Primate Group Decision Making Simulation

Software Requirements and Design Documentation

Author: Julian Zappala



Contact: jxz@cs.nott.ac.uk

School of Computer Science University of Nottingham Jubilee Campus Wollaton Road Nottingham, NG8 1BB

1 Introduction

This document details software requirements and design documentation created during the replication of a multi-agent simulation first published by Sellers et al. [2007] and concerning group decision making within a troop of chacma baboons (*Papio hamadryas ursinis*).

The software described herein has been implemented in Java using the MA-SON multiagent simulation library (version 13) which can be obtained from http://cs.gmu.edu/~eclab/projects/mason/

2 Software Requirements Specification

The objectives of this work include the implementation of a generic model for group decision making, and the application of this model to the replication of the model created by Sellers et al. [2007]. Below, these objectives are translated into a software requirements specification. Firstly, requirements specifically concerned with the provision of group decision making are stated. Following this, requirements of the work to be replicated are given; this work is described in some detail.

2.1 Requirements for a model of group decision making

The software is required to provide a simple, yet flexible mechanism permitting a group of individuals to arrive at a consensus decision. A group, in this context, is considered to be an aggregation of two or more conspecific individuals; the population (of all individuals) is an aggregation of one or more groups. For the purposes of this work it is reasonable to assume that group membership is static, and that the group, as a whole, will remain cohesive.

The consensus decision shall be made in respect to a finite and known set of choices from which each individual may express a single preference, creating a set of individual preferences. Individuals shall not be bound to provide a preference, and may elect not to do so; as such, the software should provide for unshared, partially shared and equally shared consensus decisions. A policy should then be applied to this set of preferences to determine a single, group preference representing the group decision; each individual is then informed of this decision.

The policy must permit the quorum threshold to be specified, and in doing so permit decisions to be made by either a majority, sub majority, or super majority of individuals. In the case of a sub majority threshold more than one subset of preferences (a subset here meaning a set of similar individual preferences) may meet this threshold; given this situation the consensus decision should be determined as the larger of the eligible subsets, ties may be resolved either randomly or heuristically, as appropriate. Where the quorum threshold is not achieved by any subset, i.e. consensus has not been reached the individuals should be informed of this.

On receiving notification of the group's decision each individual should then incorporate this information into their final choice. Where the group decision is aligned with that of the individual they may proceed with their preferred choice, however, where there is disparity between the "will" of the individual and that of the group the individual should concede to the group decision, revising its preference accordingly.

2.2 Requirements of the work to be replicated

The requirement is to re-implement an extant multi-agent simulation of group decision making; the work of Sellers et al. [2007]. The specification given here has been drawn both from the published material, and also directly from the (freely available) source code of their simulation which throughout this work has been treated as a working specification. Only details relevant to the specification of the software required for this study are considered here; for a detailed description of the study conducted in Sellers et al. [2007] the reader is encouraged towards the original work.

The simulation composes two models; the first representing the environment, and the second representing an agent (baboon). The environment model is based on field data collected from the baboons' natural habitat, the De Hoop Nature Reserve in South Africa, whilst the agent model treats each baboon as an independent agent, and includes a physiological model of the creature's energetic profile; in this sense the baboon model may be considered to be bio-realistic. The requirements of these two models are given below.

2.3 The environmental model

The environment is approximated by a series of squares representing a 200m \times 200m area each of which may contain at least one, and at most six different habitat types (acacia woodland, burnt acacia woodland, climax fynbos, grassland and vlei). Additionally, each square is characterised as containing none or more of the following features: a water source, a sleeping site or a cliff/other refuge. These squares are arranged in a grid and are identified using Cartesian coordinates.

A habitat type represents plant-life, consumable by an agent, and is described by the following properties, each of which may assume a distinct floating point value for each calendar month:

| Property | Description |
|----------------------------|--|
| Visibility | The visibility of an agent in this habitat |
| | (metres) |
| Productivity | The maximum food availability |
| | $({ m grams/metre}^2)$ |
| Foraging Rate Factor | Consumption by a stationary agent in |
| | one second (grams) |
| Move Foraging Rate Penalty | Deducted from the above if the agent is |
| | moving |
| Recovery Rate | Daily rate of re-growth |
| | $({ m grams/metre^2/day})$ |
| Mean Energy Value | Energetic value of 1 gram (Joules) |

Table 1: Properties of Habitat types used by Sellers, Hill, and Logan [2007]

Where a square contains more than one habitat type the square assumes the arithmetic mean of each of the properties (above) across those habitat types; the assumption is that each habitat type is present in an equal proportion.

For each square two further derived properties are defined; "Risk" is inversely proportional to visibility and is reduced to one tenth of this value where the square contains a sleeping site or a refuge. "Current Resource" represents the current food availability within a square and defaults to the maximum food availability, calculated as the product of the productivity by the area of the square.

At the end of each day the environment is maintained, during which re-growth occurs and the properties of each square are updated. To provide diversity amongst similar squares the risk and current resource of each square are scaled randomly by a factor of 1 ± 10^{-8} .

2.4 The baboon (agent) model

Broadly, each agent is capable of performing a number of activities; these are foraging, drinking, socialising and resting. At a defined interval, measured in seconds, the agent selects an activity to be performed for the duration of that interval. Activity selection is determined by a number of factors, including the time of day, the current environment, and the strength of the agent's desire for that activity. Where an agent wishes to perform an activity which is not feasible in its current location (for example the agent may be thirsty, but water may not be present) it may wish to move towards a location more favourable for that activity; all movement is treated as "travel-foraging". The activity which an agent ultimately engages in is further constrained by its obligation to concede to the will of its peers; this does mean, for example that an agent whose desires are fully satiated, and therefore would naturally prefer to rest, may be coerced into movement by the group as a whole. More specifically, each agent is described by a number of properties, these being:

| Property | Description |
|-------------------------------|--|
| Comfortable Speed | Velocity of movement (metres/second) |
| Rest Power | Energy consumption when at rest (watts) |
| Move Power | Energy consumption whilst moving |
| | (watts) |
| Social Power | Energy consumption whilst socialising |
| | (watts) |
| Forage Power | Energy consumption whilst foraging |
| | (watts) |
| Search Radius | Extent of local knowledge (metres) |
| Daily Nominal Energy Budget | Daily energy intake target (Joules) |
| Daily Nominal Social Time | Daily socialising target (seconds) |
| Nominal Time Between Drinking | Target interval for drinking (seconds) |
| Relative Food Importance | The relative importance of each of these |
| | activities |
| Relative Social Importance | determining the rate at which the |
| Relative Drinking Importance | corresponding desire will increase |
| Foraging Move Threshold | the relative degree by which an |
| | alternative location |
| Socialising Move Threshold | may be more favourable than the present |
| | location; |
| Resting Move Threshold | prompting a desire to move. |
| Rest/Social Risk Threshold | The absolute value above which these |
| | activities may not occur |

Table 2: Properties of agents used by Sellers et al. [2007]

Further, for each activity where a target is defined (foraging, socialising and drinking) the agent maintains two derived properties. A "Score" representative of its current progress towards these targets and a "Desire" representative of the strength of its drive to perform a given activity.

When an agent engages in a targeted activity its score for that activity is increased proportionally to the duration for which the activity is pursued. Similarly, when an agent does not engage in an activity the respective scores are reduced in the same manner. Each desire is determined as ratio of the score, to the daily target for that activity, and is scaled according to the associated relative importance. The goal of each agent is to maintain each desire at the minimum possible level; in this regime a desire is increasingly satiated as it tends towards zero.

Particular activities are notable for their specific traits; as foraging takes place the energy available in the appropriate environment square is depleted. Drinking is treated as an instantaneous activity, which may occur opportunistically whenever water is available, even if the desire to drink is fully satiated; an agent may therefore also engage in any one of the other activities, as is appropriate, for the entire duration of the time interval in which drinking occurs. The model does not place any constraint on the number of agents required to engage in social activity; in effect, it is possible for a single agent to socialise.

As the simulation commences, all agents are placed in a randomly selected sleeping site. Throughout the day the agents may roam freely, although always as a group; they must, however, not stray beyond a point whereby they can return to a (any) sleeping site by nightfall; the daily hours of daylight for each calendar month are provided. Agents are assumed to always know the location of all sleeping site and water sites, but their knowledge of the best locations to forage (determined by available energy) and rest (determined by risk) is constrained to a circular area prescribed by either the search radius property or by the maximum distance they may travel before nightfall, which ever is the lower.

If all desires are satiated the agent will wish to rest, otherwise, within the constraints detailed above, the agent selects its activity at random from a distribution weighted in proportion to the strength of each desire.

3 Design

The software requirements were considered in conjunction with the objectives creating a design which addresses both the generic requirements of the objectives and the specific requirements of replicating the work of Sellers et al. [2007]. The final design, in simplified form, is illustrated in figure 1 and detailed below.

3.1 The Generic Model

Aspects understood to represent common concepts within software of this nature have been abstracted into a series of "manager" classes giving rise to the following:

TimeManager – This class provides various methods related to the passage of time throughout the day/night, and throughout the calendar year.

GeometryManager - This class maintains two, overlaid, two dimensional planes; a coarse-grained grid, suitable for describing an environment and a continuous space suitable for occupation by agents. Each plane is addressed through the Cartesian co-ordinate system; the class provides methods suitable for translating to/from grid and continuous planes.

PopulationManager – This class provides a container for the composition of groups (themselves a composition of agents) into a population; permitting either all agents, any group of agents, or an individual agent to be addressed.

EnvironmentManager - This class provides a container for the composition of environmental features and also provides methods for searching within such features.

Logger – An abstract class for generating structured, text based log files.

3.2 The Group Decision Making Model

The group decision making process is provided for by a generic architecture which can be tailored to the needs of a specific scenario. This is achieved through the following classes:

Group – This class provides a container composing agents into a cohesive, conspecific group; the class is generic and reusable.

Policy – This class provides the logic necessary to formulate a consensus group decision. The quorum threshold is defined by the "proportionToWinVote" attribute. A Policy may be reused in association with appropriate Decision and Activity classes.

Decision – This class represents the data structure used to describe a decision. The class is intended to represent either an individual or the group decision.

Activity – This enumeration describes the set of items about which a decision may be made.

3.3 The Agent & Habitat Models

The following classes are provided to support the agent (baboon) and habitat models necessary to replicate the work of Sellers et al. [2007]:

Agent – This class provides the baboon model, and implements methods required by the group decision making model.

FloraKind – This class represents the data structure used to describe the various plant-life required in the simulation.

HabitatArea – This class composes FloraKinds, and other descriptive properties to define a region of the environment. In this simulation, each square of the grid based geometry is represented by a separate instance of HabitatArea.

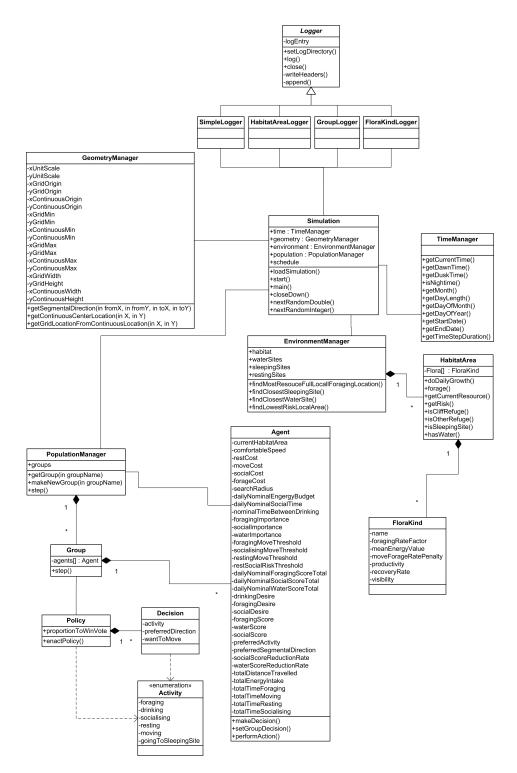


Figure 1: UML Diagram of the software design (simplified) $\frac{8}{8}$

References

WI Sellers, RA Hill, and BS Logan. An agent-based model of group decision making in baboons. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 362(1485):1699–1710, 2007. doi: doi:10.1098/rstb.2007.2064.