I Like What I Know: How Recognition-Based Decisions Can Structure the Environment

Peter M. Todd[1*] and Simon Kirby[2]

[1] Center for Adaptive Behavior and Cognition
Max Planck Institute for Human Development
Lentzeallee 94, 14195 Berlin, Germany
ptodd@mpib-berlin.mpg.de
http://www-abc.mpib-berlin.mpg.de/users/ptodd

[2] Language Evolution and Computation Research Unit
Department of Theoretical and Applied Linguistics
University of Edinburgh, Scotland, UK
simon@ling.ed.ac.uk
http://www.ling.ed.ac.uk/~simon/

Abstract. Cognitive mechanisms are shaped by evolution to match their environments. But through their use, these mechanisms exert a shaping force on their surroundings as well. Here we explore this cognition-environment interaction by looking at how using a very simple cognitive mechanism, the recognition heuristic for making choices, can result in some objects in the world being much more often recognized and "talked about" than others. An agent-based simulation is used to show what behavioral factors affect the emergence of this environmental structure.

1 Introduction: Cognition Creating Environment Structure

Some things are more famous—more recognized, more talked about—than others. This unsurprising fact about the world can be useful when making decisions: If we have the choice between a recognized and an unrecognized option, we can make a rapid (and often adaptive) decision based on recognition alone. What is more interesting is that some things are *much* more famous, or successful, or noteworthy, than others. The few most productive scientists are cited far more than the majority of their colleagues, and the few most successful bands sell far more records than their competitors. If we plot such quantities as number of citations, or records sold, or city population, or Olympic gold medals won, against the rank-ordered list of authors, bands, cities, or countries, we see a characteristic J-shaped curve (where the "J" is on its side), indicating very few objects with very high values and most objects with rather low values [5]. This striking aspect of environment structure appears to emerge commonly in domains that

* We are indebted to Martin Dieringer for his work programming this model.

J. Kelemen and P. Sosík (Eds.): ECAL 2001, LNAI 2159, pp. 166–175, 2001.

people think about, talk about, and make decisions about. The question we want to explore here is, how does this structure arise?

The obvious answer is that there is some underlying structure in the world that our decision-making and knowledge simply tracks. Some researchers just produce more important work, and some cities are located in more resource-rich locations. But these pre-existing differences do not explain the entirety of this phenomenon in all domains—the action of intelligent agents making choices and communicating information about their world can also shape the distribution of knowledge and choices in the world. This is a specific instance of the more general question of how cognition and environment act to affect each other. Cognitive mechanisms are shaped by their environments, both through evolutionary selection across generations and through learning and development within lifetimes. But by making decisions that guide actions which in turn alter the surrounding world, cognitive mechanisms can also shape their environments in turn. This mutual shaping interaction between cognitive structure and environment structure can result in coevolution between the two over extended periods of time. What will the dynamics of such feedback loops be, and what patterns of structure in both cognition and environment will emerge?

We can observe the outcomes of these interactions in the world around us, for instance in the J-shaped distributions described above. But to understand *how* such interactions work to produce environment structure, we would like to start out with systems that are initially unstructured, particularly lacking pre-existing differences in importance between objects, and then observe how structure can emerge over time. Thus, to get a handle on the questions of environment-behavior coevolution, we can construct simple simulation models of decision-making agents interacting over time with each other and the rest of their environment. We begin here with a model incorporating perhaps the simplest decision heuristic, the *recognition heuristic*. With this model we can study how the use of a simple heuristic shapes the distribution of choices in (and knowledge about) an environment full of different options. In particular, we are interested in whether and in what circumstances the recognition heuristic can itself create a clumpy environment, in which some options are recognized, and chosen, much more often than others (e.g., following a J-shaped distribution). Such environment structure can in turn make the recognition heuristic ecologically rational, that is, able to make beneficial choices. In the rest of this paper, we indicate past work on how cognitive mechanisms can shape their environments and introduce the recognition heuristic as a particularly simple cognitive mechanism to study, describe our model of recognition-based decision making in a group of moving, communicating agents, and discuss some of our initial results on the factors affecting the emergence of environment structure via simple cognition.

2 Background: Behaviors, Environments, and Recognition

Behavior by definition affects the environment (including the behaving organism) in some way, usually in a small way, but sometimes in a manner that has lasting consequences for further behavior. Important theories of how organisms can affect their own selective environments (including their cultures) have been developed, notably dual inheritance (gene/culture) ([2]) and niche construction ([10]) models. These models however typically rely on very simplified behavioral strategies analyzed at a population level. Because agent-based simulation enables the study of population-level effects of decision-making mechanisms operating at an individual level, it allows more explicit models and analyses. This approach has proven particularly fruitful in studying the evolution of language—by simulating populations of agents communicating with an evolving syntactic mechanism, one can gain insights into the way that this mechanism shapes, and is shaped by, the coevolving language being used in the population ([7], [8], [9]).

But language mechanisms are somewhat complex, and we wanted to start here with a simpler cognitive mechanism to explore environment construction. The simplest kind of choice—numerically, at least—is to select one option from two possibilities, according to some criterion on which the two can be compared. The simplest form of this choice occurs when the only information available is whether or not each possibility has ever been encountered before. In this case, the decision maker can do little other than to rely on his or her own partial ignorance, choosing recognized options over unrecognized ones. This kind of "ignorance-based reasoning" is embodied in the recognition heuristic [4]: When choosing between two objects (on some criterion), if one is recognized and the other is not, then select the former. For instance, when deciding which of two cities is larger, say Palo Alto or Weed, California, the recognition heuristic would lead (most of) us to choose Palo Alto (which is the correct choice).

Using the recognition heuristic will be adaptive, yielding good decisions more often than would random choice, in those environments in which exposure to different possibilities is positively correlated with their ranking along the decision criterion being used. Such correlations are likely to be present in species with social information exchange where important environmental objects are communicated and unimportant ones are ignored. For instance, because we talk and hear about large cities more often than small cities (because more happens in them, more people come from them, etc.), using recognition to decide which of two cities is larger will often yield the correct answer (in those cases where one city is recognized and the other is not). Use of the recognition heuristic appears widespread: For instance, Norway rats prefer to eat foods they recognize as having smelled on the breath of another rat over unrecognized foods [3], even when more useful information is available to guide their choice (such as that the rat whose breath they smelled was ill from eating that food—for a model to explain this ignorance of further information, see [11]). Humans make decisions in accordance with the recognition heuristic both in the lab (e.g., choosing recognized cities as larger than unrecognized ones—see [4]) and outside (e.g., choosing stocks to invest in based on recognition of the company, a strategy which can be adaptive in a financial sense as well—see [1]).

3 A General Model of Agents Deciding with Recognition

In the most general setting for a model of agents acting in a world and making decisions on the basis of recognition, we can consider a space of options and a population of agents, each of whom can choose one option at a time. The options are identified just by a unique name, and so the agents can only use name recognition (or lack thereof) to choose between them. We can think of each agent as making a series of choices between two named options presented simultaneously. The environment-structure outcome we are interested in is how the population of agent-choices is distributed over the set of available options at some point in time—more specifically, are choices made by agents evenly distributed over options, or are they clumped or J-distributed?

Even with this very simple setup, we have some variations to consider. We can affect the decisions made in at least four main ways:

– 1. How recognition is determined and remembered over time
– 2. Whether and how recognition information is communicated between agents
– 3. How options are selected to be presented as choice-pairs to an agent
– 4. How choices are made between pairs of presented options

Before we discuss how specific instantiations of these factors affect emergent environment structure, though, we need to make the simulation scenario itself a bit more specific.

4 A More Specific Model: Agents Affecting “City” Sizes

To make a more concrete and intuitively “recognizable” model, we start with a simulation that can be related to the city-size choice task mentioned earlier. Now the options are essentially place-names (locations) in a grid, and agents make choices between these locations and then move to the chosen place. More specifically, agents live on a grid-world (multiple agents can inhabit each location) and have a (recognition) memory of places they have been and have heard about from other agents. Time proceeds in steps, and at each step an agent has a chance of “talking” with other agents about some of the locations it recognizes (and hearing from them about what they recognize in turn). Every so often, each agent is presented with a choice between two locations selected from the grid-world, and the agent must choose one of these locations on the basis of its recognition knowledge, using the recognition heuristic. (Thus if it recognizes one location and not the other, it will choose the former; otherwise the choice will be made at random.) The agent then moves to the chosen location. Even less often, the agent has a chance of forgetting about some of the locations it recognizes. (This can equivalently be thought of as the old agent dying and a new child agent being introduced who first learns some reduced set of what its parent knew.) By instantiating this all in a two-dimensional space, we can monitor how agents are distributed over this space over time, and look for actual population clusters. Given this overview, we can now turn to a consideration of the specific instances possible for the four factors listed in the previous section.

4.1 How Recognition Is Determined and Remembered

An agent could recognize a location after visiting that location, hearing about it from another agent, or seeing it as one of the options presented in its occasional choice-pairs. Here we always incorporate the first source of recognition knowledge, and turn on or off communication to test its effect; we leave out the third indirect form of experience (though even such minimal exposure can have a strong impact in humans, causing what has been called the "overnight fame" effect—see [6]).

It became clear in early runs of our system that individual recognition knowledge must also disappear somehow, or else each agent ends up recognizing nearly all locations and can no longer use the recognition heuristic to choose between them. In this model, recognition memory is reduced at regular (user-specified) intervals as mentioned earlier. Different sets of locations to be forgotten can be selected: a fixed percentage of recognized locations chosen at random; or locations that have not been experienced recently; or the least-often experienced locations. The latter two can also be combined in a decaying memory trace, which we will not elaborate on here—instead, we will focus on the first two forms of forgetting.

4.2 How Recognition Knowledge Is Communicated between Agents

It is possible that agents acting independently on the basis of their own individual experience, choosing to go to locations that they personally recognize by having visited before, could suffice to create emergent environment structure (though we have not seen this happen yet). But it is more likely, and more realistic, that communication between agents will enhance any clustering of choices in the space of options (here, locations)—the "social computation" enabled by a communicating population of simple decision-making agents should lead to greater environmental impact (as has been found in simulations where the interactions of many generations of simple language learners enable syntax to emerge—see [8]). Thus we allow agents to tell each other about locations that they recognize. We must specify who can talk to whom, how often, and about what. In these studies we let agents occupying the same location talk to each other; we set how many other agents a given agent can listen to during each time step, and how many locations each of those other agents can mention to the given agent (who then enters those locations into its memory as "recognized via communication"). Each agent can mention its allotment of locations in the same orders that are possible for forgetting (but reversed): random, newest-experienced first, or most-experienced first. (Note that with communication included this model bears a resemblance to Boyd & Richerson's model of conformist transmission [2], but here knowledge, rather than choices, is communicated, and individual decision making is much more the focus.)

4.3 How Options Are Selected for Presentation as Choice-Pairs

Humans and other animals face choices between recognized and/or unrecognized options all the time—but how are the options we simultaneously face actually determined? At times they may be randomly encountered, or generated in a more systematic way by the environment; in other situations, other individuals may present us with the things to choose between. Similarly, the agents in our model can encounter their two options (at user-specified time intervals) in a few ways. The two locations can be selected at random from the entire grid-world, or in a locally-biased fashion by specifying a distance-factor that determines how rapidly the chance of seeing a given location as an option falls off with distance between the deciding agent and that location (according to $probability = 1/distance^{distancefactor}$). Or the two locations can be chosen from those recognized by other agents in the current location, instantiating another form of local communication, or in proportion to how many agents are at each location in the grid-world, representing "advertising" being sent from cities near and far saying "move here!" (with more advertising coming from larger cities).

4.4 How Choices Are Made between Presented Options

Once a pair of options is presented to an agent, it must use its recognition knowledge to make a choice. This choice can be made using the recognition heuristic as discussed earlier, or it could be made randomly, to provide a basis of comparison with the impact of the recognition heuristic. (Additionally, the "anti-recognition heuristic" could be used, always picking unrecognized over recognized options, which would result in greater individual exploration rather than conformity. This is certainly useful to have in societies and other systems facing an explore/exploit tradeoff, but which we will not explore further here—see [2], [13].) Choices could also be made on the basis of recency or amount of recognition (as for forgetting and communication). Finally, we can allow agents to refuse to make a choice when there is no clear basis for selecting one option over another, for instance, when both are recognized or both unrecognized; such refusal would allow their rate of use of the recognition heuristic to increase, which could enhance its power to shape the environment.

5 Initial Results

Given some combination of the factors just described, we can specify an initial state of the world (how many agents, where they are located, how old they are, what they know if anything at the start—here, we use random positions for some number of agents with random ages and no starting knowledge) and then watch over many timesteps how the environment (distribution of agents and knowledge) changes. (The system is implemented in Java and runs at a good clip in real time for populations of several hundred agents.) There are a few important values to monitor as the simulation progresses: First, we look at the distribution of agents per location, graphed with the locations along the x-axis ordered from left to right by "population size" so that we can see when this distribution becomes more or less J-shaped. Second, we look at the correlation between the recognition for each location (that is, the number of agents that recognize it) and its population size (the number of agents on it)—the higher this correlation is, the more effective and valid is recognition knowledge for choosing the more-populous location to move to (though keep in mind that there is no selective pressure or other advantage for agents to go to more or less crowded locations). And third, we keep track of variables regarding the agents and their knowledge: the mean number of locations recognized, the mean number of decisions made, the number made when recognition distinguished between the choices, and the number of choices to move to the larger of the pair of locations.

We have not yet explored all combinations of the behavioral factors available in this model, but the initial indication is that substantial clustering of knowledge and choices is not so easy to come by. Consider an illustrative example run, in which we used a 10x10 world, 200 agents, a 200-timestep "memoryspan" (meaning that forgetting occurred for each individual every 200 timesteps, plus or minus a 20-timestep standard deviation), a 10-timestep interval between choices and moves (with standard deviation of 4), randomly selected locations as choice options, 50% chance of communicating on each timestep, and 1 location heard from 1 agent every time this communication takes place. The results of this run of 1000 timesteps are shown in Fig. 1. Here we see that the population at this moment is essentially J-distributed, with very few locations with the largest number of agents (10) and most locations with only a few agents (3 or fewer), but this changes from timestep to timestep and can often appear less skewed. The population vs. recognition correlation over time begins low (around 0) in the initial random scattering of individuals, and then as they move around this correlation rises above 0.4 by timestep 300 before falling and rising again (the reason for these apparent cycles is not yet clear). The average number of locations recognized by each individual asymptotes around 38, with 30 of these recognized from hearing about them from other agents, and only 8 coming from "personal" experience (just the proportions one would expect from the rates of choice/movement and communication). Finally, the average number of decisions made so far by each individual (counted since the last time they hit their memoryspan and forgot past experiences) levels off around 9, with half of those made using the recognition heuristic (as would be expected when nearly half of the locations are recognized by each individual), and half again of those being choices to move to the larger of the two locations.

It is this last value—the proportion of recognition-based decisions that result in movement to the higher-population location—that is crucial. Here, with half of those decisions going each way, we see that recognition knowledge has no effect on movement to more populous locations, meaning that the agents will not be able to cluster strongly. This is because the locations that agents recognize are mostly ones they have heard about from others, and hence are usually locations that the others also heard about, meaning that the last time anyone was actually personally *in* the location could have been a long time ago—and any cluster of individuals that were there then could have long since disappeared. Thus, for clusters to build, the information communicated between agents has to be made more timely. There are two main ways to do this: change what the agents remember (and so can talk about at all), and change which locations they choose to talk about at any instant.

When we change both of these factors, making agents preferentially forget the oldest-experienced locations and communicate the most-recently-experienced ones, we do indeed see more environment structure emerge. Now the population vs. recognition correlation averages around .4 (peaking at .6), and a few locations end up with 12 or 13 agents in them. The average recognition per agent falls to 11 locations, because most of what they hear about now they have already heard before. As a consequence, the number of recognition-based decisions also falls, to 2.5 out of 9, but the proportion of these choices to larger-population locations rises slightly, to 60%. This is enough to produce a more-structured environment, but the effect is not as large as might be possible, and our search continues for what factors may be able to increase the environment structure further.

6 Conclusions and Extensions to the Model

In our simple model of the impact of recognition-based decision making on the environmental distribution of choices and agents, it appears that introducing time-dependencies into the memory and communication of agents is sufficient (and perhaps necessary) for allowing population clusters to emerge. This fits in well with the ecological rationality perspective that cognitive structure and environment structure should match to produce adaptive behavior [12]—given that the environment changes over time, agents' knowledge should also change so that old possibly outdated information is no longer relied on.

So far, our search for the factors that most enable environment construction has been carried out by hand, manipulating the factors ourselves and running new simulations. But this process could potentially be sped up by enabling the coevolution of the recognition heuristic and the structure of the environment, allowing agents to inherit the exact strategies they will use—both their decision strategies (e.g., recognition-based or anti-recognition-based) and their communication strategies (e.g., tell everyone everything you've seen, or the most recent locations, or none). One question of interest in such a coevolutionary setting is whether there is selection to "the edge of chaos" (or at least to a cognitively advantageous sweet-spot) such that agents end up recognizing around half of

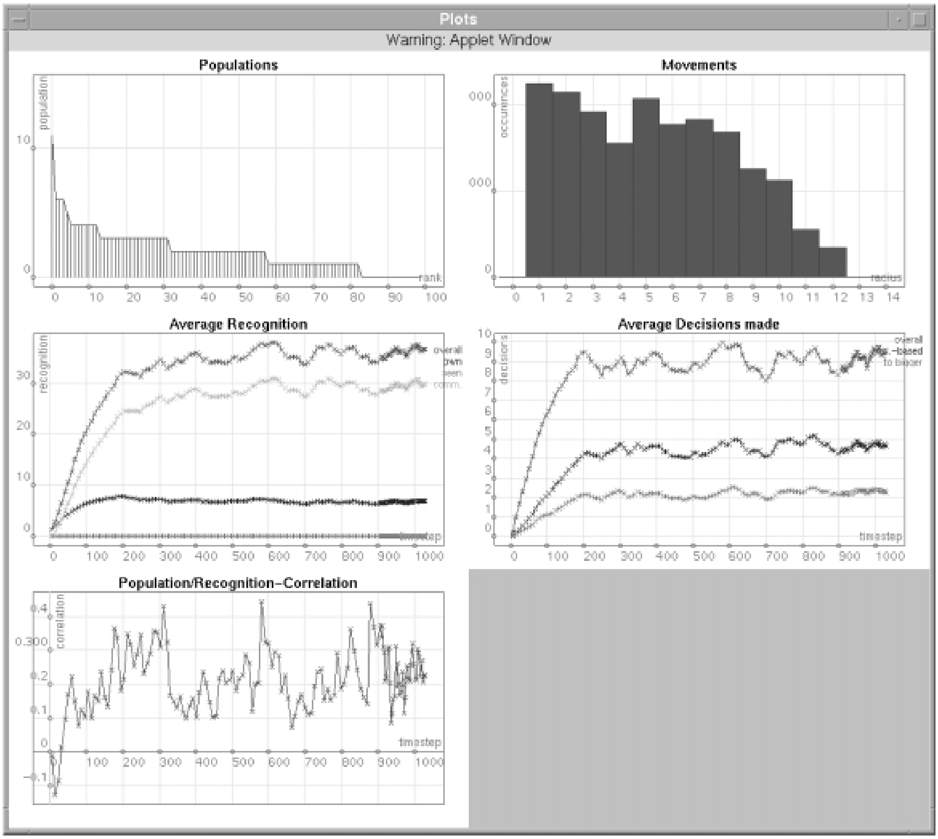

Fig. 1. Statistics for a sample run after 1000 timesteps, showing the distribution of agent population per location against ranked location for this timestep (basically J-shaped), distribution of frequency of movements by their distance, mean recognition knowledge per agent over time (showing from top to bottom total number of locations recognized, number recognized via communication, and number recognized via own experience), mean number of decisions made per agent over time (top to bottom: total, recognition-based, and number that were to the larger-population location), and correlation between population-size and recognition of each location over time.

the population-clusters they could move to and thus can continue to make good recognition-based decisions. Finally, we can consider the ecological rationality of the recognition heuristic in these situations by introducing some fitness differential into the environment—that is, by making some locations "better" or "worse" somehow than others, to see if (and how fast) the population of agents discovers this additional environment structure. If we build in a few "bad" locations where the death rate is higher than elsewhere, for instance, then presumably fewer agents will leave this location and spread the (neutral) word about it, so over time it will be less visited. We could then also give agents the ability to communicate a valence, good or bad, about the locations they have visited (which could be perfectly remembered by audience members, or eventually forgotten, which seems to be the fate of quality tags in other domains), to see how this alters the physical (population-based) and cultural (knowledge-based) environments of the population as a whole. These agents may not know much, but they like what they (and others) know, and that can be enough to change the structure of their world.

References

1. Borges, B., Goldstein, D.G., Ortmann, A., Gigerenzer, G.: Can ignorance beat the stock market? In: Gigerenzer, G., Todd, P.M., ABC Research Group: Simple Heuristics That Make Us Smart. Oxford University Press, New York (1999) 59–72
2. Boyd, R., Richerson, P.J.: Culture and the Evolutionary Process. University of Chicago Press, Chicago (1985)
3. Galef, B.G., Jr.: Social influences on the identification of toxic foods by Norway rats. Animal Learning & Behavior **18** (1987) 199–205.
4. Goldstein, D.G., Gigerenzer, G.: The recognition heuristic: How ignorance makes us smart. In: Gigerenzer, G., Todd, P.M., ABC Research Group: Simple Heuristics That Make Us Smart. Oxford University Press, New York (1999) 37–58
5. Hertwig, R., Hoffrage, U., Martignon, L.: Quick Estimation: Letting the Environment Do the Work. In: Gigerenzer, G., Todd, P.M., ABC Research Group: Simple Heuristics That Make Us Smart. Oxford University Press, New York (1999) 209–234
6. Jacoby, L.L., Kelley, C., Brown, J., Jasechko, J.: Becoming famous overnight: Limits on the ability to avoid unconscious influences of the past. Journal of Personality and Social Psychology **56** (1989) 326–338
7. Kirby, S., Hurford, J.R.: Learning, culture and evolution in the origin of linguistic constraints. In: Husbands, P., Harvey, I. (eds.): Fourth European Conference on Artificial Life. MIT Press/Bradford Books, Cambridge MA (1997) 493–502
8. Kirby, S.: Syntax out of learning: the cultural evolution of structured communication in a population of induction algorithms. In: D. Floreano, J.D. Nicoud, F. Mondada (eds.): Advances in Artificial Life. Lecture notes in computer science, Vol. 1674. Springer-Verlag, Berlin Heidelberg New York (1999)
9. Kirby, S.: Spontaneous evolution of linguistic structure: An iterated learning model of the emergence of regularity and irregularity. IEEE Transactions on Evolutionary Computation (in press).
10. Laland, K.N., Odling-Smee, J., Feldman, M.W.: Niche Construction, Biological Evolution and Cultural Change. Behavioral and Brain Sciences **23** (2000) 131-146
11. Noble, J., Todd, P.M., Tuci, E.: Explaining social learning of food preferences without aversions: An evolutionary simulation model of Norway rats. Proceedings of the Royal Society of London B: Biological Sciences **268** (2001) 141–149.
12. Todd, P.M.: The ecological rationality of mechanisms evolved to make up minds. American Behavioral Scientist **43** (2000) 940–956
13. Wilson, S.W.: Explore/exploit strategies in autonomy. In: In P. Maes, M.J. Mataric, J.-A. Meyer, J. Pollack, S.W. Wilson (eds.): From animals to animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior. MIT Press/Bradford Books, Cambridge MA (1996) 325–332