

Implementing a "Fast and Frugal" Cognitive Model within a Computational Social Simulation

William G. Kennedy (wkennedy@gmu.edu)

Krasnow Institute for Advanced Study, George Mason University, PO Box 4400
Fairfax, VA 22030 USA

Jeffrey K. Bassett (jbassett@cs.gmu.edu)

Volgenau School of Information Technology, George Mason University, PO Box 4400
Fairfax, VA 2030 USA

Abstract

Large-scale social simulations require a cognitively credible but computationally efficient cognitive architecture to support simulations with thousands to tens of thousands agents. In previous work developing and experimenting with a large-scale social simulation, we successfully employed an ad hoc cognitive model, a formula-based decision function. We explain why we now believe that a "fast and frugal" cognitive architecture to be superior based on its indistinguishable computational efficiency and much better cognitive plausibility.

Keywords: social simulation; cognitive architectures; agent-based modeling.

Introduction

When simulating thousands to tens of thousands of people making individual decisions as part of a large-scale social simulation, their decision-making process could be modeled in many different ways. A challenge is to maximize cognitive credibility of the process while minimizing the computational intensity involved.

What we mean by "cognitive credibility" is face validity of the observable behavior of agents, i.e., it looks right, and that their decision-making processes are recognized by the cognitive science community. The ideal, of course, would be to match human decision-making data, but such data is very rarely available in social science. The "computational intensity" to be minimized is the relative amount of computing resources required for the modeling of human decision-making compared to the total amount of computing resources involved in the overall social simulation. The latter may be quantified and measured, but the former is not easily quantified.

Social simulations have begun relying on agent-based models (ABM) as a useful approach for modeling many individual agents, but a common tool for modeling human decision-making has not emerged. With social simulations involving so many agents, research quality, cognitive architectures such as ACT-R (Anderson 2007; Anderson et al 2004), Soar (Laird 2008), and CLARION (Sun 2006) are normally considered too computationally intensive and too flexible for modeling the numbers of agents in social simulations. Practitioners therefore typically use ad hoc approaches to modeling human cognition.

In this report, we discuss replacing our original ad hoc mathematical approach to modeling agent decision-making with a "fast and frugal" cognitive architecture approach (Gigerenzer 2007).

Our Original Approach

For an ongoing project to study the origins and continuation of conflict in east Africa, we needed to model the activities and interactions of the pastoralists (herders) and agriculturists (farmers) of the area. The simulation represents the 150km by 150km area as a 2D set of parcels at the level of one square kilometer. Agents represent family units including their associated herds and the time step is one day. Our experiments typically involve runs over 5 to 100 years.

We used an ad hoc mathematical approach to modeling the cognitive behavior of each agent. Farmers were modeled as stationary agents who acted only to protect their family's resources and the rest of the farm's operations were driven primarily by the weather. Herders were modeled as mobile agents who needed to decide on a daily basis where to move their herd and family. In our published work (Kennedy et al 2010a, Hailegiorgis 2010; Kennedy et al 2010b), we used a polynomial that evaluated four situational factors to determine which parcel of the nearby land was the best for the herder to move to. The factors were the herd's water needs, its vegetation needs, the potential for getting into conflict with another agent, and the distance to the subject parcel. (Conflict in our model was the result of two agents occupying the same parcel of land.) To determine which of the possible locations to move to, every location within the agent's field of view (10km) was evaluated using a multivariate polynomial with terms for the four factors, each with arbitrary scaling coefficients as shown in Eq. 1.

$$Q = a*T + b*H + c*C + d*(1/D) \quad (\text{Eq. 1})$$

Where: Q is quality of the subject parcel

T is thirst need

H is hunger need

C is whether the parcel is occupied and would result in conflict (binary)

D is distance from parcel to water supply

a, b, c, & d are scaling coefficients

The model of the herders' decision making using this formula-based behavior was considered reasonable by the anthropologists on our team and peer reviewers of the papers submitted, accepted, and presented in three venues. The computational intensity was reasonable on our systems even with thousands of agents, but the approach does not seem very cognitively plausible. And we do not want the community to suffer from the same plausibility challenges the field of economics has. Therefore, we considered more cognitively plausible cognitive architectures used in cognitive science.

The "Fast and Frugal" Approach

In 2007, Gerd Gigerenzer published the book, "Gut Feelings", discussing the research behind the concept of intuitive reasoning (Gigerenzer 2007). Although the book is primarily aimed at unconscious cognition, he also describes a "fast and frugal" heuristic for representing human decision-making. He discussed its application in improving emergency room decisions, explaining judges' bail decisions, and implementing Herbert Simon's bounded rationality concept (Simon 1966). The "fast and frugal" heuristic approach continues to be the topic of current discussions (e.g., Marewski et al 2010; Hilbig 2010).

This approach to cognition is a little different from the traditional rule-based approach. This approach considers the factors affecting a decision sequentially in the order of their importance rather than in

parallel. Each rule is also focused a little differently from the standard approach. Commonly, rules identify all the conditions in the environment necessary and sufficient to determine which possible action to take. Here, the approach is to ask whether the agent has sufficient information to act. If so, act. If not, add consideration of the next prioritized factor. In other words, rather than saying humans weigh several factors simultaneously to make a decision, the reported research argues that humans rank order factors and consider the factors sequentially until they have enough information to act (explaining why car buying decisions could hinge on the number of cup holders, all other factors being judged balanced and, therefore, indistinguishable).

In our original model, our agents considered four factors at once and we adjusted the coefficients to tune the behavior to be appropriate. The new "fast and frugal" approach, those factors are ordered by importance. The ordered questions of the rules in new model are:

1. Is the situation dire?
2. Is there conflict near by?
3. Does the herd need watering?

We handle the original fourth factor, the distance to the potential location, as part of the action side of the rule. The "move" command optimizes movement based on the current situation, either moving away from conflict, toward water, or toward vegetation.

This new approach considers the same four factors as was done previously, but prioritizes the movement conditions explicitly rather than through the weights. Although the previous four-factor formula was more flexible, its wide range of possible actions based on the weights was not cognitively plausible. The new approach is less flexible, but the flexibility has been traded for cognitive feasibility.

We have implemented the "fast and frugal" approach replicating the decision-making process used in our previous published experiments. Both approaches were implemented using the MASON system (Luke et al. 2005), a Java-based simulation environment used in the previous studies.

Implementation

In our specific domain, all herders have several high-level concerns that need to be monitored and balanced against one another. These fall into a natural hierarchy with situations that are most likely to harm the herd being considered first, while the most mundane issues are considered last.

High-Level Rule Decision High-level decisions are implanted using the "fast and frugal" approach. Figure 1 gives the basic decision structure including actions used implementing the ordered questions developed to model the herder's decision-making. The answers to these questions trigger certain low-level motor behaviors.

Low-Level Motor Behaviors We have implemented a small number of low-level motor behaviors as actions to be performed under specific circumstances. While we refer to these as low level, they can actually be complicated and involve some (automatic) decision making as well, although these usually involve spatial evaluations.

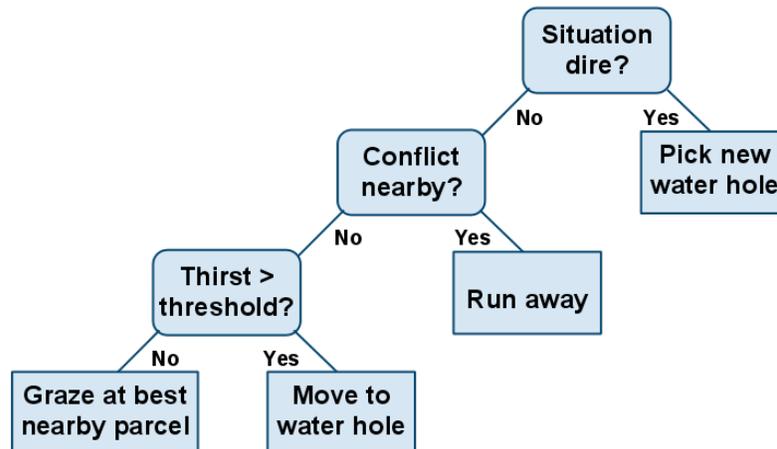


Figure 1: Fast and Frugal herder behavior.

Pick New Watering hole Each agent centers its activities around a local watering hole. When necessary, agents decide to abandon their current base watering hole and move on to another. Our agents keep a list of the five most recently visited watering holes to avoid oscillating between two. They will not return to any of these until they have been forgotten and removed from the list. When choosing a new watering hole, the closest hole to the agent that is not in the list of recent watering holes is the one chosen.

Run Away Agents move away from the center of the conflict in their view. The center is calculated by averaging the locations of all the conflicts that occurred in the previous time-step that were within ten parcels of the agent. A seven parcel long vector is calculated in the direction away from the center of conflict, and the agent moves to the parcel that contains the most vegetation that is within a radius of three parcels from the end of that vector.

Move to Watering hole If the watering hole is within a single day's movement range from the agent (ten parcels), the agent simply moves onto the watering hole. Otherwise, a seven parcel long vector is calculated in the direction of the watering hole, and the agent moves to the parcel containing the most vegetation that is within a radius of three parcels from the end of that vector.

Graze The agent moves to a nearby (within ten parcels) unoccupied parcel in order to graze the herd. When considering which parcel to move to, a balance is struck between the amount of vegetation available, and the distance to be traveled. When searching for high vegetation parcels, ones that are further are discounted based on their distance from the agent.

Comparison

We have compared the two approaches to the modeling cognition of thousands to tens of thousands of agents. We first discuss computational intensity similarities. The real difference between the two approaches appears to be primarily in the day-to-day behavior of the agents while grazing near a watering hole or migrating between watering holes. After comparing their observable behavior, we discuss their relative cognitive plausibility.

Computational Performance Differences

We attempted to compare some of the performance (computational intensity) differences between the two cognitive approaches, but there were a number of difficulties in getting reliable numbers to compare. MASON can display the frame rate (number of time-steps per second) at which the simulation is currently running. When the graphics are displayed, the two approaches seem to be very similar. Runs were performed with a several different population sizes, and the two cognitive methods always showed very consistent frame rates. However, most of the runtime is consumed in drawing the graphics. When graphics are turned off, the simulation runs much faster, but the frame rates then have such a high variances that it becomes very difficult to get any sort of accurate comparison of the methods.

Day-to-Day Foraging Behavior Differences

Figures 2 and 3 compare formula and rule based behaviors of three agents on their fourth day of foraging around the same watering hole. The green background is an indication of the vegetation available in the approximately 10 by 10 km area. The blue square is the parcel with a watering hole. The solid circles show the current locations of the agents, while hollow circles show previous locations of agents, connected by lines indicating the paths that were taken.

Figure 2 shows an image of the grazing behavior of the equation-based behavior model. As can be seen, most of the agents tend to huddle around the watering hole, with occasional brief forays out to more distant parcels in order to find high quality food sources. The forays seem to be in general, and as shown here, a trip to the closest best area and then a slow return to the watering hole as the evaluation balances need for vegetation and water. The

Figure 3 shows the grazing behavior using the "fast and frugal" rule-based approach at the same watering hole at the same day. In this case, the agents tend to take much longer forays outlying areas, generally following a path through the further richer grasslands, and returning to the waterhole on day five.

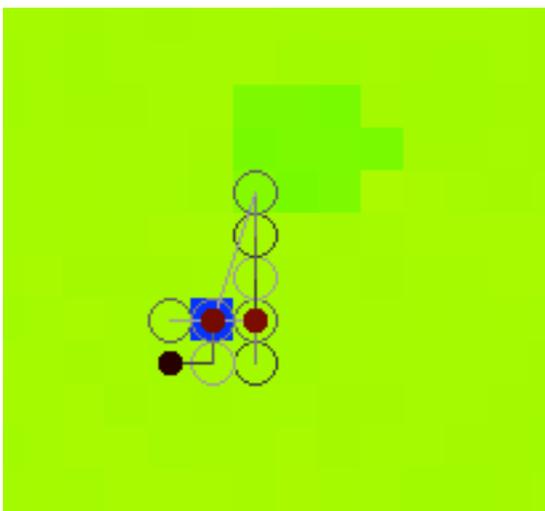


Figure 2: Formula-based Grazing Behavior

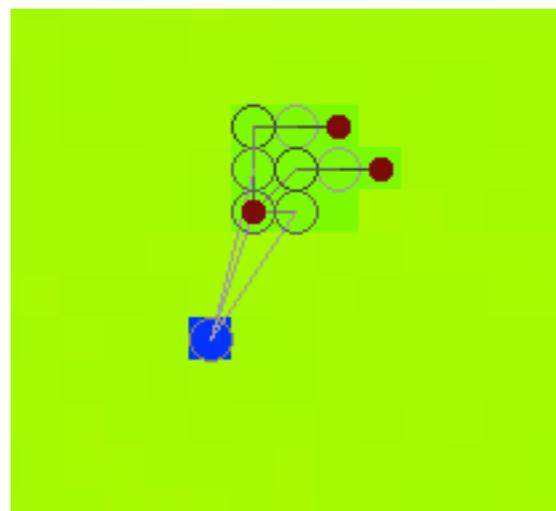


Figure 3: Rule-based Grazing Behavior.

Migration Behavior Migration behavior occurs when a herder agent changes base watering holes due to the watering hole running dry or some other reason. This can often involve many herding family units moving at nearly the same time toward a new watering hole. Each agent has a memory of their five most recent watering holes and moves to the closest water source that was not among those five. Depending on the spatial distribution of the watering holes, this process can produce cyclic behavior or a generally linear behavior. The cyclic pattern arises when six or more watering holes located near each other. The longer distance behavior arises if the sixth watering hole is not near the first watering hole, for example along a river where the previously visited watering holes are in a line with the last far from the first. Both behaviors occur among the pastoralists of the region and both the formula-based and the rule-based approaches have this macro level behavior.

Where the behaviors differ is on the migration paths followed. Figure 4 shows the tracks and current location of several agents migrating on the same path using the formula-based approach. Figure 5 shows the traces of agents using the rule-based approach. In both figures, several agents (solid red circles) are shown moving between watering holes (small, dark blue squares). The destination water hole is off screen to the south. Again, the image includes empty circles showing the location of agents during the previous five days indicating the herder's tracks.

Notice that the formula-based movement in Figure 4 is following two paths and at two speeds. Some agents move one parcel per day and some move 10 per day. The difference is based on whether the agent is more hungry or more thirsty. In both cases, the base watering hole has been changed causing the basic direction of the movement, i.e., toward the next watering hole. If the agent is more thirsty than hungry, it will move at the maximum speed to get to the next watering hole. If the agent is more hungry than thirsty, then it moves only one parcel per day. When the movement of two agents would result in both agents in the same parcel, the conflict factor of the evaluation formula-based approach causes the second agent to avoid the first. This results in path variations.

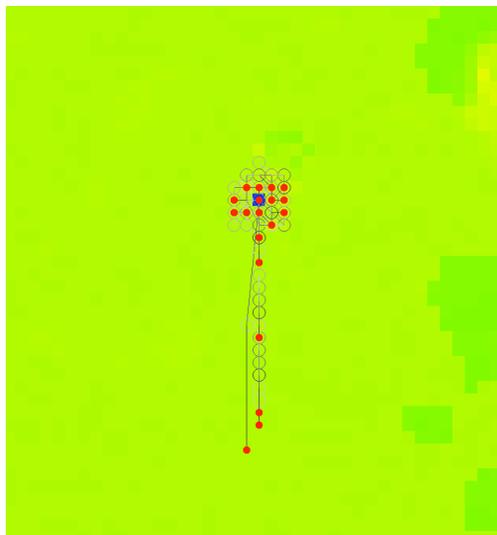


Figure 4: Formula-based Migration Behavior

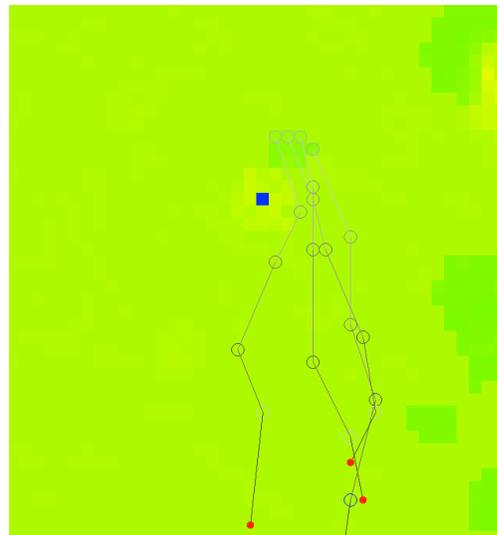


Figure 5: Rule-based Migration Behavior

Figure 5 shows rule-based migrations. In this approach, the agents also move avoiding conflict, but are more separated (shown here) or form irregularly shaped groups with agents spread out along the general direction rather than straight lines (not shown).

Discussion

This paper discusses two approaches to modeling herder behavior. Our original approach was based on combining four factors: need for water, need for vegetation, avoidance of conflict, and minimizing travel costs. We implemented that approach through a four-factor mathematical evaluation of each possible next place to move the herd within the herders' vision or daily movement range. That approach produced workable macro-level behavior but we had concerns for its cognitive credibility and its micro-level behavior. So, we considered other more cognitively plausible approaches. The "fast and frugal" rule-based approach offered a more cognitively credible decision-making process.

The comparison of the two approaches was based on two questions. The questions were whether this approach could be implemented in a computationally cost-effective manner and whether the resulting behavior would have better micro-level, observable behavior. Our implementation of the "fast and frugal" approach proved to be straight forward and apparently without a noticeable difference in computational efficiency. We believe it to be certainly less than the computational cost of using ACT-R, Soar, CLARION, or any research quality cognitive architecture. However, we have yet to perform quantitative comparisons.

The second question was whether the resulting behavior would appear more credible than the formula-based behavior. Traces of the rule-based approach appear more reasonable than the formula-based approach. The day-to-day grazing around a watering hole using the rule-based approach seems more reasonable by not having the herds moving at returning to the water every other day. During migrations, the rule-based approach again appeared far more reasonable because it did not have the linear formations of agents but had much more natural grouping of herders during the migrations.

Finally, as an additional benefit, the rule-based approach is far more transparent in its operation than the formula-based approach. Although the rule-based approach is not as flexible as the formula-based approach with its four scaling parameters, the restriction in the expressivity of decision criteria is a strength, not a weakness, of the representational approach.

What we will not know until the previously published experiments are replicated is whether there are any effects in the overall simulation results due to the change in cognitive modeling. That remains to be evaluated as future work. We also did not explore modifying the parameters of the formula-based approach to improve the face-validity of the resulting behavior. That, too, is future work.

Conclusions

We have concluded that with apparently little cost, we could greatly improve the cognitive credibility of our computational social simulation by replacing a formula-based decision-model for each agent with a "fast and frugal" rule-based approach to the cognitive modeling of agent decision-making. Our new approach provides more cognitively credible day-to-day movement decisions and migration behavior without distracting and unreasonable modeling artifacts.

Acknowledgments

This work was supported by the Center for Social Complexity at George Mason University and by the Office of Naval Research (ONR) under a Multidisciplinary University Research Initiative (MURI) grant no. N00014-08-1-0921. The authors would like to acknowledge input from the Mason-HRAF Joint Project on Eastern Africa (MURI Team), particularly Tim Gulden for his original concept and implementation of the formula-based approach to agent decision-making. The opinions, findings, and conclusions or recommendations expressed in this work are those of the authors and do not necessarily reflect the views of the sponsors.

References

- Anderson, J.R. (2007). *How Can the Human Mind Occur in the Physical Universe?* Oxford University Press.
- Anderson, J.R., Bothell, D., Byrne, M.D., Douglas, S., Lebiere, C., and Qin, Y. (2004). An integrated theory of mind. *Psychological Review* 111(4), pp. 1036-1060.
- Gigerenzer, G. (2007). *Gut Feelings: The Intelligence of the Unconscious*. New York: Penguin Books.
- Hailegiorgis, A.B., Kennedy, W.G., Balan, G.C., Bassett, J.K., and Gulden, T. (2010) An Agent Based Model of Climate Change and Conflict among Pastoralists in East Africa. *Proceedings of the International Congress on Environmental Modeling and Software*. Ottawa, Canada.
- Hilbig, B.E. (2010) Reconsidering "evidence" for fast-and-frugal heuristics. *Psychonomic Bulletin & Review*, 17, 923-930.
- Kennedy, W.G., Hailegiorgis, A.B., Bassett, J.K., Coletti, M., Balan, G.C., Rouleau, M., & Gulden, T. (2010a) An Agent-Based Model of Conflict in East Africa and the Effects of Watering Holes. *Proceedings of the Behavior Representation in Modeling and Simulation Conference 2010 (BRiMS)*, pp 274-281. Charleston, SC.
- Kennedy, W.G., Hailegiorgis, A.B., Rouleau, M., Bassett, J.K., Coletti, M., Balan, G.C., Gulden, T., and Cioffi-Revilla, C. (2010b) MASON HerderLand: Modeling the Origins of Conflict in East Africa. *Proceedings of the First Annual Conference of the Computational Social Science Society*, Tempe, AZ.
- Laird, J. (2008) Extending the Soar Cognitive Architecture. *Proceedings of the Artificial General Intelligence Conference*. Memphis, TN.
- Luke S., Cioffi-Revilla, C., Sullivan, K., and Balan, B. (2005) MASON: A Multiagent Simulation Environment. *Simulation*, 81(7), 517-527.
- Marewski, J.N., Gaissmaier, W., Gigerenzer, G. (2010) Good judgments do not require complex cognition. *Cognitive Processing*, 11, 103-121.
- Simon, H. A. (1966) *The Sciences of the Artificial*. MIT Press: Cambridge, MA.
- Sun, R. (2006). The CLARION cognitive architecture: Extending cognitive modeling to social simulation. In: Ron Sun (ed.), *Cognition and Multi-Agent Interaction*. Cambridge University Press, New York.