# Adapting the Environment Instead of Oneself

David Kirsh*
University of California, San Diego

*This article examines some of the methods used by animals and humans to adapt their environment. Because there are limits on the number of different tasks a creature can be designed to do well in, creatures with the capacity to redesign their environments have an adaptive advantage over those who can adapt only passively to existing environmental structures. To clarify environmental redesign, I rely on the formal notion of a task environment as a directed graph in which the nodes are states and the links are actions. One natural form of redesign is to change the topology of this graph structure so as to increase the likelihood of task success or to reduce its expected cost, measured in physical terms. This may be done by eliminating initial states, hence eliminating choice points; by changing the action repertoire; by changing the consequence function; and, lastly, by adding choice points. Another major method for adapting the environment is to change its cognitive congeniality. Such changes leave the state space formally intact but reduce the number and cost of mental operations needed for task success; they reliably increase the speed, accuracy, or robustness of performance. The last section of the article describes several of these epistemic or complementary actions found in human performance.*

*Key Words: adaptation; task environment; redesign; epistemic actions; complementary actions*

## ıtroduction

A creature has at least three logically distinct strategies for improving its fitness.[1] It can *adapt to* the environment, *migrate* to new surroundings, or *adapt the environment itself.* In this article, I shall examine the third strategy, redesigning the environment.

The problem of when to adapt, when to redesign an environment, and when to search for a new habitat is broad enough to be treated as a general fact of life. Humans,

Department of Cognitive Science, University of California, San Diego, La Jolla, California; kirsh@cogsci.ucsd.edu

1 Although I am posing the problem as one facing an individual creature, we could equally well pose the problem for populations, intelligent agents, groups, or any unit that can be thought capable of adapting to an environment, leaving an environment, or modifying an environment.

 

no doubt, are among the few creatures who (explicitly) reason about this problem; but it is certain that other animals tacitly face a similar problem and that evolution has wired in a partial strategy for solving it. The reason we may expect creatures to have built-in strategies for redesigning their environments is that creatures with some active control over the shape of their environment will have an adaptive advantage over those who can adapt only passively to existing environmental structures. This follows because the number of tasks in which a creature can be designed to do well is limited. It is unlikely that a design which allows a creature to be exceptionally fleet-footed on flat terrains also will confer exceptional agility on rocky terrains. Design trade-offs must be made. If, however, a creature could somehow change the rocky terrain it visits (by making paths), or if it could augment its design with prostheses or tools (e.g., carry cups of water in dry regions, use "fishing sticks" to poke narrow crevices), its existing design becomes more adaptive.

We do not have to look far to find examples of how animals and humans modify their environments in adaptive ways. Beavers dam ponds, birds build nests, ants farm aphids, chimps leave useful nut-cracking stones in commonly used places (Kummar, 1995), squirrels collect nuts for winter, and Egyptian vultures drop stones on ostrich eggs. In each case, the environment is warped to the creature's capacities rather than the other way around. This is to be expected within a classic adaptationist approach, for in the struggle for existence, organisms with these favored behavioral tendencies will outreproduce their competitors with less favored behavioral tendencies; if these tendencies are heritable, the distribution of tendencies will change over generational time, with the favored tendencies becoming more common. Hence, classic adaptationism predicts that environmental modification will occur (Butler, 1995).

Nonetheless, evolutionary arguments do not explain *why* these traits are adaptive. For that we must look to economic and ethological arguments. Thus, because foodstuffs are scarce in winter but plentiful and cheap in the fall, it is wise to stock up in the fall (standard inventory control principles). Because the probability of finding an ideal nesting site is small, at some point it becomes more cost-effective to build a nest in a suboptimal spot than to continue searching (investment analysis). And because good nutcracking stones take time to find, it is better to leave them where they will be most useful for everyone concerned than to discard them in arbitrary ways (amortize the cost of search). When such explanations are offered, we believe we have understood the phenomena better but, if we look closely, we note that behind such explanations there is an appeal to a notion of task environment that is distinct from the notion of the selective environment present in evolutionary arguments. It must be so, for the economic principles appealed to all show that a given behavior X is preferred to behavior Y with respect to a certain goal. They cannot show that behavior X is better than Y, all things considered. Thus, leaving around a nutcracking stone is desirable for the task of nutcracking, but it may not be a good tactic overall because it may have a *side effect* of cueing predators. What is good for a task in isolation may not be good when all tasks are considered together.

My objective here is to explain this notion of task environment in order to extend the range of analyses available to understand active redesign. I will draw on ideas from computer science, economics, the theory of problem solving, and the study of interactivity to clarify what environmental redesign is and suggest ways we might measure its benefits.

This article is organized into five major sections. After this introduction, I distinguish the notion of a task environment from that of the environment more broadly construed. Task environments are substructures within a more all-encompassing selective environment. This distinction supplies us with the descriptive apparatus to pose the problem of environmental redesign in a natural way: Namely, how is a creature to change the structure of at least one task environment so that its overall performance—that is, its performance summed over all task environments—is improved. This global improvement may be achieved by enhancing performance in one or more task environments without reducing performance in any other (Pareto optimization) or by enhancing performance in one or more environments to more than compensate for any reductions incurred in others.

Because a task environment is a substructure within a larger selective environment, it is necessary to provide effective criteria for deciding which parts of the larger environment fall within a given task environment and which parts fall outside it. This is undertaken in section 3. Using the language of computer science, I define the structure of a task environment to be a directed graph in which the nodes are states and the links are actions. Because not all actions available in the larger environment are allowable as moves in a particular task environment, we distinguish actions that are *internal* to a task from actions that are *external*. For example, actions that modify the structure of a task can be achieved only by performing actions that are external to the task. This distinction makes it possible to claim that creatures perform actions for the sake of redesigning a task environment.

The remainder of the article focuses on some of the different types of *task-external* actions available to both human and nonhuman animals for redesigning task environments. Two broad categories are distinguished: (1) external actions that change the topology of a state space and that we may expect to find prevalent in animal populations, and (2) external actions that change the cognitive congeniality of a space and that are an important and understudied feature of human and cultural adaptation.

## 2 Environments Versus Task Environments

The term *environment*, in normal biological parlance, refers to everything exogenous to a creature that may affect its physiology, experience, or death. This includes the full range of external factors that may affect an agent's experience and motivation—factors affecting its internal state broadly understood—as well as all factors determining its possibilities for action and the consequences of its actions. For our purposes here, it will be sufficient to use the term *environment* in a slightly narrower manner to refer to *the totality of cues, constraints, resources, and dangers* in a creature's world that determines its success, where *success* means its differential reproductive success. *Environment*, then, means the "selective" environment (cf. Brandon, 1990).

One feature of the notion of selective environment is that it abstracts from most of the microstructure in a creature's niche. If the environment is the totality of external cues and resources that can make a difference to reproductive success, there is no distinction between those cues and resources that are relevant to one type of task (e.g., food collection) and affect the success or failure of particular foraging strategies, and those cues and resources that are relevant to another task (e.g., predator avoidance) and affect the success or failure of particular defensive strategies. Although the creature inhabits only one world, for certain types of analysis it is useful to circumscribe that one world into subdomains, each of which is relevant to the success or failure of particular task strategies. Following psychologists, I shall refer to these subdomains or microenvironments embedded in the larger environment as *task environments*. A creature's selective environment, then, is a superposition of task environments.

Once we distinguish the selective environment (broadly construed), from particular task environments within it (microenvironments), it is easy to state one problem that evolution, learning, or intelligence must "solve" for a creature. It must determine how to conform to the optimality principle:

> *Optimality principle*: Efficiently allocate time and energy to the performance of different tasks so as to maximize overall "return" (Lewontin, 1978; cf. Horn, 1979).

For example, we expect that a well-adapted creature will efficiently divide its time among hunting, drinking, exploring the terrain, hiding from predators, finding and attracting mates, and the like. Given the returns inherent in each task environment, the creature must allocate its resources among its different behavioral strategies so as to maximize its overall yield.

Ethologists have been tackling this resource allocation problem for some time now, using the language and methods of economics. For instance, to decide how much time an optimal lizard ought to spend foraging versus hiding from predators, an ethologist would study lizards in their natural habitat and then plot the cost-benefit curves of foraging and the cost-benefit curves of hiding. An optimal creature ought to allocate its time and energy between foraging and hiding so that, at the margin, it is equally profitable whether the next minute is spent either foraging or hiding. This follows because if one activity—say foraging—were more profitable, an optimal creature would invest more time in foraging until the danger of being out in the open would lower the profitability of foraging, making hiding a more attractive option. In this way, it is possible to identify an optimal allocation strategy (Koza, Rice, & Roughgarden, 1992).

The calculations here are based on simple microeconomic principles of optimization. The hard work, of course, is to discover the underlying facts about the lizard environment that permit the cost-benefit curves to be constructed. For instance, it is necessary to determine both the probability of predation as a function of distance from the hiding spot and the probability of catching prey as a function of distance. This is now a well-accepted methodology in ethology. Once again, the entire enterprise assumes that activity-salient features of the environment can be identified. Features of the environment that are relevant to an activity can be distinguished from features that are irrelevant.

If we accept this microeconomic way of posing the problem of adaptation, we can easily restate our question of environmental redesign. If a creature already is allocating its time optimally among all the task environments in which it operates, then the only way it has of increasing its yield, assuming it does not change its physiology or one of its behavioral strategies—a classic process of self-adaptation that we assume requires generational time—is to increase the yield of one of its behavioral strategies. This gives rise to a principle of superoptimality:

> *Superoptimality principle*: If a creature already is allocating its time optimally among its different tasks, the only way its overall welfare can be increased is if the payoff function for one of its strategies increases.

Because the payoff function is determined by the environment,[2] this principle means that one of the creature's task environments must change. Assuming that the forces of change are not stemming from the global environment, this implies that either the creature migrates to a new habitat where it is easier to achieve one or more

2 Payoff functions always are defined with respect to a creature's physiology and behavioral strategies. Thus, a physiological change having the effect that a creature's need to drink or eat is lower would change the shape of the payoff curve for eating and drinking, as thirst and hunger thresholds now will be different. Similarly, the payoff function also is tagged to particular behavioral strategies, as a change in how a creature behaves (particularly if this involves a change in its capacities) requires identifying the consequences of new actions or action sequences.

of its goals, or it redesigns its environment. (See also Fletcher, Zwick, & Bedau, "Dependence of Adaptability on Environmental Structure in a Simple Evolutionary Model," in this issue for a related view.)

In microeconomics, individual agents cannot migrate to new economic environments, cannot change their basic internal structure (physiology, technology) or behavioral capacities and, significantly, cannot alter the payoff function by redesigning their environment. This last constraint holds because it is assumed that individual economic agents are unable to alter the structure of their task environments. Any one firm cannot change the price or supply of raw materials, the market price of its product, or change the available methods of production. This is not to say that the totality of firms making up an industry is unable to influence prices or affect technology. Collectively they can, but this power to reshape the environment does not reside in individual agents. It is a macroeffect: The shape of the payoff function lies outside the control of any one agent. It is a given of the environment and so not a manipulable variable (Simon, 1955).

The same assumption applies in classic ethology. For ethological analysis to get off the ground, it is assumed that an individual predator cannot significantly change the prey population (input supply)—although, of course, there are aggregate or macroeffects of the entire population of predators that result in predator-prey cycles, trends, and so on. Nor can individual predators change the metabolic benefits (price) it obtains from eating prey, nor the hunting techniques it has at its disposal. These are the givens of the adaptive context in which it finds itself. Evolution can change metabolism, culture can change hunting techniques and, in the course of a single life span, predator-prey cycles may change the relative proportion of prey to predators, but all these changes must be regarded to be macrolevel changes, events that help to shape and define the microenvironment of individual predators.

Because the behaviors in which I will be interested often occur on a very short time scale (relative to more gradual, smooth evolutionary changes) and often with limited effect (just affecting the environment of one creature, or possibly a few), the environmental changes brought on by populations are not my primary concern. This does not mean that such changes in task environment structure cannot be explained using the analyses I will propose. It means, rather, that there is a class of environmental redesigns that are likely to slip beneath the filter of evolutionary selection.

Despite the importance of the assumption that agents do not change their task environments, this idea is not enshrined in the theory of natural selection, which is concerned with population changes in the selective environment. There was nothing in Fisher's (1930) original mathematical analysis of natural selection that required agent environment independence, although Fisher himself relied on the assumption to prove his fundamental theorems. The reason Fisher assumed that environments would be the same both before and after adaptation was to prevent having to change the differential goodness of an attribute as it spread through the population. For example, if a gene for increased appetite and stomach size enters the population and confers an adaptive advantage on its owners, Fisher (1930) assumed that as the gene spread through the population, the differential goodness of larger stomachs and appetites would not diminish even though a consequence of more animals with larger stomachs might be that there is less vegetation available for other members of the population. Put differently, because every environment has a limited capacity to accommodate its population (its carrying capacity, $K$), we might well expect that increased appetite alters the carrying capacity of the environment, leading to a smaller sustainable population and also to a reduced selective advantage (reproductive rate $r$) as the population nears $K$. Fisher did not alter $r$ and $K$, but again there is no mathematical reason, short of simplicity, to require these to be constant.

Not surprisingly, there have been several efforts to accommodate the density dependence of attributes. (See Royama, 1992, for a good account of density-dependent parameter models.) Such models permit the spread of an allele to alter the carrying capacity of the environment. This is clearly a step in the right direction. However, one assumption limits all these models: They all assume that $r$ and $K$ vary smoothly with the spread of the attribute. Accordingly, there are no threshold effects or jumps in $K$ or $r$ as a result of reaching certain population densities. This is not entirely realistic. Such jumps might happen, for example, if greater appetite leads to greater defecation, which in turn either increases the vegetation yield nonsmoothly because of threshold effects, hence nonsmoothly increasing the carrying capacity, or if greater defecation decreases the yield nonsmoothly through threshold toxicity, suddenly killing off vegetation and so nonsmoothly decreasing carrying capacity.

Introducing agent environment codependence is an important first step in allowing individual agents to have a significant effect on the structure of their environments. However, as noted, there remains an assumption in all these models that such codependence is highly constrained, leads to smooth changes in environmental structure and, most importantly of all, is the kind of codependence for which there can be genetic selection. If, as I believe, there exist examples of environmental redesign that are idiosyncratic to particular agents and that sometimes result in nonsmooth changes in fitness, research at the level of evolutionary selection will require complementary studies to explain these phenomena. Again, this does not imply that evolutionary selection is irrelevant; the capacity to exploit elements of the environment opportunistically to enhance task performance is a valuable capacity to pass on. But it does imply that the specific ways individual creatures have of redesigning their environments may be localized in time and space and so not fully explicable from a selectionist standpoint. To make this case plausible, it is necessary to define exactly what we mean by a task environment.

## 3 What Is a Task Environment?

The adaptationist program encourages an engineering approach to organisms: In the ideal case, well-adapted creatures ought to converge on optimal designs for particular tasks (McFarland & Houston, 1981). The point of partitioning the global environment into a collection of task environments is to provide us with a theoretical abstraction, the task environment, that is sufficiently well defined that we can evaluate the performance of competing organism designs or competing algorithms for action by using concepts drawn from computer science. We want to be able to compare the efficiency of different algorithms for carrying out particular tasks.

Formally, a task can be represented as a directed graph, in which the nodes of the graph denote choice points (i.e., possible states), the links represent transitions or actions, and a privileged set of states represents the possible ways of completing the task. A successful effort at the task can then be understood as a trajectory, or path, through this graph structure, starting from an initial state and ending at one of the states satisfying the goal condition. Typical tasks we might hope to represent as trajectories include caring for cubs, building a nest from local debris, collecting termites from a 6-foot mound, damming a stream, avoiding a charging predator, mating, and so on.

In the human world, the tasks for which a directed graph representation might be constructed range from highly structured activities, such as playing solitaire, solving an algebraic problem, or making a curved surface in a graphics program, all cases where there are a small number of possible actions at each choice point, to less formal tasks, such as cooking, cleaning, driving to work, and even writing an essay, for which the actions available at an arbitrary choice point are more difficult to enumerate and success is more difficult to measure. In studying behavior in these tasks, a researcher attempts to determine the topology of the state space and to discover a plausible metric to permit comparing the goodness of different plans or algorithms for performing the task. Typical metrics might be the minimum number of actions required to get from the current to the goal state, or the amount of energy required to reach the goal, or the reliability of the paths to the goal. It then is possible to rank different algorithms with respect to how much energy they require for task completion, or how reliable they are, or how many actions they use. It is believed that this methodology will scale to large tasks too, as shown by the vast literature on administrative behavior (Simon, 1976), workplace design (Kroemer, 1993), human factors (Teichner, 1971), and human computer interaction (Diaper, 1989).

To be more precise, let us follow Simon (1973) and say that a task environment is well defined if the following conditions apply:

- There is a well-defined *initial state* in which the agent begins.
- There is an operationally defined *goal condition*, satisfaction of which represents success at the task or which allows us to determine the degree of success the creature has achieved.
- There is a set of discrete, task-relevant *choice points*—nodes in the state space—although these may be arbitrarily dense.
- There is a determinate but finite set of goal-relevant actions (including the null action) available at each choice point—*the option or feasibility set*—which is a subset of the agent's task-relevant *action repertoire*.
- There is a *consequence function* that determines for each action taken in a state the new state that will be produced.
- There is a *metric* that specifies for each state the cost required to reach the next state (or possibly the goal), as measured in steps, energy, reliability, or time.

Action selection can then be seen as the application of two successive filters (Elster, 1979). Given a choice point, filter the action repertoire to yield a feasible set, then filter the feasible set to yield a choice set by applying the metric to the consequences flowing from each feasible action and invoking a decision rule to select the best. If the choice set is determined by a simple decision rule (filter 2), such as maximize expected return, each action in the choice set will have the same expected utility. If the choice set is determined by a different decision rule, such as minimize the worst possible injury or cost, each action will serve the conservative function of leading to states that are not likely to be costly, even though they may have very different expected benefits. The action actually carried out is any one of those in the choice set and is assumed to be arbitrarily chosen, although often there will be only one action in the choice set.

To see how we can use the directed graph notation to understand behavior, consider how we would explain a bird constructing its nest. We would begin by defining the goal condition as, say, the construction of a stable structure with certain shape and size, and having certain thermal and tensile properties. The initial state would be the site chosen on which to build, such as a branch or small hole, as well as the distribution of useful resources or debris—the local inventory—to be found in the region over which the bird is willing to scavenge (e.g., sticks, feathers, paper, saliva, and dirt within a radius of 50 m). Choice points then are introduced by assuming that there is a set of construction-relevant actions available at various points in the nest-building process. For instance, if the nest walls need to be heightened and the bird's beak is empty, the bird faces the choice of selecting which of a small set of appropriate nearby inventory items should be *fetched*. If an item is already in the beak, then the bird must choose among *placing* the item on one of a small set of points or surfaces on the emerging nest, rejecting the item and *dropping it*, or storing it for later use by *placing it to one side*. As each action is completed, the state of the environment changes and the bird faces a new choice point determined by the demands of the task. The distance from the goal, meanwhile, might be estimated by how many items of the inventory will be required to finish the nest or by how much energy will be needed. The consequence of any action is given by stating how the nest and bird each differ after tamping, or after the bird fetches a stick, and so on. By using a directed graph to represent the nest-building process, we distinguish trajectories of nest building that are likely to be successful from those likely to be unsuccessful. When this analytical framework is used in conjunction with empirical observation of these trajectories, we hope to infer the strategies, plans, and behavioral programs that animals actually use.

One advantage of developing an account of behavior that treats it as algorithmic is that we can consider some of the computational properties of the algorithm regulating it. (See Harel, 1987, for a simple account of these properties.) How much memory is required to follow this algorithm? How many steps will be required in the worst case or in the average case? How robust is the algorithm to interference? to noise? If we can determine these properties of behavior control strategies, we can explain the adaptive advantage of different strategies.

The factor that complicates this simple picture is that we cannot *simply* infer these computational facts. The robustness of a program may depend on how it is implemented. A production system implementation (Newell, 1973) of a skill or strategy that consists of a set of rules that trigger only when the creature and environment are in a certain state may have one level of robustness and performance, whereas a recurrent network implementation of the same skill or strategy may have another. Hence, we cannot go directly from the behavioral analysis of programmed behavior to the computational properties of the "program" that regulates behavior, because our analysis will depend on how we think the creature is programmed. The reason to be hopeful about this line of research, though, is that there is a well-established belief in computer science that, at an abstract level, we can discuss the absolute complexity of a task or problem in a manner that abstracts from arbitrary aspects of an algorithm and implementation (Harel, 1987). Thus, even if we do not know whether a creature has a working memory capable of storing $M$ items, we still can know that the creature must be capable of keeping track of $M$ items if it is to perform the task and that it will have to perform at least $N$ (mental or physical) manipulations on those $M$ items. How it does this storing is implementation-dependent. It may somehow encode markers in working memory and manipulate them internally, or it may have a less symbolic memory and rely on visual strategies for tracking or quickly isolating recently seen items in the environment and then physically may manipulate those external items. We cannot say in advance of detailed research. Nonetheless, we can rank different algorithms by their "memory" requirements, and we can rank different tasks by their complexity in the sense that the most efficient algorithm capable of solving them will have certain complexity measures (Papadimitriou, 1995).

### 3.1 Justifying a choice of a state-space representation

The concept of a task environment that I have been presenting is an abstraction laid over the interactions between an agent pursuing some goal, which we call its *task*, and the physical environment in which it is acting. The point of the abstraction is to reveal the constraints on goal-relevant activity inherent in agent environment interaction, so that we can explore the subgoal structure of the task and thereby explore the costs and benefits of pursuing different paths in the task state space.

To justify a particular state-space analysis of a task, it is necessary first to justify one's choice of action repertoire, choice points, option set, consequence function, metric, goal condition, and initial state. Of these, the most crucial choices are the action repertoire, option set, and consequence function, as these determine the nonmetric (or qualitative) topology of the space. How then do we choose the states and actions that allow us to interpret activity as a trajectory in a particular task state space?

There are really two issues here. First, how do we decide how to classify the states and actions occurring in task performance? Second, how do we decide whether an action observed while a creature is engaged in a task is the type of action that qualifies as a move *within* that task space as opposed to being an action that just happens to occur in the same time frame as the task?

A word about the first problem: Because a task environment is purely a theoretical construct, its adequacy depends entirely on its success in explaining behavior. For instance, in describing the task environment of ping-pong, we are free to choose a set of states and state transitions (actions) that we think will clarify the structure inherent in the game. Thus we might choose *impacts* as the states, and divide state transitions (actions) into two sorts: *strokes*, which are the transitions caused by players, and *bounces*, which are the transitions caused by the ball hitting a nonmoving surface, such as the tabletop or the player's body or the floor. Each transition has two parameters: spin and nonangular momentum. Alternatively, though, we might characterize the game more qualitatively. There are several types of actions: *forehands, backhands*, and *smashes*, each with parameters such as *topspin, underspin*, and *flat*. The states are described from the perspective of the player currently stroking the ball and are characterized qualitatively by such terms as *topspin deep court, smash down the line*, and so forth. The advantage of such a qualitative account is that it coheres more with how players themselves think about the game and also lends itself to describing strategies we know players know. Nonetheless, such an account is no less a theoretical construct imposed on the continuous set of changes occurring in a real ping-pong game than is the first description. The choice between descriptive languages must be based on the quality of the task analysis it permits. This is all the more true when the animals performing the task lack a system of descriptive concepts for events and situations.

This leaves us with problem two: How do we decide which actions of a creature's action repertoire belong to a task and so are internal to the task environment, and which actions are, properly speaking, not part of performing the task and so lie external to the task environment? Three natural criteria present themselves as necessary:

1. *Measurable effect on performance*: Actions should have a measurable effect on the time, energy, or number of moves an agent would need to complete the task. Hence, chirping while searching for suitable nest-building twigs, or stretching while hiding in a cave, would be task-external actions, as they have nothing to do with nest building or hiding per se and should have no impact on task performance.
2. *Not too general*: The action is not so prevalent, or general, that it has nothing specifically to do with the task. For example, breathing, blinking, and perspiring are actions that, if omitted, may drastically affect performance. Nonetheless, their contribution is more to the general condition of the agent than to anything task-specific. They are background actions that contribute to the creature's normal state of being task-ready but are not themselves part of any task in particular.
3. *Not too specific*: The action is not so idiosyncratic, so specific to one creature, that any other creature also performing the action in the same task circumstances would neither improve nor worsen its position in the state space. For example, suppose we were to find particular creatures for whom such actions as twitching their whiskers while chasing prey, scratching their ears while hiding, or licking their paws while playing with mice can help them to hunt or hide or play. When prohibited from doing these things, these particular animals perform less competently at these tasks. Hence, the actions all have a measurable effect on performance. They satisfy condition 1. They also satisfy condition 2, since they are not necessary to keep the animal in its normal state. Still, it would be odd to view these actions as part of a general strategy for hunting, hiding, or playing. Given our understanding of the tasks in which they occurred, they seem to have more to do with the history and idiosyncrasies of the particular creature who displays them than with the task itself. Hence, they ought to be treated as task-external.

Applying these criteria is not always an easy job. We cannot always have advance knowledge of all the actions that are potentially relevant to a goal, so it may be difficult to decide whether an action is too idiosyncratic to count as an action within a task or whether it is part of a novel strategy for task performance. Empirical studies should help but this remains a real problem. As formal criteria themselves, however, these conditions seem both necessary and sufficient to let us distinguish actions that are internal to a task from those that are external.

I now want to show that if we stick to these natural criteria, all sorts of external or "metatask" actions are possible that can have significant effects on performance. Actions taken outside a task may have *side effects* that affect performance within the task. If there are regularities in the way they may be created, the side effects can be exploited by a creature so that, in effect, the creature may adapt the task to the strategies it has rather than adapting its strategies to the task.

We may distinguish two broad families of strategies creatures possess for shaping their task environments: (1) strategies for *deforming* the topology of the state space and (2) strategies for making existing state spaces more *cognitively congenial*. Of these two, changing congeniality seems reserved for higher animals, particularly humans, so my examples of those strategies will be drawn from studies and analyses we have done with humans. However, close observation of animal populations may yet reveal that even for these highly cognitive strategies, there are rudimentary analogous strategies occurring in the wild, particularly among creatures with cultures.

Let us turn to the first category of task-external actions.

## Strategies for Deforming the State Space

The structure of a task, I have been arguing, is given by its state-space topology. It follows that to change a task environment is to change its topology. Because any modification of the choice points in a creature's environment, its action repertoire, or the consequence function (or metric defined over that environment) will add or subtract nodes, add or subtract links, or change the distance between nodes, a natural strategy of redesign is to alter one of these constituents.

To establish that a change is to a creature's advantage, let us assume that behavioral strategies can be analyzed as procedures, or *algorithms*. Relative to a particular task environment, an algorithm will have a particular set of costs and benefits. These may be measured in computational terms, such as time (number of steps or actions), or space (amount of memory, items to be tracked in the environment) required to accomplish the task. However, we may also use broader metrics, such as the amount of energy or labor required to complete the task, the robustness of the algorithm to interference and noise, and so on. It is important to be clear that, at this stage of our inquiry, the term *algorithm* will be reserved for the program or control structure organizing the actual behaviors a creature reveals in task performance. Because every attempt at performing a task will be interpreted as a trajectory of actions through a task state space, we can ask what are the properties of a program responsible for these trajectories. For lower animals, these programs may resemble a collection of highly constrained action routines (tropes). Accordingly, action is largely reactive to local conditions. As we ascend the great chain of being, however, we expect the programs to become larger, more flexible, and sensitive to nonlocal features of the environment. When humans and perhaps the great apes are the agents, their programs become sufficiently complex that we cannot understand their effects unless we also understand the regulating role of these creatures' mental operations in the more general system consisting of both mental and physical acts. At that point, we must expand our notion of algorithm to include the idea of a control structure regulating both mental and physical actions and the tight coordination between the two. This idea is expanded in section 5, where we discuss cognitive congeniality.

Because algorithms regulate activity within a task environment, we are assuming that they do not change a task; rather, they cause transitions to new states within the task. This is an observation of some import. To change the task environment requires executing an action that lies outside the normal algorithm for the task. This makes creating a better task environment resemble selecting a better habitat. The similarity is only superficial. In environment redesign, the creature remains in the same geographical region and is itself responsible for the change in environment. The global environment does not present the creature with a range of preexisting habitats, differing in salient respects, from among which the creature then chooses. Rather, the creature itself actively creates the changes from a different preexisting environment. Thus, in habitat selection, the environment is assumed to have its task characteristics independently of the creature, whereas in active redesign, the environment has been forcibly changed and may be expected to return gradually to its original state on the creature's death.

Let us now review some of the different ways individual creatures have of altering the topologies of their task environments. These include eliminating initial states, hence eliminating choice points (the "just say no" strategy and method of routine maintenance), changing the action repertoire (the methods of routine maintenance and tool use), changing the consequence function (scouting ahead), and adding choice points (tool use).

### 4.1 Just say no

The first example of a strategy for deforming the task environment is best understood in terms of a filter that sifts out undesirable choice points, particularly the initial

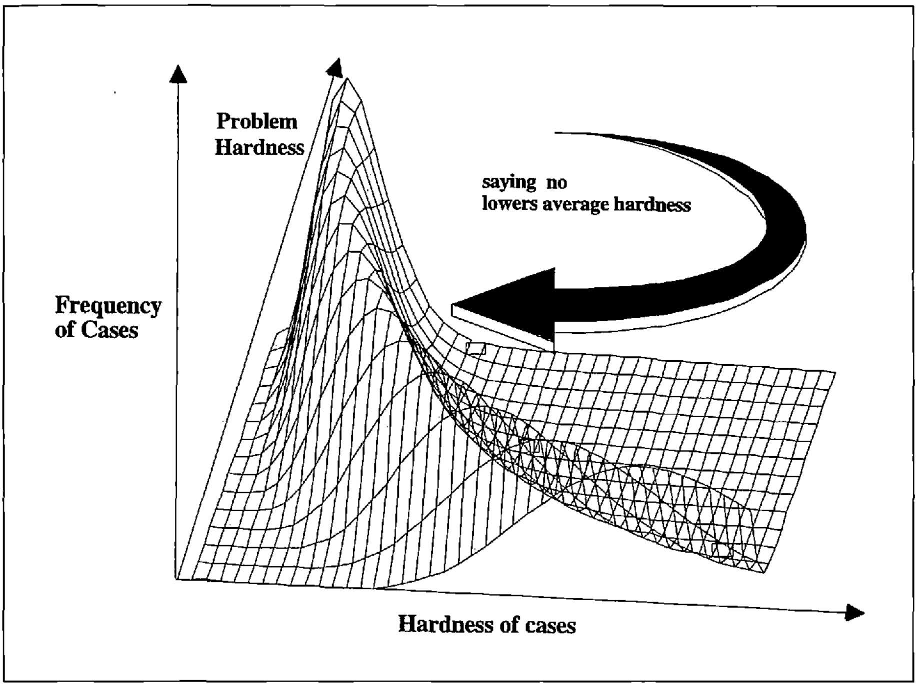


**Figure 1**
Here we see the probability density function showing the distribution of difficult and easy cases in a task environment where the "just say no" strategy is used. As difficult cases are filtered out, both average and worst-case performance improves.

choice points that a creature faces on entering a task. In the course of its life, a creature comes across a variety of task situations, a variety of *initial states*. Some of these initial states are easy to solve; some are difficult. If a creature learns to avoid the difficult situations, if it learns to refrain from attempting a task when it is hardest, it can ensure that certain *regions* of the complete state space are never visited and so are effectively pruned from the space. This has the effect of reducing the worst-case performance of its behavior routines and hence the average complexity of the task it actually attempts.

The "just say no" method bears closer scrutiny. Virtually every problem has difficult and easy instances in the sense that every algorithm designed to solve that problem will do worse on harder instances than on easier ones. The average complexity of a problem is given by the distribution of these difficult and easy cases (see Fig. 1). Although in general it is not possible to identify in advance which cases are going to be difficult and which are going to be easy, in particular problems there often are cues that signal difficult cases and, of course, a creature may remember that a particular case is hard if it has tried it once before. As the creature learns more about the cases that are difficult, it can create a personal filter that effectively reshapes the distribution of cases it will face. Without physically changing the environment, it changes its task environment.

In real life, this strategy is intuitive, pervasive, and population-wide, although we also expect to find idiosyncratic instances of it, were we to observe individuals closely. If a beta male finds itself threatened by an alpha, it is likely to growl, then avert its eyes and avoid a confrontation. Because it lacks the strength and fighting skills (fighting algorithms) to survive the harshest competition, it adopts a policy of retreat. Given the less daunting environment defined by the world to which it retreats, its fighting algorithms may suffice. This means, of course, that it will never have access to several females; but *ex hypothesis* we are assuming that in a competition it would lose and potentially incur serious injury.

Prudent retreat is not restricted to aggressive contexts. In looking for a stream to dam, a beaver judges whether the environment will prove hospitable and yield to its construction algorithm or whether it is likely to be a difficult case. If it seems that the engineering skills required are beyond its level, the beaver may simply reject the site and search for a more amenable one. The same applies to nesting birds. Sites for nests are chosen not just for their camouflage value or for their protective features but also for their affordances for nest building. If one local area does not provide hospitable sites, it is best to continue the search for better sites. A second-rate site might be adequate for a first-rate nest if a bird has the requisite architectural skills, but if it does not, the easiest way to live within its means is to reject second-rate sites. In a similar way, predators select their prey carefully, usually stalking them until the "initial conditions" of the attack phase are in a good region of their attack algorithm. If it is raining or stalking conditions are bad, rather than testing the excellence of its stalking algorithms, a predator may simply postpone hunting and wait for better conditions. Selectivity is an important component of competence. Even if a creature cannot guarantee that it will never have to face worst-case scenarios, it can strongly bias the distribution of cases and so improve the average case scenario it confronts.

A strictly human example of the "just say no" method of altering the task environment can be found in chess playing. It sometimes is maintained that chess players inhabit the same state space whether they are novices or experts. If the state space is determined by the rules of the game, it is difficult to see how it could be otherwise. Yet novice players never face most of the states that experts face. One reason is that they themselves are incapable of making the sort of moves necessary to reach certain states. This is a consequence of the competitive nature of chess. Yet another, more salient reason for us now is the way chess partners *self-select*. Players prefer to play players of comparable rank. Hence, novices face a slightly different chess environment than do experts because they never face states in which only experts could put them.

A final example of the "just say no" strategy can be found in the recent simulation studies on the iterated prisoner's dilemma by Batali and Kitcher (1996). These investigators found that when players are allowed a third action of waiting out a turn in the prisoner's dilemma task, they are able to achieve better aggregate outcomes. The average outcome for an agent playing the iterated prisoner's dilemma is greater if that player along with his or her coparticipants can abstain from play whenever he or she likes instead of having to make one of the two classic moves inside the game. By devising a policy of saying no under certain circumstances, players change their opponents' expectations of the costs and benefits of actions, leading everyone to choose actions with greater social welfare—a classic example of the "just say no" strategy, applied in a game theoretical context.

This then is the "just say no" strategy for environment change. It is not always easy to decide when it is being followed, as it so often looks like an integral part of a creature's fighting, building, or hunting algorithm. However, conceptually it is distinct. Indeed, if our analysis has been correct so far, it must be distinct from the task algorithm itself as such actions are task-external and often superoptimal: That is, when they do not have serious downstream effects, such as preventing learning (we are assuming the creature has completed its learning phase), such actions are ways the creature has of increasing the average payoff for one of its strategies. This leads to an increase in overall fitness.

To determine why actions such as avoiding confrontation, sitting on the sidelines, or looking for more hospitable sites are task-external actions and to determine that they can at times be superoptimal, we need to remind ourselves of how we decide on the action repertoire of a task. As was mentioned earlier, it is not always obvious what actions should be regarded as part of a task and what actions should be regarded as external to a task because it is not always clear whether a given action may, in principle, advance or hinder progress toward a goal. It may seem that selective acceptance of initial conditions is just such a difficult case. My reason for treating selective acceptance actions as external is that entering and leaving a task are not literally part of a task; they represent decisions at another level. If we were to admit that "just say no" actions were part of hunting or fighting or nest building, we would have to admit a range of considerations into these capacities that really have nothing to do with them. For instance, the decision to accept a nesting site must be taken in light of the availability of sites, the lateness of the year, the prevalence of predators, and so forth. All these considerations have nothing to do with the job of building, which can be treated as a fairly modular skill. Hence, "just say no" actions fall outside that modular skill. They are actions that affect the structure of a task without changing the state of the task; they do not cause transitions in that state space.

Showing that "just say no" actions often are superoptimal is more difficult. An action is superoptimal, by our definition, if it does not lower the return to actions in other tasks and it raises the expected value of activity in the current task, or else if any lowering of returns it causes in other tasks is more than offset by gains in this one. Because "just say no" actions do not require that the creature do anything other than reject the opportunity of engaging in a task, we must show that there are times when doing nothing is the best thing a creature can do. To appreciate why sitting on the sidelines can increase the yield of the current task without decreasing the yield in any other task, we need to recognize that there usually are *entry costs* to shifting from one task to another. The opportunity cost of doing nothing depends not only on the return that might be available were the creature to do something else with its time; it also depends on how easy it is to leave one task and begin another. Assuming that a creature cannot instantaneously change gears, there are going to be moments when the best thing a creature can do is to be idle, particularly if by idling the creature increases the expected return of time spent on the same task later. The upshot is that by invoking a strategy of waiting, or deferring action, a creature can improve substantially its expected yield from its current skills. Without having physically altered the properties of its habitat, it has altered the structure of one of the tasks it faces in that habitat and so has improved its prospects for surviving.

### 4.2 Routine maintenance

Learning to say no selectively to certain problem situations is one way of filtering out choice points to improve performance. Another way of achieving a similar outcome is to have a policy of maintaining one's environment so that undesirable choice points (i.e., states on which one must act) rarely arise. This is a more active policy that actually alters physical attributes of the habitat. Proverbs such as "a stitch in time saves nine," "an ounce of prevention is worth a pound of cure," and "scatter the stones before they make a pile" (Lao-Tze, 1898) all reflect the idea that preventive measures cost little when compared to the costs they save later. They often are easier to perform too. It requires less knowledge and effort to do routine maintenance on a car than to fix it when it breaks.

The value of maintenance strategies was discussed in Hammond, Converse, and Grass (1995) in the context of activity management and planning. These researchers noted that because certain resources tend to reside in specific places (e.g., clean glasses, crockery, cutlery, and cleaning equipment typically are found in cupboards, drawers, and closets), agents learn to count on these resources being in their appointed location. This a useful feature, for it means that if one is cooking or setting the table,

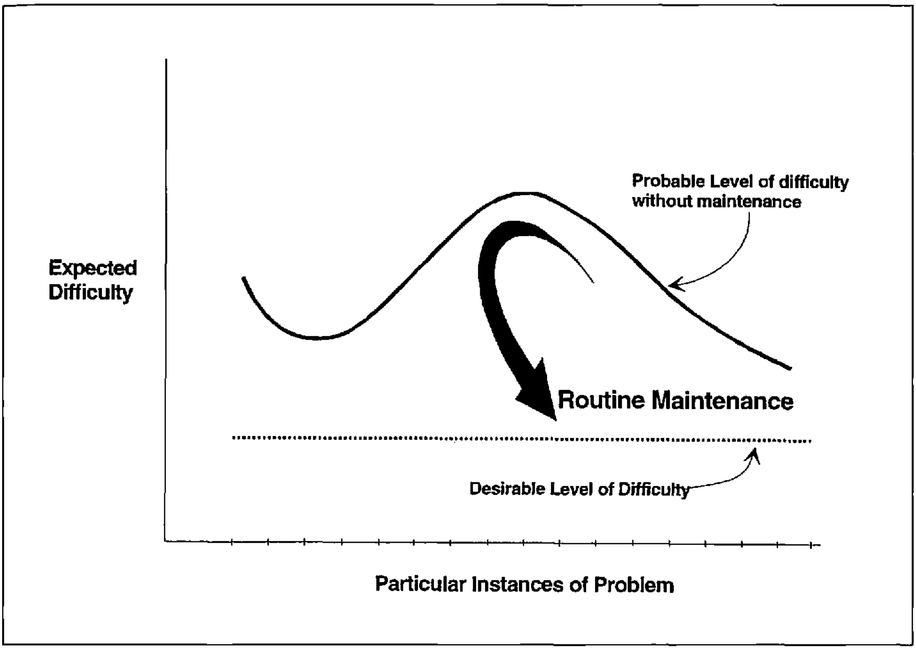


**Figure 2**
Without routine maintenance, environments have a probability of moving into states that create extra work. By maintaining its task environment, a creature reduces the probability that awkward cases will arise, thereby reducing the expected average difficulty of cases.

it will not be necessary systematically to search the work space to find the items one requires. However, such savings do not happen magically. Entropy teaches us that objects tend to scatter, so a plan or program that depends on resources being in their expected places is apt to fail unless someone in the agent's environment ensures that items find their way back to their proper place. A certain state of the environment must be maintained or enforced, by the agent itself, by some automatic mechanism, or by other members of the agent's group.

We can generalize the notion of resource maintenance. To begin, we note that environments have a probability of moving into states that are undesirable for a creature (Fig. 2). These states vary from the manageable but difficult to the impossible. In a world designed to make life easy for a creature, they would not arise at all, but inevitably they do arise because often they are side effects of the very actions taken by the creature. Thus, eating leads to digestion leads to defecation, which leads to soiling one's immediate locale—unless, of course, the odious result is buried, as domestic cats do, or the creature leaves the immediate locale to defecate, as most creatures do their nests. Related actions are taking the garbage out of the burrow.

Squirrels are known to remove rotting vegetation and, occasionally, the empty shells of their nuts from their burrows. Undesirable states also arise because of exogenous factors. Winter snows bury nuts that otherwise could be counted on to be found. Other animals gather and consume nuts. The net result is that in wintertime the probability of finding a nut just in time is too low to be reliable. It is better to pay the storage costs and the up-front labor costs to build an in-house inventory. Hence, storing nuts for winter is another instance of resource maintenance.

Given that resources are constantly in demand and not themselves always well defined, it is difficult to know the extent to which creatures practice resource maintenance. Involved in the general notion of maintenance are keeping parts or systems (or livestock, or cultivated goods) in *good working order* and ensuring that resources are located *where* they are most convenient and *when* they are most needed and, most complex of all, in *a form that is most useful.* Thus, to address the last condition alone, tidying up a work space may involve putting items in canonical locations but equally might involve redistributing the items in the same work space to remove clutter. No wonder it is hard to know when an action has beneficial resource management consequences.

Two general principles can be invoked to explain the virtues of maintenance: investment principles and artificial intelligence principles. According to standard investment theory, the decision about whether to invest time doing $A$ right now, when one is not likely to need the results of $A$ until later, depends on whether the expected benefits to be collected some time in the future outweigh the current costs multiplied by some interest factor to compensate for risk.[3] If there is no interest rate, there is no reason to prefer near-term returns on investment (instant gratification) to long-term returns (deferred gratification). Storing nuts for winter makes good investment sense because, effectively, the price of nuts goes up enough in winter (when they will be scarce and require more labor to be found) to compensate for present efforts. Accordingly, the current value of doing $A$ now exceeds the current value of anything else that might be done now. This is the simple story. However, it makes the idealizing assumption that the cost of labor is constant and arbitrarily divisible, when in fact we know that some requests on our time are more urgent than others, and we cannot always break from what we are doing to engage in other, more profitable, activities. Hence, we cannot suppose that we will always be able to do $A$ at the last minute or at least not without incurring potentially exorbitant costs. This complicates the equation, for it means that there is risk not only on the benefits side but also on the labor (i.e., cost) side. The upshot is that on economic principles, not

3 *Risk* here means that if I had chosen to do something else with my time and labor right now, I could have immediate gratification, whereas there is uncertainty about whether I will actually use the results of $A$ later, whether the changes I cause by doing $A$ now will still be present later, and whether I will even be alive later to enjoy the benefits.

only should a creature engage in maintenance actions whenever it has spare time (so doing $A$ now, when labor is cheap, is more profitable than doing anything else), but also it should engage in more maintenance actions if it is *uncertain* of when it might next get the chance.

Maintenance can also be justified on artificial intelligence principles. The reasoning is as follows: Expertise consists in knowing what to do in particular cases. The more an agent practices a skill, the broader the range of cases it learns to handle and the more compiled or chunked its skills become (Newell, 1990). The net effect is faster performance on a wider set of examples. Because the full space of possible situations is large, however, there are always cases that are novel but that will be treated as identical to known cases with possibly unfortunate consequences, or else they will require on-line adaptation (reasoning). The virtue of maintenance is that it can bias the probability against the occurrence of these novel or unfamiliar cases. We know what to do in accustomed cases, so the more cases that are familiar, the better our performance. If maintenance can cause this biased sample, it is adaptive.

For both these reasons, routine maintenance is an effective means of shaping the environment to help a creature circumvent performance-limiting circumstances and so increase average yield on its actions. It is an important superoptimizing strategy.

### 4.3 Scouting ahead

The strategies of "just say no" and routine maintenance are both key ways creatures have of taking back some degree of control over their environments. Because of such task-external strategies, creatures do not have to be passive optimizers always striving to do their best in games posed by nature; they can be active participants, modifying the rules of the games they must play. A third technique for changing the terms of a task is to scout ahead and change one's knowledge of the consequences of actions.

Imagine a lion stalking a herd of wildebeest. Should it attack from the right flank or should it circle around and attack from the left? If there is a hill nearby, a third action would be to defer attacking in order to secure extra or better information about the layout of the herd.

In decision theory, value of information analysis (Howard, 1966) provides a method for determining when it is worthwhile to pay the costs of acquiring information. As can be seen in Figure 3, the expected utilities of the actions available to a lion at the same physical location can be plotted on two separate occasions. On the first occasion, before any special information-seeking actions have been taken, the lion must operate with prior probabilities about the arrangement of large and small animals. On the basis of these prior probabilities, the clear choice is to attack

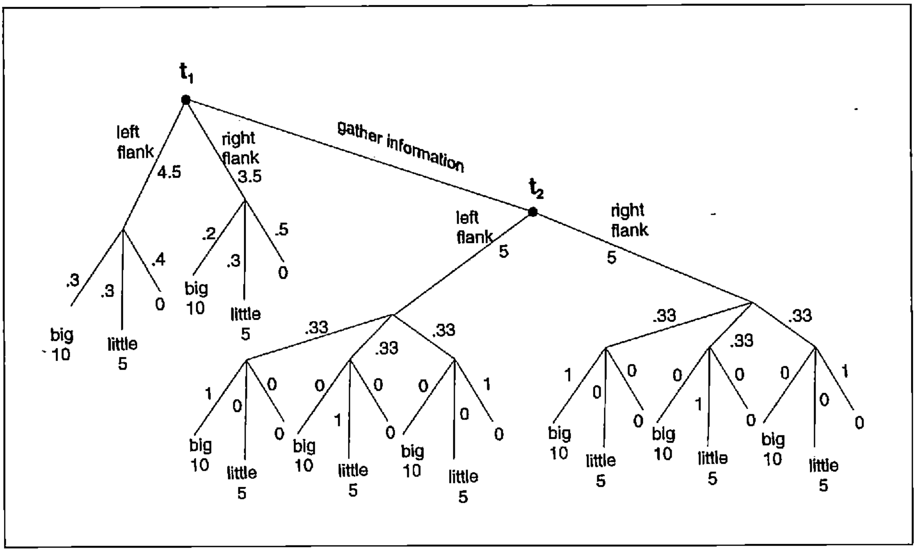


**Figure 3**
In this decision tree, we see the expected utilities of action at the same physical location on two occasions: at $t_1$ based on prior probabilities alone, and at $t_2$ after an information-gathering action has been taken. As long as the cost of information gathering does not exceed 0.5 units, it is rational to collect information before acting.

from the left flank, for the value of that attack is given by the sum of the expected utilities, which is 4.5. On the second occasion, the lion moves to a better viewing position and now has a commanding view of how this particular herd is distributed around the plain. Given this more informed idea of the organization of the herd, it is possible to know with very high confidence—virtual certainty, let us suppose—what the payoffs will be from an attack from each flank. Thus, once at $t_2$ the lion will know whether it will catch a large, small, or no animal by attacking from a particular flank. However, because neither the creature (nor us) can know at $t_1$ which situation will obtain at $t_2$, all that can be known is that once an information-gathering activity has been performed, the creature will be in one of three states of knowledge. It will know that it can secure a large animal or a small animal or no animal at all. Thus, at $t_1$ it knows that at $t_2$ there is a one-third chance of taking a large animal, a one-third chance of taking a small animal, and a one-third chance of taking no animal. Because the value of being at $t_2$, is 5 and the value of being at $t_1$ is 4.5, any information-gathering activity that costs less than 0.5 is worth undertaking.[4]

4 The values 0.33 for large, small, and no animal are by no means necessary. It is very likely that a further advantage of scouting ahead is that the lion sufficiently improves its knowledge of the layout of the herd that it reduces its chances of missing a kill. This would have the effect of changing the values to perhaps 0.4 large, 0.4 small, 0.2 no animal, thereby increasing the benefit of scouting ahead.

Scouting ahead is an interesting way of increasing control over an environment because we do not normally think of information-gathering actions as ways of altering a state space and, indeed, often they are not. Exploration might be an action *internal* to a state space, in which case any scouting action cannot alter the space, for the topology of that space will resemble that given in Figure 3. If exploration is part of a task, the action of moving to $t_2$ would not alter the state space; it would merely be a rational move within that space. This reiterates the view that state spaces are relational constructs, defined relative to an action repertoire. However, often exploration is *not* an action that is internal to a state space, in which case it is a way a creature has of altering the space.

The reason scouting ahead, as an external action, would have the effect of altering the state space is that a state space is not just a topology of connected states; it also contains a measure of the expected goodness of different states, and this measure assigns a distance between states. Any alteration in the distance between states counts as a change in the state space. An external action of scouting ahead has the effect of altering the distance between states.

If this seems odd, it may be because we are not used to regarding changes in expected utility as real changes in distance—hence, as real changes in a state space. An example that converts expected utility into expected travel time may help dispel this view. Every day when I leave the university, I must choose between two routes to travel home, the coastal route or the highway. Normally, the highway is faster (though less pleasant) and so, if I am in a rush, I take it. However, if I leave at rush hour, there is a good chance of a traffic jam on the highway route, right where there is a merge with another highway. Happily, it is possible to take a small detour and look out over the merge area. At rush hour, this action, though taking a few minutes, ends up saving me time on average because I am able then to take whichever route is faster. If I were always to take the coastal road at rush hour, my average travel time would be longer than if I take the coastal road only when there is a traffic jam on the highway. Similarly, if I were always to take the highway at rush hour, my average travel time would be longer than if I sometimes take the coastal road. By adding a scouting ahead action, I can reduce my average travel time and hence improve my average performance.

This ought to make clear that scouting ahead can alter the topology of a state space if it is an action external to the task, but why view scouting ahead as a behavioral routine that actually is external to such tasks as hunting (or driving home)? My argument is at bottom a slippery-slope argument: If we are to use the notion of task environment as an explanatory construct in understanding animal behavior, we need to draw a line between actions that are part of the task (part of hunting, in this case) and actions that are not. If we do not draw this line, then because information-gathering actions are so diverse and so difficult to determine whether they are relevant, we may be forced to include, as part of the core competence of hunting, actions that seem completely unconnected to hunting.

We can concretize this slippery-slope argument by returning to the example of a lion scaling a hill to look out over the plains. Clearly, scaling a hill is not part of the attack phase of hunting, which occurs on the plains. If scaling hills is part of hunting at all, it must be part of a more inclusive task of hunting—say, hunting and hunting-related scouting. Are there natural limits to this more inclusive task? I think not. Wherever we draw the boundaries of hunting-related scouting, there will always be new information-gathering strategies that have a measurable impact on hunting performance and yet that fall outside the task of hunting more broadly defined. For instance, who would call the action of going for a postkill exploratory walk a part of hunting? For one thing, it occurs after hunting. For another, it may involve walking to regions that are spatially distant from areas in which the lion ever, as a matter of fact, hunts. Yet a predator with good knowledge of the lay of the land, knowledge of where enclosures and open spaces are, often can use that knowledge in trapping animals during the hunting phase. As rats are known to engage in exploratory behavior unconnected to food search, other animals may be expected to as well, but then we face a dilemma: Either we must conclude that there are actions that are temporally and often spatially distant from the normal actions of a task but that nonetheless we must accept as part of the task (a conclusion that is tantamount to saying we do not know what is involved in performing the task) or we must accept that wherever we draw the boundary of a task's state space, it will be possible to find information-gathering actions that lie outside that boundary and that are capable of altering activity inside it. Sometimes these actions are superoptimal.

### 4.4 Creating new actions by using tools

Throughout our discussion, we have spoken often of the task-changing power of tools. Because a task environment is defined relative to an agent's action repertoire, any change in the actions the agent may perform changes the topology of the state space. Introducing a tool is one of the easiest ways to change an agent's action repertoire, for now it is possible to do things previously unattainable, or unattainable in a single step.

Take the case of New Caledonia crows, recently discussed by Gavid Hunt (1996). Many bird species use twigs, bits of bark and, in at least one case, cactus spines lying on the ground to aid them in their search for tasty insects and spiders. The New Caledonia crows, though, actually fashion the probes they use. One tool is made from twigs bitten from living trees. It serves as a hook. Another is a pointed probe, 20 cm long, made from the tough barbed leaves of the screw pine. In certain cases, such tools just make it easier or more reliable to fish out insects. More often, it would be impossible for the crows to search the interior of holes without the sticks. Owing to the different ways crows in different places use the tools, it is interesting to ponder how local the tool-making cultures are, but the virtue of the tools and, in part, the cultures that ensure the skills to craft them is that new actions are possible that alter the state space of insect search. These allow new and shorter paths to the goal of insect capture.

It may seem that delineating the effect of tool use on the structure of a task is as easy as rewriting the state space with a slightly changed connectivity and possibly a few extra states thrown in, but this is an oversimplification. In fact, tool use can change almost every facet of a task. Tool use may require any of the following modifications:

- Adding nodes to the state space as new things can now be done: For example, with a rock hammer, a chimp is able to crack harder nuts, so now the space of things that can be done increases.
- Reconceptualizing or refining the state space so that states and tasks once treated as unitary must now be differentiated: For instance, because of the role that rock hammers play in the activity of cracking nuts, chimps now can distinguish nuts that are crackable without a hammer from those crackable with a hammer from those that are totally uncrackable. Given the value of distinguishing these nuts, it becomes possible to define new tasks and actions, such as sorting nuts by these new categories, preparing oneself to use the tools, and the manifold activities associated with maintaining the tools.
- Adding new branches and changing the distance measures between states so that once-distant states now are reachable in fewer steps with the help of a tool: This reduces the shortest path from one state to another. For example, a tray may permit one to pick up several objects at once, thereby creating an action that behaves as a macro-operator, collapsing several actions into a single one.
- Changing the probability of reaching a state: For instance, chimp fishing poles increase the probability of securing termites, as do the manufactured probes used by the New Caledonia crows.

Given the profound impact that tool use can have on performance, it hardly needs justification as a strategy for superoptimization. Tool use and its cognate, coopting existing resources for new functions, are two of the most powerful ways of changing the state space of the task. With the act of introducing a tool, or with a behavior that imbues an existing resource with new functionality, a creature can leverage its existing capacities to new levels. Few superoptimizing strategies are as powerful.

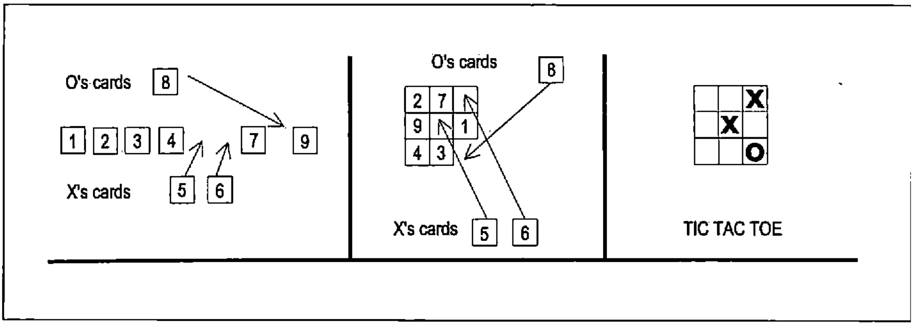


**Figure 4**
In the game of 15, there are two players. Players take turns selecting one card from the set on the table. The first to collect three cards that sum to 15 wins. The game of 15 and tic-tac-toe have isomorphic state spaces, as can be inferred from transforming the former into a tic-tac-toe-like game based on a magic square, as shown in the middle panel.

## 5 Strategies for Improving Cognitive Congeniality

Thus far we have considered actions that prune the state space of a task environment, that introduce shortcuts, and that guarantee that action trajectories will lie in hospitable regions. These actions increase the likelihood of task success or reduce its expected cost, measured in physical terms. There also is a class of task-external actions that leave the state space formally intact but that reduce the number and cost of mental operations needed for task success. These actions, which elsewhere I have called *epistemic* and *complementary* actions (Kirsh & Maglio, 1994; Kirsh, 1995a),[5] change the world in order to have useful cognitive effects on the *agent*. They reliably increase the speed, accuracy, or robustness of performance. They are yet another way a creature—usually a human creature—has of improving its fit with its environment.

Let us call the measure of how cognitively hospitable an environment is its *cognitive congeniality*. Different implementations of a state space have different degrees of cognitive congeniality. We can explore this notion using an old chestnut from the theory of human problem solving: According to Simon (1981a) tic-tac-toe and the game of 15 have isomorphic state spaces (Fig. 4), yet tic-tac-toe is a trivial game, whereas the game of 15 is not.

5 The expression *epistemic action* was chosen to describe certain types of actions that lead agents to epistemic states they might otherwise have reached by internal computation. In "Tetris," for example, Paul Maglio and I found that players seem to prefer to rotate "Tetris" pieces externally rather than rotate them mentally. They can more quickly reach states of knowing what a piece would look like rotated 90 or 180 degrees by rotating the piece in the world and observing the result than by performing the counterpart mental rotation. The expression *complementary action* was chosen to describe certain types of actions performed in the external world that are so timed that they interleave with mental actions as part of a more inclusive algorithmic strategy. For instance, we noted that subjects who count coins often use their fingers to point in a manner that complements visual routines (see the next section).

Imagine now that we are playing the game of 15 but we have been allowed to transform it to a magic square. On each turn, player X chooses a card and flips it over on its place in the square, whereas player O chooses a card and takes it off the board (see Fig. 4). Clearly, this new arrangement of 15 is cognitively more congenial than the first. It is not as congenial as tic-tac-toe itself, but it is a step in that direction. Like tic-tac-toe, the new arrangement allows us to extract much critical task information perceptually rather than by mental arithmetic and so it encodes needed information more *explicitly* (Kirsh, 1990) in the environment than in the original, linear arrangement. The result is that it saves mental computation, reduces cognitive load on working memory, and makes the game relatively easy without long hours of practice.

Using this notion of cognitive congeniality, we now can consider a range of natural environments to see how creatures, most especially humans, adapt those environments to make them more cognitively congenial. Let me emphasize that I am not suggesting that nonhuman animals engage in environmental reorganization in ways that seriously resemble the rather sophisticated methods we find among humans. Even among humans, gross rerepresentation of the sort occurring in tic-tac-toe and the game of 15 is rare outside of paper and pencil contexts, so it would be a surprise to find close analogs to external rerepresentation among animals. However, a fairly sizable class of actions exists that is less cognitive and also improves congeniality, which I will discuss. It is important to appreciate that even for animals, environments can be ranked along certain nonphysical dimensions, such as their cognitive congeniality, and that this ranking may show up in patterns of habitat selection where it is evident that creatures are exercising a selective function over the various environments through which they pass.

Study of the cognitive congeniality of an environment is still in its earliest stages. Questions such as: What is the maximum amount of working memory required to perform the task? What is the maximum amount of mental computation required to decide what to do next? How much task-relevant information is encoded in the environment, and how easily is it recovered? are central questions for the field. To date, there is no general theory to report (see Kirsh, 1995b). Nonetheless, it is apparent that cognitive congeniality is a key attribute in dozens of design fields, ranging from architecture and human computer interaction to product design and industrial engineering. As we learn more about animal cognition, it may well have an important role to play in ethology too. In the remainder of this article, I shall consider some of the techniques humans have for improving the cognitive congeniality of their environments. It is my belief that weak correlates to these can occasionally be found in the animal world.

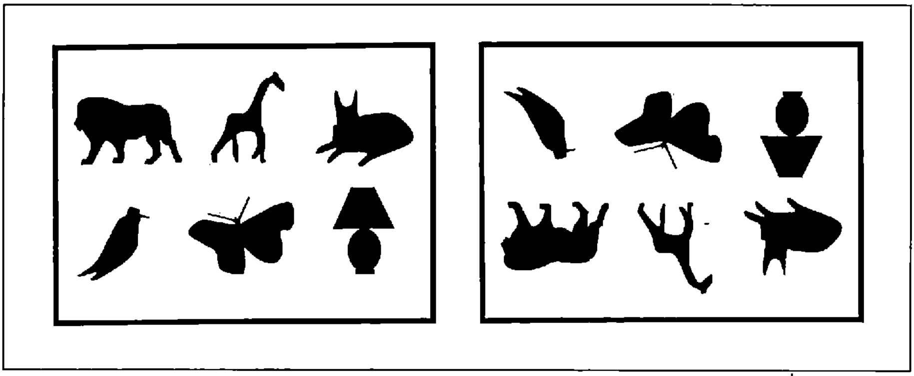

**Figure 5**
ZeeZee twice solved this wooden cutout puzzle in its normal orientation. Then she was asked to play with it upside down. After twice trying the wrong locations for pieces, she turned the puzzle right side up and quickly solved the remainder, as if by memory.

## 5.1 Complementary strategies

A complementary strategy is an interleaved sequence of physical and mental actions that results in a problem being solved—a computation being performed—in a more efficient way than if only the mental or physical actions alone are used. An example discussed in Kirsh (1995a) is pointing at coins while counting them. If subjects are asked to count, without either touching the coins or pointing to them, 30 coins consisting of nickels, dimes, and quarters, the result is that more than 50 percent of the time the subjects give the wrong answer. If they are allowed to point to the coins, that error rate is reduced to approximately 35 percent, and, if they are allowed to move the coins freely, the error rate falls to nearly 20 percent.[6] The hypothesis is that if an agent learns how to manipulate resources in its environment in a timely and constructive manner, it is able to solve cognitive tasks with less working memory, less visual spatial memory, less control of attention, or less visual search than would otherwise be required. Complementary actions are part of a strategy for restructuring the environment to improve the speed, accuracy, or robustness of cognitive processes.

Recently, I observed a good example of a complementary strategy in an 18-month-old child. As shown in Figure 5, ZeeZee was playing with a simple wooden puzzle. She had already played with this same puzzle twice before and, judging by the speed and accuracy with which she now assembled the puzzle, she had apparently memorized where each piece was to be placed. The little bird went in the lower left, the cat in the upper right, and the giraffe in the upper middle. I decided to test

6 This last result stems from pilot studies done in the Interactive Cognition Laboratory at the Department of Cognitive Science, University of California, San Diego.

how she would solve the same puzzle when it was rotated 180 degrees, so that pieces normally belonging in the bottom row now belonged upside down in the top row (see Fig. 5). Would she solve the puzzle this time by adapting her memory of where each piece went, would she fail to recognize the transformation and so use on-line reasoning to solve it as if from scratch, or would she do something different?.

Her first action was to try the little bird in the lower left corner, its customary place had the board been right side up. When the piece did not fit, ZeeZee looked the board over and tried placing the bird in the middle space along the bottom row, once again in the wrong orientation, as if now trying unsuccessfully to solve the puzzle but confused by her incorrect memory. She then went straight to the upper right slot (the correct position) and placed the piece in its appropriate orientation after a little effort. At this point, she turned the entire puzzle 180 degrees, returning it to its normal orientation, and quickly placed the remaining pieces in their proper (and apparently memorized) positions.

What type of action is this sort of board rotation? It is not a normal task-internal action, as no amount of reorientation can bring us closer or farther from the goal of having all the pieces in place. Nor is it a task-external action of the sort we have already described, as the topology of the state space remains unchanged: Any state accessible before rotation is accessible after rotation. Additionally, the number of placement actions and the physical energy required to place each piece in its position are identical whatever the puzzle's orientation, assuming arbitrary orientation of the pieces on the ground. Hence, there has been no change in the physical distance separating states. Evidently, it is a different type of action, an action that reorganizes the environment for mental rather than physical savings. By performing the metatask action of rotating the whole board, ZeeZee was able to stop her effortful on-line problem solving and return to her original rote strategy. This means that she could once again solve the puzzle using long-term memory rather than the working memory needed in on-line problem solving. She seems to have brought the world into conformity with her mental model rather than to adapt her mental model to accommodate upside down cases.

ZeeZee's board rotation is a simple example of a complementary action or strategy. Here is one that is slightly more complex than ZeeZee's. Once again, it is an interactive technique that saves mental rather than physical effort. In Figure 6, we see a scatter of 20 sticks of similar diameter but differing lengths. The agent's task is to identify the longest. Because longest is a globally defined property, we cannot be sure that a given stick is the longest without checking all the others, but here we are faced with a choice: Shall we move sticks as we visually check them, or shall we leave them untouched?

Many strategies are possible, but virtually all good ones involve moving the sticks.

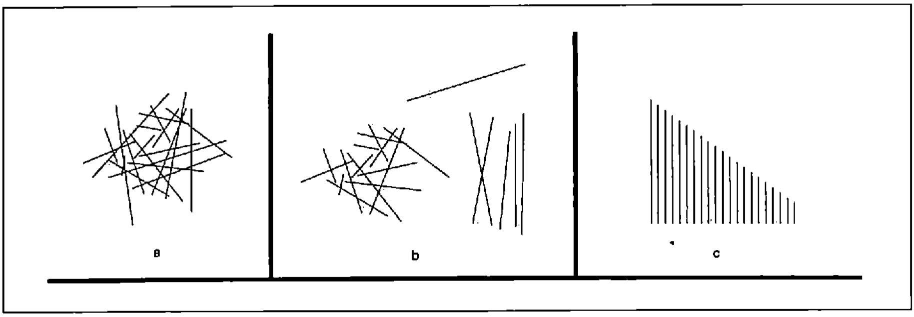


**Figure 6**
Subjects were asked to pick out the largest stick in (a). To solve the problem, subjects almost invariably move the sticks, as in (b). Part (c) shows how rearrangement can trivialize otherwise time-consuming global judgments, such as deciding whether all sticks are different in size.

For instance, one algorithm is to pick up the first two sticks, compare their length, keep the longer and discard the shorter into a reject pile, then continue comparing and discarding until all sticks in the resource pile have been checked. The stick we are left holding must be the longest.

Clearly, this is an algorithmically effective strategy. We know that we will check all and only the sticks we need to compare. Nonetheless, it runs longer than necessary. A better algorithm exploits our ability to make good guesses about the longest stick remaining in the resource pile. This time we pick up the two largest-looking sticks and discard the shorter until we see a major difference in length between the stick in our hand and the longest remaining stick in the resource pile. After a few sticks, we are certain. An even more efficient strategy is to grasp the three or four largest-looking sticks and then push their bottoms against the table. The one stick that pokes out farthest is the longest.

As theorists, what are we to make of all this activity? Should we regard these various actions of picking up candidate sticks and placing them in piles as actual moves within the state space of selecting the longest? Should we regard pushing the sticks against the table as a task-internal action? Rather, should we see them both as task-external actions? Because one might just visually scan the sticks, mentally comparing each, there is nothing intrinsic to the task of selecting the longest that requires either the sort of organizational behavior we observe in creating distinct piles or the analog computation of pushing the sticks against the tabletop. Nonetheless, sticks must be picked up and then put down. Therefore, intuitively, there is nothing task-external about actions that involve moving sticks around and placing them in piles, although the action of pushing them in groups against the table's surface does strain the intuition. The real problem is to decide how to describe these actions, for if we include only "objective" spatial elements in our action descriptions—so that actions of picking up are different if they involve sticks in different locations and orientations, as are actions of placing down—then the state space will be large enough to describe, *at a detailed level*, every individual sequence of actions that subjects display, but the level of description will be so fine that we will be unable to describe meaningfully the strategies the subjects are using. For instance, the interesting regularity at a strategic level is that subjects seem to make meaningful piles of sticks, or they intentionally exploit physical properties of the sticks and tabletop to perform an analog computation. Such regularities would be lost if we confine our descriptions of actions to straightforwardly physical characteristics of actions.

However, if we define the action repertoire using concepts that human subjects themselves use in describing the structure of their work space,[7] we create the problem that two subjects will have to be described as operating in different task environments if they conceptualize their environment in different ways. This is not acceptable. Either we must give up the notion of task environment as a useful explanatory construct or grant that there is a core task environment shared by all performers of a task, and that a variety of actions may be performed that do not fall within that task environment, narrowly construed, but that alter its cognitive congeniality. Once an environment has been so altered, strategies can be devised that exploit various of these nonintrinsic properties of the task environment. Thus, although stick-sorting tasks cannot be performed without shifting sticks about, there is no need to manage the spatial organization of the environment in the sense of partitioning it into discard regions, candidate regions, and the like. Such actions make the task easier to accomplish but are undertaken for their effects on the agent's understanding of the task and for the "cognitive affordances" they create rather than to make literal progress in the task. They are complementary actions, actions performed externally for their effects on internal computation.

In interpreting certain actions in this way, we are following a growing tradition of constructivism in learning theory (Duffy & Jonassen, 1992), and situated cognition in cognitive science (Hutchins, 1995; Norman, 1988; Suchman, 1986) by regarding many actions as having more to do with *cognitive scaffolding* than with step-by-step advancement to the goal. For instance, the point of creating a discard pile is to encode information about the state of our algorithm. It is to help us keep track of the sticks we have checked and the ones that remain. Indeed, the point of creating a

7 To confirm the idea that people conceptualize regions of their work space, we ran several subjects on the *select longest stick* task, not only recording their speed and accuracy, but also asking them after the task was over to describe what they were doing. All our subjects drew a distinction between discard areas and resource areas. In a small follow-up study, we further tested this idea by "accidentally" laying down a battery from the video camera on the discard area. This caused some of the sticks between the two piles to become mixed. As predicted, subjects reorganized the space to maintain the segregation of sticks.

discard pile can be understood only if we see the action as part of an algorithm that is being executed partly in the world and partly in the creature's head.

Assuming, then, that a class of actions exists that may improve task performance and that lies outside a state-space formalism—a class of actions that improves judgment, decision making, planning, and execution—all that remains is to show that such metatask actions are superoptimal. This is easy: Because the point of most complementary actions is to reduce cognitive loads and so improve performance, it is no surprise that without such actions judgments tend to be error-prone. In our pilot study, when subjects were not allowed to point, touch, or reorganize the sticks, and their task was to identify the longest stick of 20 distributed (as shown in Fig. 6a), they took more than 55 percent more time and made more than three times as many errors as when they were allowed to manipulate the sticks any way they liked. Similar results hold when subjects are required to pick up the largest stick but otherwise leave the arrangement intact. When the prohibition against movement was lifted subjects invariably relied on moving the sticks into piles, and their accuracy rose as indicated. This suggests that if accuracy is taken into account, the easiest way to improve performance is to allow certain task-external actions.

## 5.2 Actions that encode information externally

If we attempt to be more specific in the function of complementary actions, we very soon distinguish a large family of such actions that are concerned with externally encoding information about ongoing mental activity. The stick-sorting algorithm, for instance, uses the distinction between reject and resource piles to encode information about which sticks have been checked and which sticks have not. Having a reject pile lets us proceed in the algorithm without having to remember which sticks have already been checked.

A second example of this same pervasive activity can be found in card games. In ordinary games, such as gin rummy, bridge, and especially pinochle, it is common for players to organize their hands continually into arrangements that encode information about their game intentions. For instance, in Figure 7, we see four different ways of encoding the same set of cards dealt in a game of 14-card gin.[8] From a purely pragmatic viewpoint, there is no reason to group cards. Grouping has no effect on the goals you can reach, so from a task-environment perspective, each grouping designates the same state in the state space. Hence grouping must be a metaaction. Why do players bother? What function does it serve?

The simplest analysis of grouping is that it is done to encode plan fragments.

8 The goal of 14-card gin is to find two groupings of 4 cards and two groupings of 3 cards. Cards form an admissible grouping when they are 3 or 4 of a kind or 3 or 4 in a row of the same suit.

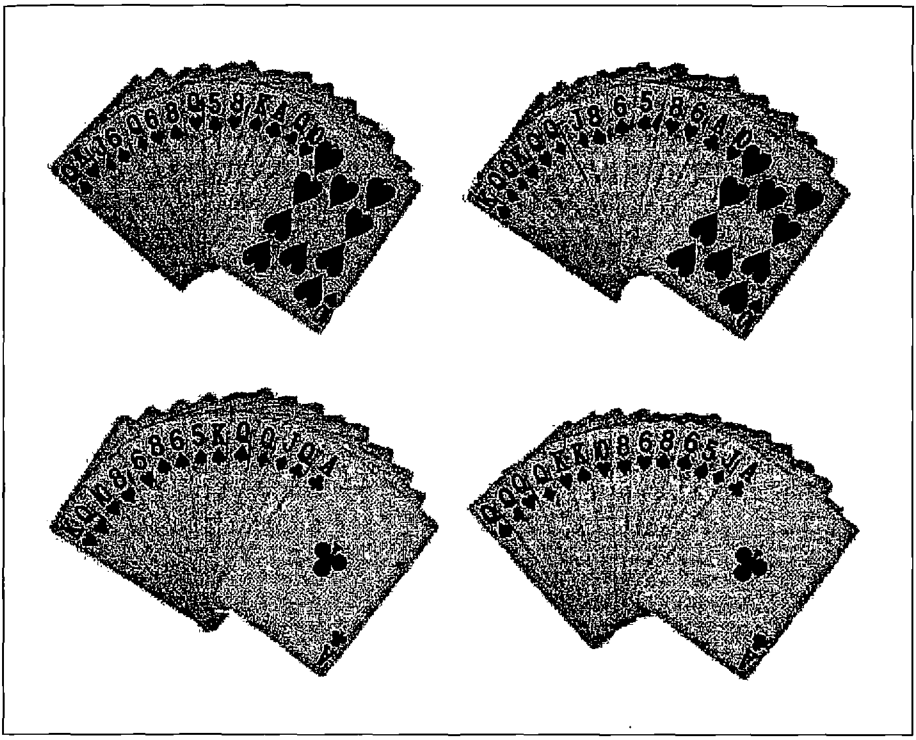

**Figure 7**
These four figures represent three different ways of organizing the same set of cards. Part (a) shows the cards as originally dealt. The other figures show the rearrangements made by three subjects.

One of the key cognitive tasks a card player faces is to sketch out a rough plan of the subgoals he or she will attempt to achieve. In gin rummy, for instance, as shown in Figure 7, there are many possible completions for which one can aim. An obvious explanation of the differences between the various organizations shown is that different players have chosen different strategies. Of course, we cannot be certain that we are right in our interpretation of what is encoded in each hand. We must guess at the encoding scheme each player is using or believe what each says when we ask. Moreover, players may make mistakes in encoding according to their own encoding scheme. Within these limits, characteristic of an interpretive science, we believe we can tell from the way the hands are laid out (and from the player's response to questioning) when a player has overlooked certain possible continuations that others have noticed. In Figure 7, player *c* has overlooked the fact that there are four queens, which *b* and *d* noticed.

The reason I am elaborating specific techniques we observe in card playing, counting sticks, and solving simple puzzles is that these actions are a central, if neglected, element of human activity. In card playing, it is obvious what the advantages of continual re-sorting are. Because game intentions are encoded externally, the player need not remember them; they can be read off the cards. This savings will reappear every time a change in intentions is made. Moreover, if the cards are well laid out, the time required to judge whether a target card can serve as a completion is faster than if the cards are poorly laid out. Players using good layouts also produce fewer errors,[9] leading one to suppose that effective goal encoding helps make execution of a plan more reliable and speedy. As long as we accept that card playing is not an unnatural task—a task unlike any we might be called on to perform in noncard contexts—we have reason to suspect that the kinds of complementary actions shown in card play have their counterpart throughout everyday life.

To sum up, my point in discussing interactive strategies typical of human-situated activity, is that there is a second family of strategies that agents possess for making their environments more hospitable. In addition to deforming the topology of their state spaces and hence the physical effort required to traverse states, these strategies may alter the cognitive properties of their environments and thus save mental effort. In some of the examples just mentioned, the method of changing the cognitive properties of environments was to redesign the appearance of the task sufficiently to change the complexity of the task. We know that how a problem is represented can have a major impact on the time and space required to solve the problem (Gigerenzer & Hoffrage, 1995). With a good representation, a problem may be easy to solve, requiring little search but, with a bad representation, the problem may be almost impossible to solve in that form and may require inordinate amounts of search, calculation, and recall of states. Once we view creatures as carrying out algorithms partly in their heads and partly in their environments, we must recognize that particular environmental layouts permit algorithms that lead to major savings. Even if these savings are not always evident in the time necessary to complete the task, they often will show up as significant improvements in accuracy, robustness, and reliability, all factors that matter to creatures.

## 6 Conclusion

I have been arguing throughout this article that organisms have two rather different ways of improving their fitness. The first, and most familiar, is by adapting themselves to their environments. In equilibrium, this leads to an optimal allocation of time between the different adaptive activities that make up an organism's life. The

9 This conjecture has been confirmed in pilot experiments carried out in our laboratory.

second way of improving fitness is by redesigning the environment to make it more hospitable: That is, organisms can adapt the environment to fit their existing skills and capacities. This leads to superoptimization for a creature already in equilibrium. Superoptimization requires the environment to be altered so that existing skills and techniques of survival yield greater returns in at least one task environment without sacrificing returns in others (Pareto optimization) or so that any lower returns in other task environments are more than compensated for by increases in others.

To implement an analysis based on superoptimization, it was necessary to introduce the notion of a task environment (drawn from the theory of problem solving) and the notion of a behavioral strategy interpreted as an algorithmic process. Two broad methods for improving algorithmic performance in a task environment were distinguished: deform the topology of a task's state space or improve the cognitive congeniality of a task's state space. Both have the effect of altering the average task complexity of performance.

It was suggested that the cognitive capacities needed to bring changes at the level of cognitive congeniality exceed those of most animals, although there may exist analogs to such environment-changing activity, which naturalists may discover once they have more detailed models of the cognitive processes underlying animal skills. Chief among these strategies are complementary actions. Complementary actions are those actions performed for the sake of simplifying computation. In humans, we find complementary actions everywhere. When a person uses his or her finger to point to a phone listing in the phone book, that individual is executing a complementary action because the use of one's hands saves one from having to remember the precise location of the target amid a set of distracters. Without the sort of interactive help that comes from manipulating environmental resources, including hands, many of the tasks we perform easily would be beyond our abilities.

Actions that deform the topology of a state space, by contrast with transformations of congeniality, are common in the animal kingdom and no doubt will be more widely appreciated once naturalists begin explicitly looking for them. Some are obvious. The actions of introducing a new tool and putting an existing object to new use, for instance, clearly are ones that may enhance the performance of existing strategies and permit variations that increase the yield of activity undertaken in that task environment. Less obvious strategies have to do with filtering out hard initial states of a task (the "just say no" strategy), maintaining the environment in a felicitous condition (routine maintenance), and undertaking exploratory actions (scouting ahead). The point of each of these strategies is to change the expected payoff of actions; it is to deform the task structure to make it easier or cheaper to complete the task successfully.

In the end, the value of analyzing activity along task analysis lines will depend on the insight it gives us into the principles that structure behavior. I have suggested that one useful approach is to distinguish actions that are internal to a task from actions that are external to it. These task-external actions are special only in the sense that as researchers we note that they force us to revise our models of behavior. Instead of assuming that most actions that occur in the time frame of a task are part of a strategy for solving the task, we may begin to consider whether some of those actions are external to the strategy, designed specifically to modify the task. If this proves to be a constructive way of looking at human and animal behavior, then evolution may select for both effective behavior control strategies and effective task redesign strategies.

**Acknowledgments**

It is a pleasure to thank Peter Todd for his excellent comments and generous efforts to help improve this article. I also thank Eric Welton, Rik Belew, and Matt Brown for their thoughtful commentary. This research is being funded by the National Institute of Aging, grant number AG11851.

## References

Batali, J., & Kitcher, P. (1995). Evolution of altruism in optional and compulsory games. *Journal of Theoretical Biology, 175*, 161–171.

Brandon, R. N. (1990). *Adaptation and environment.* Princeton, NJ: Princeton University Press.

Butler, D. R. (1995). *Zoogeomorphology: Animals as geomorphic agents.* Cambridge, UK: Cambridge University Press.

Diaper, D. (Ed.). (1989). *Task analysis for human computer interaction.* Chichester, UK: Ellis Horwood.

Duffy, T. M., & Jonassen, D. H. (Eds.). (1992). *Constructivism and the technology of instruction: A conversation.* Hillsdale, NJ: Lawrence Erlbaum Associates.

Elster, J. (1979). *Ulysses and the sirens: Studies in rationality and irrationality.* Cambridge, UK: Cambridge University Press.

Fisher, Sir R. A. (1930). *The genetical theory of natural selection.* Oxford, UK: Clarendon Press.

Gigerenzer, G., & Hoffrage, U. (1995). How to improve Bayesian reasoning without instruction: Frequency formats. *Psychological Review, 102*, 684–704.

Hammond, K. J., Converse, T. M., & Grass, J. W. (1995). The stabilization of environments. *Artificial Intelligence, 72*, 305–327.

Harel, D. (1987). *Algorithmics: The spirit of computation.* Reading, MA: Addison-Wesley.

Horn, H. S. (1979). Adaptation from the perspective of optimality. In A. T. Solbrig, S. Jain, G. B. Johnson & P. H. Raven (Eds.), *Topics in plant population biology.* New York: Columbia University Press.

Howard, R. A. (1966). Information value theory. *IEEE Transactions on Systems Science and Cybernetics, 2*, 22–26.

Hunt, G. R. (1996). Manufacture and use of hook-tools by New Caledonian crows [letter to editor]. *Nature, 379*, 249.
Hutchins, E. (1995). *Cognition in the wild.* Cambridge, MA: MIT Press.
Kirsh, D. (1990). When is information explicitly represented? In P. Hanson (Ed.), *Information, language, and cognition: The Vancouver studies in cognitive science* (vol 1). Oxford, UK: Oxford University Press.
Kirsh, D. (1995a). Complementary strategies: Why we use our hands when we think. In *Proceedings of the Seventeenth Annual Conference of the Cognitive Science Society.* Hillsdale, NJ: Lawrence Erlbaum Associates.
Kirsh, D. (1995b). The intelligent use of space. *Artificial Intelligence, 73*, 31–68.
Kirsh, D., & Maglio, P. (1994). On distinguishing epistemic from pragmatic actions. *Cognitive Science, 18*, 513–549.
Koza, J. R., Rice, J. P., & Roughgarden, J. (1992). Evolution of food-foraging strategies for the Caribbean Anolis lizard using genetic programming. *Adaptive Behavior, 1*(2), 171–199.
Kroemer, K. H. E. (1993). Fitting the workplace to the human and not vice versa (use of ergonomics in work area design). *Industrial Engineering, 25*(3), 56.
Kummar, H. (1995). *In quest of the sacred baboon: A scientist's journey.* (M. A. Biederman-Thorson, Trans.). Princeton, NJ: Princeton University Press.
Lao-Tze. (1898). *Tao-teh-king, Chinese-English.* (With introduction and transliteration and notes by Paul Carus.) Chicago: Open Court.
Larkin, J. H., & Simon, H. A. (1987). Why a diagram is (sometimes) worth ten thousand words. *Cognitive Science, 11*(1), 65–100.
Lewontin, R. C. (1978). Adaptation and natural selection. *Scientific American*, September.
McFarland, D., & Houston, A. (1981). *Quantitative ethology: The state space approach.* Boston: Pitman Advanced Publishing Program.
Newell, A. (1973). Production systems: Model of control structure. In W. C. Chase (Ed.), *Visual information processing.* New York: Academic Press.
Newell, A. (1990). *Unified theories of cognition.* Cambridge, MA: Harvard University Press.
Newell, A., & Simon, H. A. (1972). *Human problem solving.* Englewood Cliffs, NJ: Prentice-Hall.
Norman, D. A. (1988). *The psychology of everyday things.* New York: Basic Books.
Norman, D. A. (1991). Cognitive artifacts. In J. M. Carroll (Ed.), *Designing interaction: Psychology at the human-computer interface.* [Cambridge series on human-computer interaction, no. 4.] Cambridge, UK: Cambridge University Press.
Papadimitriou, C. (1995). *Computational complexity.* Reading, MA: Addison-Wesley.
Royama, T. (1992). *Analytical population dynamics.* London: Chapman & Hall.
Simon, H. (1955). *Behavior and the structure of the environment.* [Reprinted in H. Simon, *Models of thought* (vol 1), 1984.]
Simon, H. (1956). Rational choice and the structure of the environment. *Psychological Review, 63*:129–38. [Reprinted in H. Simon, *Models of thought* (vol 1), 1984.]
Simon, H. (1973). The structure of ill-structured problems. *Artificial Intelligence, 4*,

181–201.
Simon, H. (1981a). *The sciences of the artificial* (2nd ed). Cambridge, MA: MIT Press.
Simon, H. (1981b). The architecture of complexity. [Reprinted in *The sciences of the artificial* (2nd ed).] Cambridge, MA: MIT Press.
Simon, H. A. (1976). *Administrative behavior: A study of decision-making processes in administrative organization* (3rd ed). New York: Free Press.
Suchman, L. (1986). *Situated cognition.* Cambridge, UK: Cambridge University Press.
Teichner, W. H., & Olson, D. E. (1971). A preliminary theory of the effects of task and environmental factors on human performance. *Human Factors, 13*(4), 295–344.

## About the Author

**David Kirsh**

David Kirsh is currently an associate professor in the Department of Cognitive Science at University of California, San Diego. He received his BA from the University of Toronto in Economics and Philosophy and his DPhil from Oxford University in the philosophy subfaculty researching the Foundations of Cognitive Science. He then spent 5 years at the Artificial Intelligence Laboratory at the Massachusetts Institute of Technology as postdoctoral fellow and then a research scientist. In addition to an enduring interest in foundational issues, Dr. Kirsh is working on a theory of interactivity in natural and artificial environments.