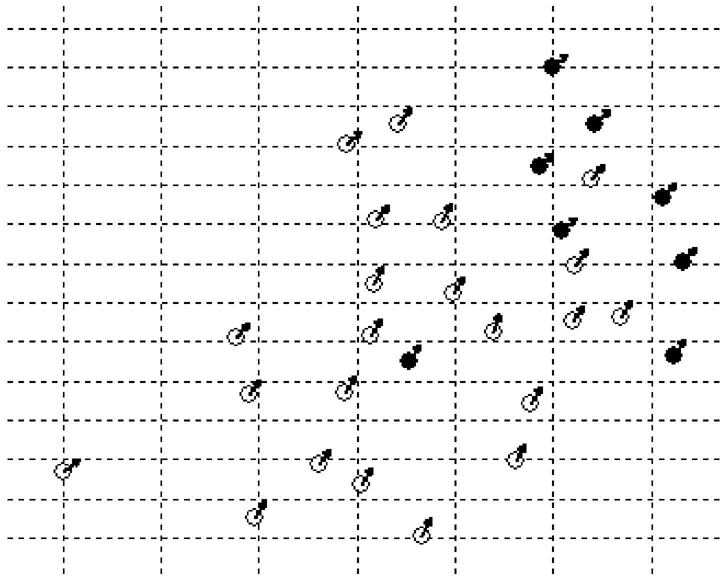

Figure 11: A group using movement-alignment rules at 20 periods after movement alignment has switched on. Dark circles represent the goal-directed agents. Goal-directed agents are the most dominant ones.

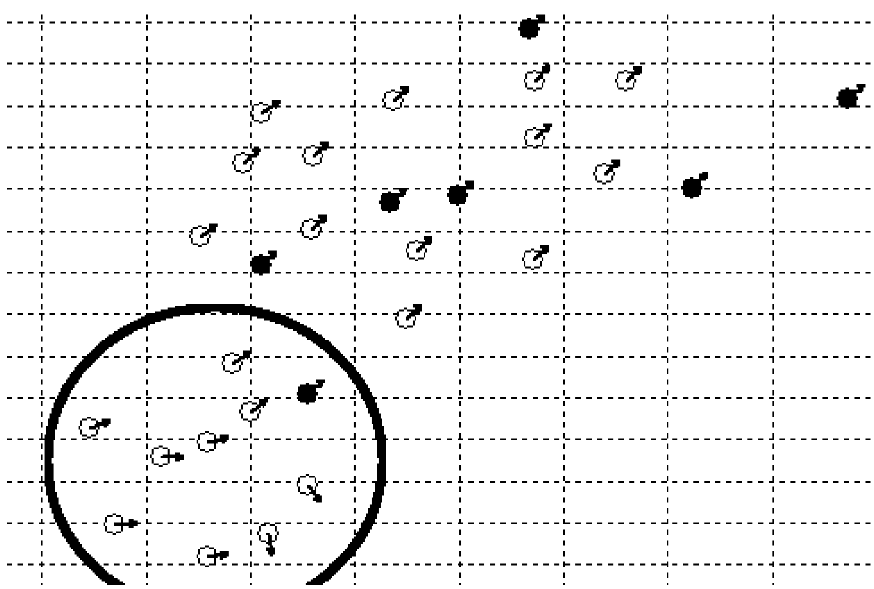

Figure 12: A group with the beginning of a possible breakup, indicated by the circle.

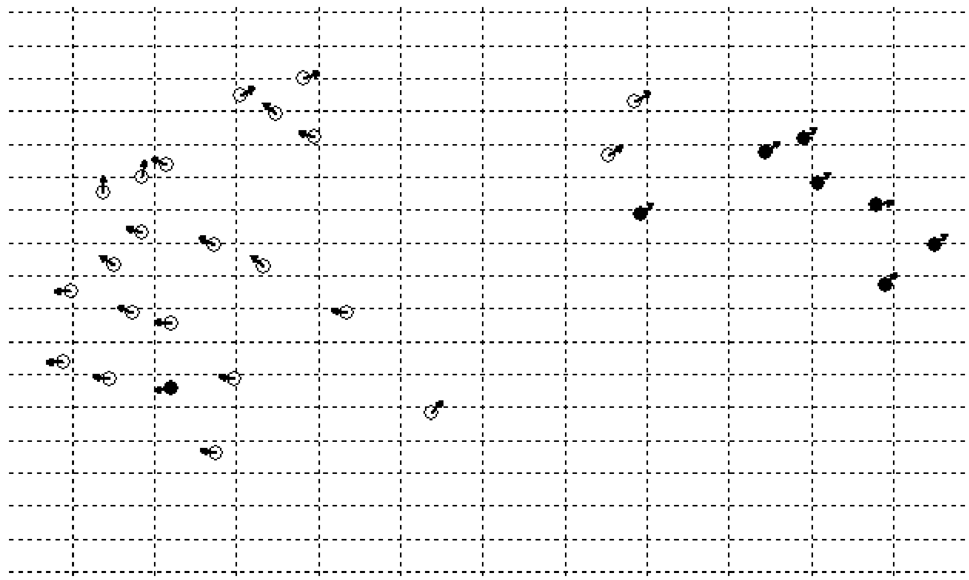

Figure 13: A typical breakup. Dark circles represent goal-directed agents. Goal-directed agents are the most dominant ones.

B Program documentation

B.1 Program structure

DomWorld was originally written in object Pascal (Borland Pascal 7.0) and has since been rewritten in Delphi, adding a graphical interface.

The program is object-oriented, with the following class structure, which can also be seen in Figure 14:

SimulationMain This is the enveloping system class, which contains both the simulation world and the observer class which is responsible for data collection and statistical analysis. SimulationMain takes care of the communication between these classes and runs the simulation. It tells the world to run every activation, and passes the result to the observer. After a number of activations (called periodFactor and defined in the constants file) the observer is asked to run a set of analyses on the collected data.

World The world class contains all information related to the simulation world. It contains the list of agents, and most of the functions concerned

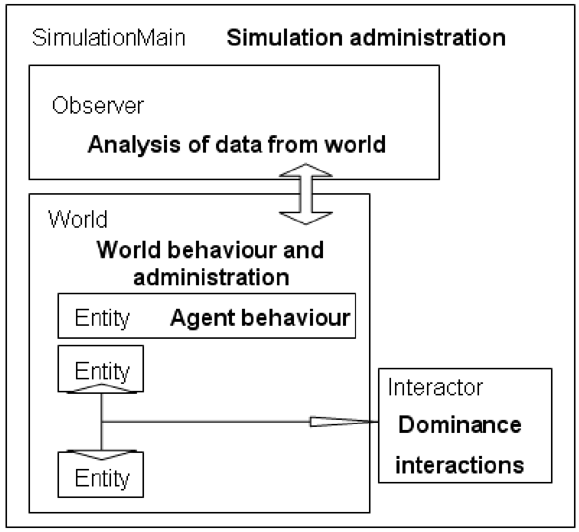

Figure 14: Domworld's class structure. Bold text indicates classes' function.

with the details of running the world, such as the distribution of the agents over the world and the timing scheme, selecting which agent is next to act.

Entity The entity class implements the simulation's agents. They have coordinates, orientation, dominance and all other variables relevant to their operation (see also Table 2). They maintain a list of pointers to all agents within their various vision ranges and to their nearest neighbour to improve efficiency when checking their various neighbourhoods. When two entities interact they do so through an *interaction* object. The Entity class also implements the agents' behaviour, in the (oddly-named) function TimeSchedule. An outline of this function's operation can be found in the *program compilation and use* section further on.

Interaction The interaction class implements the dominance interactions between agents. It contains an attacker and a defender, which are pointers to entities. It contains functions to determine the winner, and to adjust the entities' dominance accordingly.

Observer The observer class implements all the functions used in the analysis of group structure, as well as a bevy of statistical analyses. All the group structure measures discussed in section 3.3 are implemented in this class. The class methods also includes the functions used to write the gathered data to file.

Data member	Function
Id	Unique identifier
PosX, posY	Coordinates on world map
Nose	Heading
Dominance	Dominance
StepDom	Intensity of aggression
FleeDist	Distance to flee after losing
ChaseBase	Distance to give chase after winning
SearchAngle	Angle to turn when seeking others
AttackStrategy	Either obligate, risk-sensitive, ambiguity-reducing [7]
PersSpace	Personal space
CoordView	Coordination view
NearView	Near view
MaxView	Maximum view
ViewAngle	View angle

Table 2: Data members for Entity class

B.2 Program interface

In porting DomWorld to Delphi, a visual interface has been added. There is a main system window (see Figure 15) from which several other windows can be accessed through the menu bar. The main window also shows the results of the analyses that take place at the end of the simulation. There are indicators which show the current run name, period and activation, and there is a "run for" box which can be used to run multiple simulations one after another.

The windows that can be accessed from the menu bar are the following. Firstly, several matrix information windows can be accessed from the "matrices" menu. These all look similar to the one shown in Figure 16(a). The available matrices show who has attacked who, who has entered in whose personal space without attacking, who has won how many interactions with whom, the frequencies with which the agents' various behaviours (for example search, attack or walk) have occurred, and finally what the agents' dominances are.

Secondly, the settings window can be accessed from the "parameters" menu. This is pictured in Figure 16(b). The three tabs can be used to change most of DomWorld's settings.

Thirdly, two windows can be accessed from the "visualisations" menu. The first is the world chart, which shows the agents on a two-dimensional

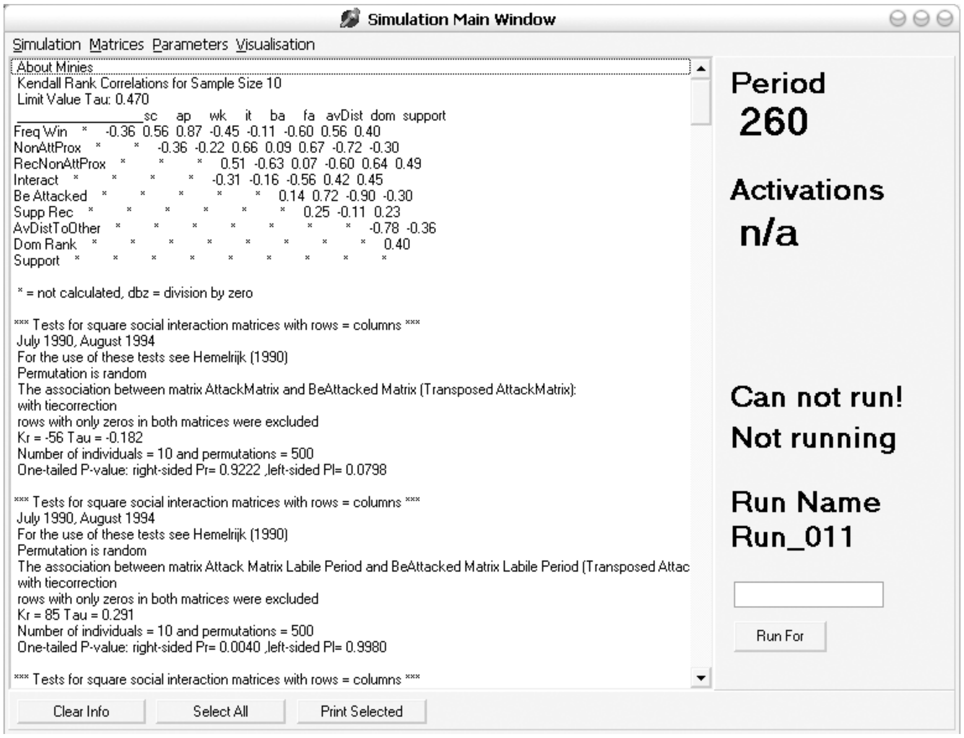

Figure 15: Domworld's main window

grid, along with several controls which can be used to control the amount of information displayed along with the agents. It is pictured in Figure 17(a). This window is especially handy for visual debugging. The agents move in realtime, and the map can be dragged and resized on the fly.

The second window available from the "'visualisations" menu is the dominance chart. It is pictured in Figure 17(b). It shows the progression of the dominance for each of the agents over time.

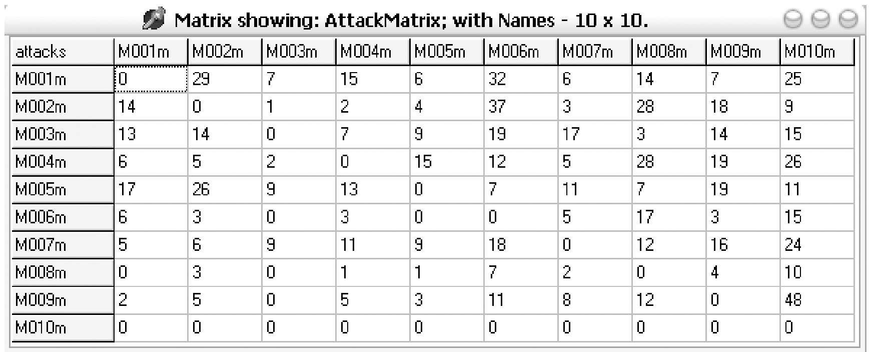

Matrix showing: AttackMatrix; with Names - 10 x 10.

attacks	M001m	M002m	M003m	M004m	M005m	M006m	M007m	M008m	M009m	M010m
M001m	0	29	7	15	6	32	6	14	7	25
M002m	14	0	1	2	4	37	3	28	18	9
M003m	13	14	0	7	9	19	17	3	14	15
M004m	6	5	2	0	15	12	5	28	19	26
M005m	17	26	9	13	0	7	11	7	19	11
M006m	6	3	0	3	0	0	5	17	3	15
M007m	5	6	9	11	9	18	0	12	16	24
M008m	0	3	0	1	1	7	2	0	4	10
M009m	2	5	0	5	3	11	8	12	0	48
M010m	0	0	0	0	0	0	0	0	0	0

(a) Example of matrix display window. Shown is the attack matrix.

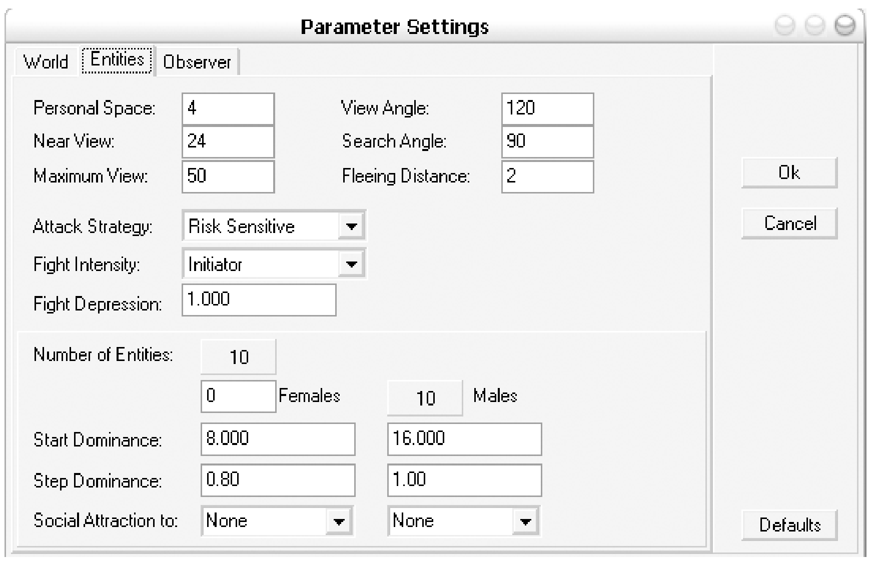

(b) Domworld settings window. Shown is the tab with settings for entities

Figure 16: Domworld visual interface - matrix and settings window

B.3 Program compilation and use

The program is a Delphi project, which means that it is Windows-only. It consists of a project file, a series of files with program code (generally speaking, one for each class and several implementing utility functions and containing constants) and a “form” file defining each of the project windows. Compilation is as simple as opening the project file in the Delphi suite and giving the “compile” command.

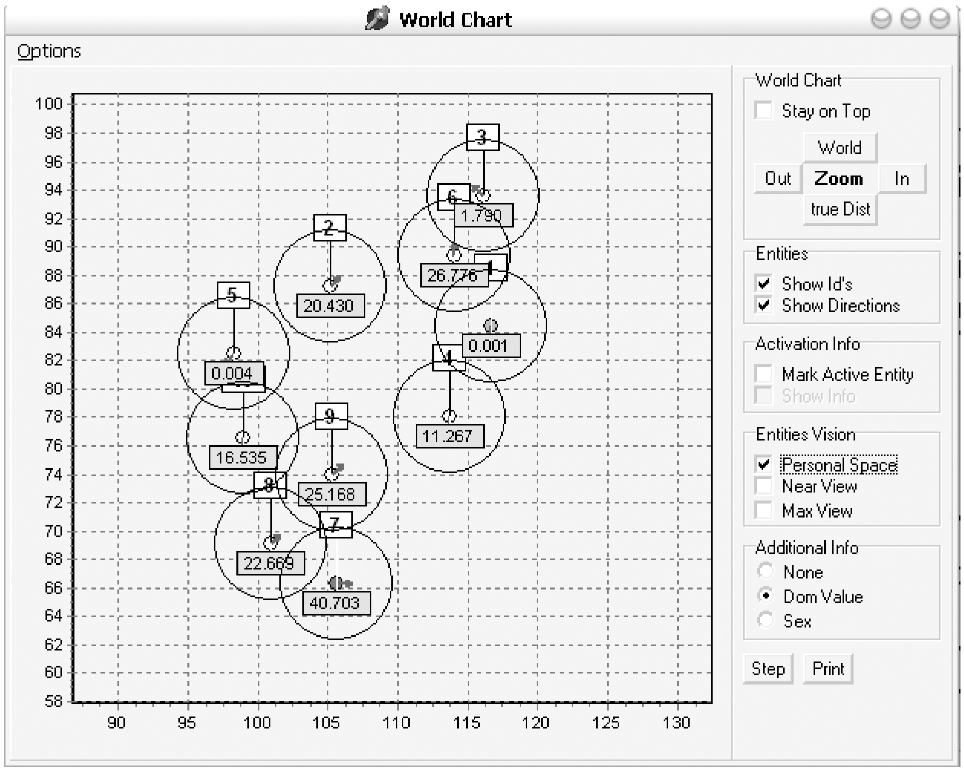

(a) Domworld's world display.

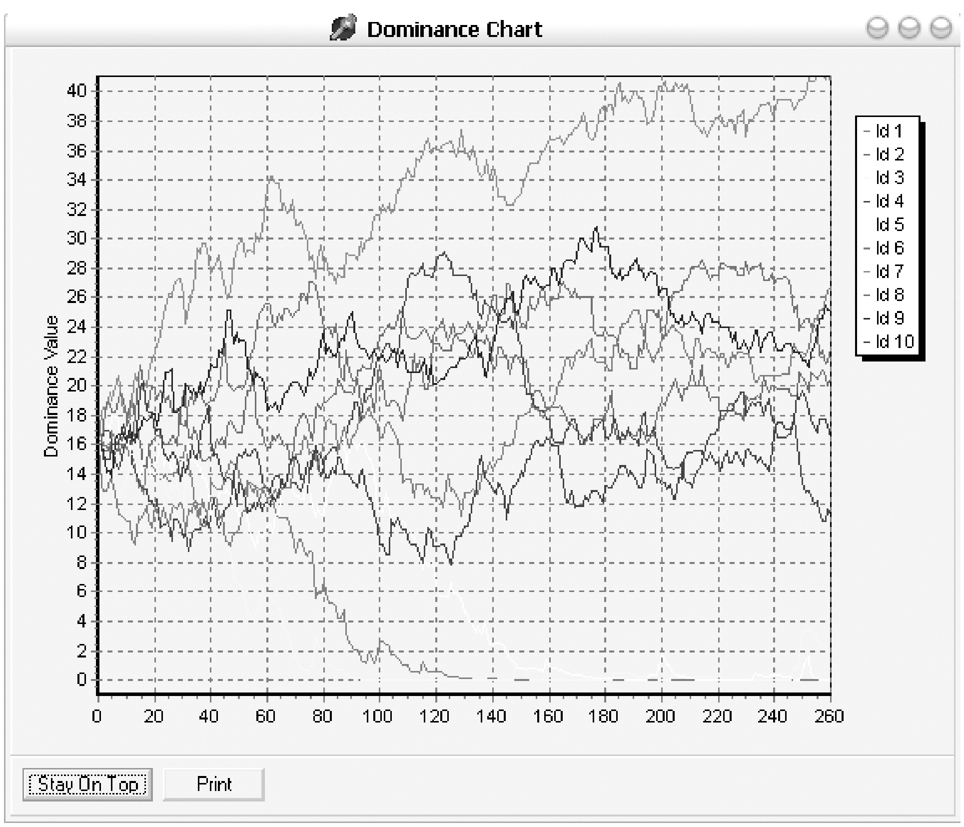

(b) Domworld's dominance chart display.

Figure 17: Domworld's world and dominance chart displays

Although there are several optimisations, each agent's behaviour is based largely on their neighbourhood, which means that the amount of operations to be executed rises geometrically with the number of agents. Using groups of over approximately 30 agents was not viable on the hardware used (an AMD Athlon XP 2400+ with 512 Mb of RAM running Windows XP).

The operation of the Timeschedule function is shown in the following block of pseudocode. This is the heart of the simulation, implementing the agents' behaviour.

Algorithm 1 TimeSchedule

```
Find nearest neighbour
if nearestneighbour = nil then
   Turn SearchAngle {no neighbour in view}
   Move
else
   if distance > nearview then
      {nearest neighbour in MaxView}
      if coordination enabled then
         Attraction
      else
         Move to neighbour
      end if
   else if (coordinationenabled and distance > CoordView) or distance > PersSpace then
      {nearest neighbour in NearView}
      if coordination enabled then
         Attraction
      else
         Follow nose
      end if
   else if coordinationenabled and distance > PersSpace then
      Coordinate {nearest neighbour in CoordView}
   else
      Dominance interaction {nearest neighbour in PersSpace}
   end if
end if
```