

Learning in the Cultural Process*

Edwin Hutchins

Brian Hazlehurst

Department of Cognitive Science
University of California, San Diego
La Jolla, California
USA

Copyright 1990 by Edwin Hutchins and Brian Hazlehurst

This paper reports a set of computer simulations that demonstrate a form of adaptation that we believe to be characteristic of human intelligence.

One of the central problems faced by biological and artificial systems is the development and maintenance of coordination between structure inside the system and structure outside the system. That is, the production of useful behavior requires internal structures that respond in appropriate ways to structure in the environment. The processes that give rise to this coordination are generally considered adaptive.

Biological evolution, individual learning, and cultural evolution can all be viewed as ways to discover and save solutions to frequently encountered problems—they are processes that generate coordination between internal and external structure.

Creatures that can learn are likely to have a greater range of responses which, as Hinton and Nowlan (1987) have shown, can lead to learning actually guiding evolution. Hinton and Nowlan imagine a population of creatures, in which the behavior of each individual is specified by some number N of alleles. The creatures inhabit a world in which there is a fitness spike associated with only one particular pattern of those N alleles. If all of the alleles are genetically fixed, the chances of any individual finding the fitness spike is low and no search strategy beats random search. Thus, for creatures in which behavior is entirely genetically determined, the process of

*A version of this paper appeared in D. Farmer, C. Langton, S. Rasmussen, & C. Taylor (Eds.), *Artificial Life II* (1991). Addison-Wesley.

discovering and saving good solutions is blind and relatively slow.¹

Now suppose that rather than having all the alleles' settings hardwired, some can be learned (by guessing in Hinton and Nowlan's scheme) during the lifetime of the individual. In this case, any individual whose hardwired alleles correspond to a partial description of the fitness spike has a chance of guessing the rest of the solution. Creatures that are genetically predisposed to learn (guess) the solution to a particular problem in the environment, by virtue of having correct settings on all of the hardwired alleles, are on average more fit than those who cannot. These more fit creatures may place more individuals who are also predisposed to learn the solution in the next generation. Hinton and Nowlan show that allowing some of the alleles to be uncommitted in value, and thus learnable, has the effect of putting shoulders on the fitness spike such that the population can "hill climb" to the best genetic solution. This scheme is much more efficient and rapid than a random search for the optimal genome. Still, from the perspective of the population, learning is slow because the products of individual learning have only very indirect effects (mediated by selection on random variation) on subsequent generations. (See Hinton and Nowlan, 1987 and Belew, 1989 for details).

Culture is a process that permits the learning of prior generations to have more direct effects on the learning of subsequent generations. As predicted by Wilson (1985), the presence of cultural factors may create selective pressure for the ability to learn itself. Hutchins performed a simple demonstration of this effect by adding a cultural bias factor to the simulation of Hinton and Nowlan. The bias makes offspring of individuals who have learned the solution more likely to learn the solution themselves. This has the effect of increasing the relative frequency of alleles that code for learnable responses. That is, adding a cultural effect increases the steady state proportion of uncommitted or learnable alleles in the

¹The probability of an individual being genetically predisposed to matching the genome associated with the fitness spike is $1/2^n$, for two-valued alleles. In this case, an evolutionary search for the spike is ineffective since there is no feedback regarding a "close" fit, and therefore no opportunity for co-adapted alleles to retain partial solutions. Furthermore, variation in the genomes of descendents means that even if the solution were found, it would be extremely unlikely to be retained in the population.

population. A replication of Hutchins' demonstration and a more complete analysis of this effect was subsequently performed by Belew (1989).

If culture permits the consequences of learning by a prior generation to have direct effects on the learning of a subsequent generation, then could a population, over many generations, be capable of discovering things that no individual could learn in a lifetime? This should be true in spite of the fact that the direct products of individual learning (internal structures) last at most a lifetime. Let us consider this problem in the framework of the coordination of internal and external structure presented above.

In that scheme we had two kinds of structure to be coordinated: external structure - a physical environment, for example; and internal structure - the organization of a nervous system, for example.

Imagine a world in which there is a useful regularity in the environment too complex for any individual to learn to predict in a single lifetime. That is, given the rate at which internal structure can be rearranged, it is either not possible or extremely unlikely that any individual will achieve coordination between external and internal structure. How could a useful form of interaction with such a regularity ever be learned? Hinton and Nowlan have demonstrated one method in which parts of the solution are learned genetically so that individuals in future generations are born partly organized and therefore require less learning in order to master the regularity. But again, that process is very slow. Hutchins' addition of a cultural bias factor showed that culture can guide the ability to learn which can in turn guide evolution. However, any model that reduces all of culture to a single scalar is clearly missing many of the central aspects of cultural phenomena. In particular, culture involves the creation of representations of the world that move within and among individuals. This heavy traffic in representations is one of the most fundamental characteristics of human mental life, yet since it is a phenomenon not entirely contained in any individual, it has largely been ignored by cognitive science. If each individual is capable of learning something about the environmental regularity and then **representing** what has been learned in a form that can be used by other individuals to facilitate their learning, knowledge about the regularity could accumulate over time, and across generations.




FIGURE 1. The relations of environmental, internal, and artificial structure. The arrows represent propagation of constraints upon interactions involving these different kinds of structure. Constraints may be propagated by many means. We use the term “coordination” to refer to the satisfaction of constraints, no matter by what mechanism constraint satisfaction is achieved.

This introduces a third kind of structure: structure in the environment that is put there by creatures. This is **artificial structure**. Our inventory now includes **natural structure** in the environment, **internal structure** in the organisms, and **artificial structure** in the environment. These structures are related to each other through time as shown in Figure 1.

In a cultural world, the internal structure of an organism is shaped by (must achieve coordination with) two kinds of structure in the environment—natural and artificial structure.

What form might an artificial representation of the knowledge of the aforementioned natural regularity have?

First, it must itself be a kind of regularity in the environment. Barring mental telepathy, a mind can only influence another by putting some kind of structure in the environment of the other mind. Taking this maxim seriously highlights the importance of the medium and process of transmission of cultural knowledge. Earlier studies of cultural evolution have not directly addressed this issue (Boyd and Richerson, 1985; Cavalli-Svorsa and Feldman, 1981), choosing instead to speak of cultural traits as if they were abstractions

without physical form. The danger of that approach is that one ignores the artifactual world and overlooks its capacity as a learning system in its own right.

A second requirement for the artifactual representation of natural regularity is that the artifacts be strictly symbolic. They must contain no direct (that is, structurally non-arbitrary) information about the mappings of states of the world onto each other.

It is important to approach this subject with the understanding that culture is not a thing or any collection of things—it is a process. In the human sphere, myths, tools, understandings, beliefs, practices, artifacts, architectures, classification schemes, etc., alone or in combination, do not in themselves constitute culture. Each of these structures, whether internal or external, is a *residue* of the cultural process. The residues are, of course, indispensable to the process, but taking them to be culture itself diverts our attention from the nature of the cultural process. In the simulations we present, the artifacts should not be taken to be the culture of the community. Instead they and the internal structures that form in interaction with them are residues of a cultural process.

Consider the scheme diagramed in Figure 1 as a case of intergenerational cultural process. Getting the internal structures into coordination with the natural regularity of the environment requires three kinds of learning: 1) direct learning of the natural regularity in the environment, 2) mediated learning about the natural regularity from the structure of the artifactual descriptions of it, and 3) learning a language that will permit a mapping (in both directions) between the structure of the natural regularity and the structure of artifactual descriptions of it. We return to this in the discussion of the simulation.

Before turning to the details of the simulation we would like to motivate its organization with a "just-so" story about cultural learning.

LEARNING THE RELATION OF MOON PHASE TO TIDE STATE

Up until about two hundred years ago, the hills on which the U.C.S.D. campus is located were inhabited by California Indians. We know from the ethnographic and archaeological record that they hunted deer and rabbits and collected greens in the canyons inland from the present site of the campus. We also

know that when the tides were low, they collected shellfish in the many tide pools along the coast. That much is well established. Imagine the sort of problem small groups of hunter-gatherers might have faced. Shellfish are a rich source of protein and are easy to get when the tides are low, so when the tides were favorable, it might have been worth moving the whole band to the beach. On the other hand, it would be a waste of energy to go all the way to the beach if the tides are not favorable. Furthermore, it is impossible to determine whether the tides are actually going to be low by just looking for a moment—one might be looking at a time when the tide is at an intermediate level. It would therefore be very nice if there was a reliable way to predict the potential for obtaining shellfish without having to go to the cliffs over the beach and watch for many hours in order to determine the tide state.

The phase of the moon provides such a predictor. When the moon is either full or new, the gravitational forces of the sun and moon are in phase and together generate large tidal variations. So both very high and very low tides occur on the same day. Figure 2 shows this relationship. This regularity, we imagine, would have been advantageous for the members of this society to learn. Of course, they already had a language that contained words for the states of the tide and phases of the moon. The problem here is to learn a set of mappings between states of the natural world—to learn an association between phases of the moon and states of the tide.




FIGURE 2. The regularity between moon phase and tide state results from a phase relationship between the moon and sun's gravitational effects upon the earth.

Since it takes many hours of watching the ocean to determine the actual state of the tide, and since the sky is not always clear, the opportunities for matching an observed state of the moon with an observed state of the tide are few, possibly too few for any individual alone to learn to predict this regularity. But we can imagine that, over time, the community of people could learn something that no individual alone could learn. Since the language is already well-developed, each member of the community learns, as part of growing up, a shared set of mappings between phases of the moon and words for phases of the moon. The same is true of a set of mappings between states of the tide and words for states of the tide. We assume that the phenomena of phases of the moon accompanied by labels for these phases, and of states of the tide accompanied by labels for

these states, are frequently available. It is only the conjunction in experience of the phase of the moon and state of the tide that is limited.

THE SIMULATION

The behavior of a community of individual "citizens" through time is simulated. Citizens are composed of connectionist networks that have the ability to learn from both the natural and artifactual structure in the environment. The latter constitutes a type of symbolically mediated task learning. A "generation" is a time step in the simulation during which each citizen in the population has an opportunity to learn the task. Three stochastic factors account for variability in that learning: a) the quality of the artifact chosen for study, b) the quality of direct experience with the environment, and c) the random set of task network connection weights assigned at "birth." After learning from both the natural and artifactual structures, each citizen generates one artifact, gives birth to one "novice" citizen, then dies. All citizens have the same network architecture and there is no passing of genetic information between generations. Each novice begins life with a random set of connection weights. The only contribution made by an individual to successive generations is a cultural one—a produced artifact.

Each novice of the next generation chooses an artifact from those produced by the previous generation. This choice can be made randomly, or selection can be introduced by having members of the younger generation probabilistically choose an artifact biased by the artifact author's "success." An artifact produced by an author who has learned much about the task, (and who therefore could be said to "know" much about the task solution or environmental regularity) is more likely to be studied by someone in the next generation. The consequences of these two kinds of choice are discussed below.

THE ENVIRONMENT

The representation of the phases of the moon and states of the tide are given in Figure 3. Notice that the representation chosen for this regularity is a continuous version of "exclusive or" (XOR). There are 28 different moon phase/tide state pairs that constitute direct sensory information about the environment. Each of the 28 pairs was generated by dividing the lunar orbit into twenty eight segments and encoding the moon

phase and tide state for what roughly corresponds to the state of affairs for each day of the lunar month.

Each element of the vector representing the moon phase is a real number between 0 and 1. The first element encodes how much of the left half of the moon is visible, the second value how much of the right half is visible, from an idealized earth. Using this representation, every instance of the two-element, moon phase vector describes a unique point on the unit square. Notice that the four vertices of the square represent the four major moon phases; new, first quarter, full, and third quarter, encoded by 00, 10, 11, and 01 respectively.

The tide state is encoded by a single real number between 0 and 1 that is generated by a transcendental function of the angle between moon and sun with respect to earth (see Figure 3). In particular, each side of the unit square is associated with either a decreasing tide variance (between new and first quarter moons, and between full and third quarter moons) or an increasing tide variance (between first quarter and full moons, and between third quarter and new moons).




FIGURE 3. Representation of the environment. Two-value moon phase shown in **bold**, and one-value tide quality shown in *italics*.

THE LANGUAGE

The simulated citizens must discover how to tell when the tide is “good” (maximum variance) and when the tide is “bad” (minimum variance) using the moon phase as a predictor. The language “spoken” by our citizens is shown in Table 1.

This scheme represents two “lexicons,” characterizing two different classes of events; the moon phase and the tide state. Of course each lexicon is not restricted to the prototypical words listed in Table 1. Each placeholder (“bit”) in the symbolic descriptors takes on a real value, providing for a theoretically infinite set of such descriptors. The prototypical words represent an externally defined language that is known (by us) to sufficiently characterize the simulated world’s behavior.

ENVIRONMENT	SYMBOLIC REP	PHYSICAL REP
	Prototypic Moon Lexicon	Moon Phase
New moon	"1000"	00
First quarter	"0100"	10
Full moon	"0010"	11
Third quarter	"0001"	01
	Prototypic Tide Lexicon	Tide State
Large-variance tide	"01"	0
Small-variance tide	"10"	1

TABLE 1. Citizen Language.

THE ARTIFACTS

An artifact is composed of four pairs of **symbols**. The first element of each pair is a symbol for a phase of the moon and the second is a symbol for a state of the tide. In the artifact creation phase, each citizen symbolically encodes responses to the moon phases represented by the vertices of the unit square. Figure 4 exemplifies a "perfect artifact," one that describes a perfect association of moon phases to tide states.

This perfect artifact provides us with a method for evaluating artifact quality utilizing a simple distance metric. Artifact quality is defined by the mean squared difference between the corresponding second elements (tide symbols) of a given artifact/perfect artifact comparison. In other words, it is a measure of the difference between the given artifact's tide symbols and those of the perfect artifact, and is thus the extent to which the artifact is a good symbolic representation of the environmental regularity between moon phase and tide state.

	<u>SYMBOLS</u>	
Pair for new moon	1000	01
Pair for first quarter	0100	10
Pair for full moon	0010	01
Pair for third quarter	0001	10
	<input type="checkbox"/>	<input type="checkbox"/>
	Moon	Tide
	Phase	State

FIGURE 4. A perfect artifact.

CITIZEN ARCHITECTURE

Each citizen is composed of three feed-forward, back-propagation networks: two “language” nets and one “task” net (see Figure 5). Each language net is a standard auto-associating network that is trained to reproduce on its output layer whatever pattern was applied to the input layer. Once trained, the language net provides a mapping between a symbolic description of an event and the event itself. By concatenating instances of these two classes of information (symbolic and physical representations) into one bit string that can be applied to the input units of a language net, the network (after suitable training on this association) can reliably generate: a) a “symbolic representation” from the experience of an event and, b) an experience of the event from a symbolic representation, via the network’s ability to do pattern completion. Using this scheme, each class of symbols (each lexicon) requires its own language net. Thus each citizen has two language nets for translating the artifacts’ encodings of moon phases and tide states, respectively. The task network is a six-unit, one-hidden-layer, XOR network.




FIGURE 5. Citizen Architecture. Fully connected language nets translate the representational structure of artifacts into “vicarious experience” from which the task network can learn.

THREE KINDS OF LEARNING

There are three kinds of learning that take place in the culture process simulations. First, there is *learning of the language*. In the current implementation, this is simply the process of training the citizens’ language nets to associate symbols with events, as described in the last section. For the moment we have left aside the rather interesting problem of how our citizens might come, by consensus, to utilize a shared lexicon suitable for describing the events in their world. (See Hutchins, 1990 for a simulation study of this phenomenon.) Here, we take language learning for granted and simply endow citizens with language ability through auto-associative training on the prototypical lexicons.

Second, there is *direct learning from the environment*. This kind of learning is one employed in standard connectionist modeling. Given some representation of the environment, in our case moon phase vectors and tide state scalars, we give the task network a limited amount of simultaneous experience with these two so it may learn to predict one from the other. A random day is chosen from the 28 day lunar cycle and the task network is presented with the moon phase representation for this day on its input layer. The predicted tide state produced on the output unit of this network is then compared with the actual tide state representation for this day. Any error is back-propagated to adjust the connection weights in a fashion that will better perform this mapping from moon phase to tide state. This kind of learning does not involve the language faculties and is accomplished by a simple presentation of input and target directly to the task network of the citizen.




FIGURE 6. Mediated Learning. The language provides interpretations of inputs and targets for task learning.

Finally, there is *mediated learning*. This learning is characterized by the utilization of language nets to transform externally encoded symbolic descriptions into “vicarious experience” of the events for which they stand. The outputs of the two language nets **produce** the input and target for the task network itself (see Figure 6). “Mediation” takes place at two different levels in this kind of learning. “Inside the skin” of citizens, language faculties mediate between symbolic descriptions and the experience of meaning while “outside the skin” artifacts, structures deposited by other citizens, mediate between events in the world and information about those events.

THE LEARNING PROTOCOL

Figure 7 shows characteristic learning potentials for two of the learning scenarios described in the last section. Each trace on the plot represents the probability of learning two-bit XOR to 0.05 mean squared error criterion (error is averaged over the four cases 00, 01, 11, 10) as a function of learning trials for the labelled scenario. The two scenarios are: a) direct learning of the environmental regularity, and b) mediated learning **from a perfect artifact** utilizing trained language nets to translate the artifact’s symbols. Each trace gives estimators for the probability of learning the respective XOR task based on a random sample of 50 starting connection weight configurations. Note that direct learning involves 28 different cases (randomly presented) while mediated learning only involves the four cases on which the net is tested. Mediated learning from perfect artifacts is thus an easier task, as is reflected in the more rapid learning rates shown in Figure 7.




FIGURE 7. Characteristics of two kinds of XOR learning (see main text). Probabilities are based on observed results from random samples of size 50.

How shall we decide the amounts of mediated and direct learning to give each citizen? If all learning was direct, then culture would be irrelevant. If all learning was mediated, then the structure of the world would be irrelevant. Part of the "just so" story was intended to motivate the notion that direct experience of the environmental regularities might be much less available than mediated experience. Thus, we would like to have a total number of trials that permits individuals to learn the task once good artifacts have developed. Simultaneously, we would like the proportion of direct learning to be such that no individual could learn the regularity from direct experience alone. Figure 7 helps us decide what the learning protocol in each citizen's lifetime should be. If we have fewer than 1800 trials of direct learning, the chances of learning the regularity directly are near zero. This sets an upper bound on direct learning. Clearly, even with perfect artifacts, the total number of trials will have to be greater than 1000 to produce any reasonable learning. This sets a lower bound on **total** learning trials.

The actual protocol used called for a citizen to first get 750 epochs of training from one selected artifact. Since each artifact contains four learning instances, this amounts to 3000 trials of mediated learning. Next, the citizen received 260 trials of direct experience learning. Notice that the probability of an individual learning this task to a criterion of 0.05 mean squared error in 260 trials of direct experience is very near zero. Thus, no individual can learn the task alone (see Figure 7). If the culture can generate

artifacts that describe the regularity well, the combination of a small amount of direct learning and a large amount of artifact mediated learning should permit the individual to learn to predict the regularity.

After learning, the citizen produces an artifact by “responding” to a test of its knowledge of the four orthogonal cases; 00, 10, 11, and 01. This process requires a reverse translation of symbol-to-experience; in particular, it entails the production of symbols which stand for that citizen’s “understanding” of these events (see Figure 8). The production of a symbol for tide state is accomplished via the internal mapping from experience of moon phase to experience of tide state. The artifact thus reflects what the task net has learned about the regularity. We have deliberately excluded internal “propositional” representations that directly link symbols for moon phase to symbols for tide state. The internal models in this simulation are models of the behavior of the natural world, not models of the structure of the artifactual world. Of course, humans do learn the latter type of representation—they may even be the basis for much of human reasoning—but they introduce unnecessary complexity into this simple world.




FIGURE 8. Generating an artifact.

RESULTS

Figures 9 and 10 show the results of two simulations run with population sizes of twenty citizens each. There are two traces plotted for each run reporting the generational averages of 1) artifact quality, and 2) the average mean squared error of each citizen's task performance on the four prototypes (00, 01, 11, 10). As is evident from these figures, the two measures of the culture process track each other quite closely.




FIGURE 9. Culture process simulation with no artifact selection bias.




FIGURE 10. Culture process simulation with biased artifact selection.

The difference between these two simulations is that the one shown in Figure 10 utilized an artifact selection bias. As already mentioned, this amounts to tagging artifacts with a selection probability that is a function of the author's task competence. The probability function utilized simulates a uniform distribution of artifacts based upon the observed task competence of the authors' deviations from that of the most competent author (i.e., the one with the lowest MSE on the prediction task). Selection bias based on author competence seems like a reasonable, though simplified, analog of what takes place in real cultural process.²

Without this biasing it appears that while the system is capable of some learning, it is vulnerable to unlucky choices of artifacts to study, resulting in slower and less dramatic learning. Just as artifacts are a bridge from the internal structures of one generation to those of another, so the internal structures of individuals are the bridge between one generation of artifacts and another. If too many individuals in subsequent generations study artifacts created by poor performers, the useful structure that has been built into good artifacts can be diluted by the noise in bad ones. This can cause the community to "forget" some of what it knows about the regularity. Nonetheless, Figure 9 shows that even with random selection there is an accumulation of knowledge that affords better task performance for individuals in later generations.

Finally, Figure 11 shows the effects of participating in a cultural system on the learning abilities of individuals. In the early generations individuals learn from artifacts with no useful structure, and no member of the community is able to predict the environmental regularity. In later generations using exactly the same learning protocol, virtually all of the individuals are able to predict the regularity. This happens even though the individuals in later generations have no greater innate learning abilities than those of the early generations. Clearly this phenomenon results from retention of "successful" knowledge in the artifactual media.

² Boyd and Richerson (1985) also utilize this type of biasing in their models of culture and biology as co-evolutionary processes.




FIGURE 11. Observed population percentages of citizens who learned the task to 0.05 MSE or better during simulation with biased artifact selection (see Figure 10).

Although the simulation presented here is quite primitive, we believe it illustrates, in principle, that the cultural process can be seen, like biological evolution and individual learning, as a way to produce and maintain coordination between internal and external structures. It has been said that culture is the most important invention in the history of life since sex (Sereno, 1990). We hope that in this paper we have been able to show both why the cultural process is so important to human mental life and why in considering the cultural process we must consider the role of artificial as well as natural structure in the environment.

ACKNOWLEDGMENTS

Research support was provided by grant NCC 2-591 to Donald Norman and Edwin Hutchins from the Ames Research Center of the National Aeronautics & Space Administration in the Aviation Safety/Automation Program. Everett Palmer served as technical monitor. Additional support for Edwin Hutchins was provided by a fellowship from the John D. and Catherine T. MacArthur Foundation.

REFERENCES

- Belew, R. K. (1989). Evolution, learning, and culture: Computational metaphors for adaptive algorithms. (Tech. Rep. No. CS89-156). La Jolla: University of California, San Diego, Computer Science and Engineering.
- Boyd, R., & Richerson, P. J. (1985). *Culture and the evolutionary process*. Chicago: The University of Chicago Press.
- Cavalli-Sforza, L. L., & Feldman, M. W. (1981). *Cultural transmission and evolution: A quantitative approach*. Princeton: Princeton University Press.
- Hinton, G. E., & Nowlan, S. J. (1987). How learning can guide evolution. *Complex Systems*, 1, pp. 495-502.
- Hutchins, E. & Hazlehurst, B. (1995). How to invent a lexicon: The development of shared symbols in interaction, in N. Gilbert and R. Conte (Eds.), *Artificial Societies: The computer simulation of social life*. London: UCL Press.
- Sereno, M. (in press). Four analogies between biological and cultural/linguistic evolution. *Journal of Theoretical Biology*, 1990.
- Wilson, A. C. (1985, October). The molecular basis of evolution. *Scientific American*, 253, No. 4, pp. 164-173.