Simple Inference Heuristics versus Complex Decision Machines

Introduction to Special Issue

PETER M. TODD
Max Planck Institute for Human Development, Center for Adaptive Behavior and Cognition, Lentzeallee 94, 14195 Berlin, Germany; E-mail: ptodd@mpib-berlin.mpg.de

Key words: complex algorithms, decision making, deep thought, fast and frugal decisions, simple heuristics

1. Introduction — The Virtues of Deep Thought and Shallow Surmise

Deep thought has long been considered a virtue. For centuries, the philosophers of antiquity have been seen as paragons of their civilizations; today, we similarly revere our judges and justices who ponder present facts and past precedents to reach wise and fair conclusions. Aristotle named reason and the practical application of intelligence as virtues, and Descartes held thought to be the very definition of existence (Sahakian, 1968); recently, active thinking has been proposed to *improve* existence, actually increasing health and longevity (Holden, 1998). The ability to deliberate and ponder one's actions is also often seen as a defining characteristic solely of *human* existence, distinguishing us from the more-reflexive animals. Even when this view is challenged, as it has been increasingly of late by primate research (e.g., Whiten and Byrne, 1998 — but see Heyes, 1998 for counter-arguments), the challenges serve primarily to increase our estimation and respect for the other lucky species shown to be capable of any intricate cogitation that rivals our own.

We deep-thinking humans have also striven to emulate deep thought in the machines we build, creating artificial intelligences and decision tools that mimic or enhance the human mind. Expert systems, data-mining software, artificial theorem provers, and chess-playing computers all perform elaborate calculations or process great amounts of information in an attempt to approach and sometimes exceed human decision-making power. These complex decision machines make up one focus of this special issue of *Minds and Machines*. And because the human mind is often viewed as a kind of computer itself (Gigerenzer and Goldstein, 1996b), many would include it in this complex-cogitator category as well.

But might there not be an alternative to all this heavy thinking? After all, humans are also renowned for making snap judgments, jumping to conclusions, and ignoring the evidence. Are we always wrong to do so, when we could ponder deeply instead? Recently, research has begun to suggest that simple inference mechanisms or heuristics can be surprisingly useful (e.g., Payne et al., 1993; Gigerenzer, Todd, and the ABC Research Group, 1999) — 'fast and frugal' thinking may be

Minds and Machines **9:** 461–477, 1999.

just as effective as long and serious consideration in many situations. Thus, simple heuristics are the second focus of this issue, and as they are the new decision-makers in town we are primarily interested here in giving them a chance to prove their worth in the struggle between deep thought and quick choice.

Of course, few would take the stated opposition of simple inference heuristics *versus* complex decision machines as an intractable either-or proposition.* Rather, the authors in this issue see both approaches as potentially viable means for reaching good decisions. The questions of interest lie more in determining when and where each approach is best used by minds, machines, or both working in combination.

In this introductory essay, I will first briefly sketch some aspects of the history of the debate between simple and complex approaches to decision making, before presenting further arguments for studying simple heuristics, and an overview of four classes of heuristics that are discussed at points throughout the papers in this issue. I will then show how these papers address some of the most important themes from the simple versus complex decision-making debate, and indicate some of the problems that remain to be tackled.

2. A Brief and Spotty History of Simple versus Complex Decision Making

In the early 1800s, Pierre Simon Laplace imagined the form that the ultimate decision maker would take:

> Given ... an intelligence which could comprehend all the forces by which nature is animated and the respective situation of the beings who compose it — an intelligence sufficiently vast to submit these data to analysis ... nothing would be uncertain and the future as the past, would be present to its eyes. (Laplace, 1814/1951, p. 4).

From the perspective of Laplace's superintelligence, Nature is deterministic and certain, and proper decisions can be arrived at through sufficient calculation; but for humans, Nature is fickle and uncertain, and inferences must be made on the basis of unreliable cues. Although omniscience and certainty are not attainable for any real system, the spirit of Laplace's superintelligence has survived nevertheless in the vision of unbounded rationality. This form of rationality is often defined by its adherence to the laws of logic and probability theory, and is the basis of the assumptions of proper economic behavior embodied in *Homo economicus*, as well as forming the foundations of the theories of planning and problem solving in traditional artificial intelligence and optimal (animal) behavior used in behavioral ecology (see Goodie et al., 1999). The lofty goals and strict standards of unbounded rationality, by their very definition, require the use of complex decision machines to process all available information, without regard for costs or limitations in time,

* However, de Garis (1999), for one, sees sinister possibilities in just this sort of opposition, envisioning a battle in the not-too-distant future between those intent on building ever-more-complex decision machines and those who value more organic thought processes.

processing power, or knowledge. Again, real systems cannot feasibly live up to these requirements, but the goal of unbounded rationality implies the more complex decision processing, the better.

To reach answers and make choices in a useful amount of time, real decision systems must employ *limited* information search, whereas models of unbounded rationality assume that search can go on indefinitely. In reasonable models, search must be limited because real decision makers have only a finite amount of time, knowledge, attention, or money to spend on a particular decision. Limited search requires a way to decide when to stop looking for information, that is, a *stopping rule*. One way to stop information search is to use a stopping rule that optimizes search with respect to the time, computation, money, and other resources being spent. More specifically, this 'optimization under constraints' vision of rationality holds that the decision maker should calculate the benefits and costs of searching for each further piece of information and stop search as soon as the costs outweigh the benefits (e.g., Sargent, 1993; Stigler, 1961). The rule 'stop search when costs outweigh benefits' sounds plausible at first glance. But a closer look reveals that optimization under constraints can require even more knowledge and computation than unbounded rationality, as all of the costs and benefits of all of the possible courses of searching or not searching for each further piece of information must be calculated (Vriend, 1996).

Of course, few would argue that real humans have the time and knowledge necessary to perform the massive computations required of either of these forms of rationality. Rather, people often use shortcuts or heuristics — simple mechanisms that allow decisions to be arrived at quickly and with little mental effort. But the visions of unbounded rationality and optimization under constraints are usually presented as ideals that human or machine reasoning *should* aspire to. The upshot of these aspirations is that they make real human (and animal, and machine) reasoning look flawed and irrational in comparison — and they thus cast a pall over the usefulness and viability of the simple heuristics that people and animals employ.

Much of the negative attitude toward reasoning based on simple heuristics emerged in psychology in the early 1970s, when the 'heuristics-and-biases' research program launched by Tversky and Kahneman (1974) tainted the idea of simple mental mechanisms by attaching them to the value-laden 'bias' term in a single inseparable phrase. Within this program, heuristics were invoked as the explanation whenever errors — mainly deviations from the laws of probability — were found in human reasoning. Although Tversky and Kahneman repeatedly asserted that heuristics sometimes succeed and sometimes fail, their experimental results were typically interpreted as indicating some kind of fallacy stemming from the use of heuristics.

Vindication of the usefulness of simple heuristics came a few years later from an unlikely source: robotics. Tired of the failures of the traditional 'good old-fashioned AI' approach to get robots to move more than a few feet across a room in an hour, some researchers gave up on attempting to have their robots build complete world-models and extensive plans following the dictums of standard theories of rationality. Instead, they tried building in hierarchies of simple rules to be followed (e.g., "go forward as long as path is not blocked"; "when path is blocked, back up and turn slightly"), and discovered that this approach yielded much faster and more robust behavior (Brooks, 1991; see also Braitenberg, 1984, for a discussion of how complexity is too readily attributed to the behavioral outcome of simple mechanisms, and Noble & Todd, in press, on the use of robot and simulation models to combat such false attributions when studying 'imitative' behaviors).

At about the same time, evidence began to grow that humans could use simple decision mechanisms to good advantage. Experts were shown to base their judgments on surprisingly few pieces of information (Shanteau, 1992). It was found that people could trade off the effort involved in making a choice against the accuracy of that choice, and choose a simple decision strategy that would achieve the desired balance (Payne et al., 1993). And simple heuristics that use only a single piece of information to make a choice between two alternatives were discovered to rival the performance of much more complex and information-hungry methods such as multiple linear regression (Gigerenzer and Goldstein, 1996a). Thus, rather than taking the use of heuristics as evidence that people can do little more than exhibit flawed instantiations of the perfect rationality they should aspire to, many now suspect that it may well be these aspirations that are flawed — reasoning, it seems, can be powerful and accurate without requiring unlimited time and knowledge.

3. How to Study Simple Decision Mechanisms

But just *how* can simple decision mechanisms make appropriate choices and inferences? What forms of bounded rationality — as opposed to unbounded rationality or optimization under constraints — work effectively, and why? The 'father' of bounded rationality, Herbert Simon, viewed his conception as having two interlocking components: the limitations of the (human) mind, and the structure of the environments in which the mind operates (see Gigerenzer and Selten, in press, for recent work following on these ideas). The first component of his vision means that models of (human) judgment and decision making should be built on what we actually know about the mind's capacities rather than on fictitious competencies. Because of the mind's limitations, humans "must use approximate methods to handle most tasks" (Simon, 1990, p. 6). These methods include recognition processes that largely obviate the need for further information search, heuristics that guide search and determine when it should end, and simple decision rules that make use of the information found.

The second component of Simon's view of bounded rationality, environmental structure, is of crucial importance because it can explain when and why simple heuristics perform well: if the structure of the heuristic is adapted to the structure of the information in the environment. Simon (1956) illustrated the importance of this component early on through an example of a hypothetical organism searching for food in a particular environment; others have made similar points both before his work (e.g., Brunswik, 1943) and at various times since (e.g., Anderson, 1990), including the extreme statement only the environment need be studied, not the mechanisms of the mind (e.g., Gibson, 1979). But in general the second part of Simon's (1956) paper title, "Rational choice and the structure of environments," has been neglected in mainstream cognitive sciences (even by Simon himself — see Simon, 1987).

To bring environmental structure back into the study of bounded rationality, we need a new focus on *ecological rationality*. Traditional definitions of rationality are concerned with maintaining internal order of beliefs and inferences. But real organisms spend most of their time dealing with the external disorder of their environment, trying the make the decisions that will allow them to survive and reproduce. To behave adaptively in the face of environmental challenges, organisms must be able to make inferences that are fast, frugal, and accurate. These real-world requirements lead to a new conception that proper reasoning must be ecologically rational, arising from decision mechanisms that are matched (that is, adapted) to the particular structure of information in the environments in which they are applied. The study of ecological rationality thus involves analyzing the structure of environments, the structure of heuristics, and the match between them.

We can now restate the questions beginning this section: How is ecological rationality possible? That is, how *can* fast and frugal heuristics work as well as they do, and escape the tradeoffs between different real-world criteria including speed and accuracy? The main reason for their success is that they make a tradeoff on another dimension: that of generality versus specificity. While internal criteria for the coherence of decisions are very general — logical consistency for instance can be applied to any domain — the correspondence criteria that measure a heuristic's performance against the real world require much more domain-specific solutions. What works to make quick and accurate inferences in one domain may well not work in another (see Bullock & Todd, this issue, for a discussion). Thus, different environments can have different specific fast and frugal heuristics that exploit their particular information structure to make adaptive decisions. But specificity can also be a danger: if a different heuristic were required for every slightly different decision-making environment, we would need an unworkable multitude of heuristics to reason with, and we would not be able to generalize to previously unencountered environments. Fast and frugal heuristics can avoid this trap by their very simplicity, which allows them to be robust in the face of environmental change and enables them to generalize well to new situations (see Martignon and Schmitt, this issue).

Robustness goes hand in hand with speed, accuracy, and especially information frugality. Simple heuristics can reduce overfitting (focusing too much on the specific details in a particular data set) by ignoring the noise inherent in many cues and looking instead for the 'swamping forces' reflected in the most important cues. Thus, simply using only one or a few of the most useful cues can automatically yield robustness. Furthermore, important cues are likely to remain important. The informative relationships in the environment are likely to hold true when the environment changes. Because of this pattern, fast and frugal heuristics that pay attention to systematic informative cues while overlooking more variable uninformative cues can ride out environmental change without suffering much decrement in performance. Thus, simple heuristics and environmental structure can both work hand in hand to provide a realistic alternative to the ideal of optimization, whether unbounded or constrained.

The research program proposed by Gigerenzer and colleagues (Gigerenzer et al., 1999) for studying the simple boundedly rational heuristics that humans and animals use involves (1) proposing and specifying computational models of candidate simple heuristics, (2) analyzing the environmental structures in which they perform well, (3) testing their performance in real-world environments, and (4) determining whether and when people really use these heuristics. The results of the investigatory stages 2 through 4 can be used to inform the initial theorizing of stage 1. The different stages of this research program rest on multiple methods, including theoretical modeling of heuristics, computer simulation of their performance, mathematical analysis of the fit between heuristics and specific environments, and laboratory experimentation. (When the goal is designing simple heuristics for artificial decision-making agents to use, rather than investigating what humans and animals actually use, the fourth step of this process can be omitted.)

3.1. Building simple heuristics from simple building blocks

To study particular heuristics in detail, computational models must be developed that specify the precise steps of information gathering and processing that are involved in generating a decision, allowing the heuristic to be instantiated as a computer program. In particular, simple fast and frugal heuristics are made up of building blocks that guide search for alternatives, information, or both, stop that search, and make a decision.

3.1.1. *Building Blocks for Guiding Search*

Decisions must be made between alternatives, and based on information about those alternatives. In different situations, those alternatives and pieces of information may need to be found through active search. The building blocks for guiding search, whether across alternatives or information, are what give search its direction (if it has one). For instance, search for cues can be simply random, or in order of some precomputed criterion related to their usefulness, or based on a recollection about which cues worked previously when making the same decision (see Section 4.2 below on one-reason decision mechanisms). Search for alternatives can similarly be random or ordered. Fast and frugal search-guiding principles do not use extensive computations or knowledge to figure out where to look next.

3.1.2. *Building Blocks for Stopping Search*

To respect the temporal limitations of the human mind (or that of any realistic decision-making agent), search for alternatives or information must be terminated at some point. Moreover, owing to the computational limitations of boundedly rational agents, the method for determining when to stop search should not be overly complicated. For example, one simple stopping rule is to cease searching for information and make a decision as soon as the first cue or reason that favors one alternative is found (as embodied in one-reason decision making — see Section 4.2). This and other cue-based stopping rules do not need to compute an optimal cost–benefit trade off as in optimization under constraints; in fact, they need not compute any costs or benefits at all. For search among alternatives, simple aspiration-level stopping rules can be used (see Section 4.4 below on satisficing search).

3.1.3. *Building Blocks for Decision Making*

Once search has been guided to find the appropriate alternatives or information and then been stopped, a final type of building block can be called upon to make the decision or inference based on the results of the search. These components can also be very simple and computationally bounded. For instance, a decision or inference can be based on only one cue or reason, whatever the total number of cues found during search (see Sections 4.1 on ignorance-based and 4.2 on one-reason decision mechanisms). Such single-cue decision making does not need to weight or combine cues, and so no common currency between cues need be determined. Decisions can also be made through a simple elimination process, in which alternatives are thrown out by successive cues until only one final choice remains (see Section 4.3 below on elimination heuristics).

These building blocks can be put together to form a variety of fast and frugal heuristics. Given that the mind is a biological rather than a purely logical entity, formed through a process of successive accrual, borrowing, and refinement of components, it seems reasonable to assume that new heuristics are built from the parts of old ones, rather than from scratch (Pinker, 1998). Following this assumption, two main methods can be used to construct computational models of fast and frugal heuristics: combining building blocks and nesting existing heuristics. Building blocks can be combined in multiple ways, though not arbitrarily: for instance, a fast and frugal heuristic for two-alternative choice that stops information search at the first cue on which the alternatives differ must also use a decision principle based on one-reason decision making. Whole fast and frugal heuristics can themselves be combined by nesting one inside another. As an example, the recognition heuristic (see Section 4.1) works on the basis of an elementary cognitive capacity, recognition memory, but it can also serve as the first step of one-reason decision heuristics that draw on other capacities, such as recall memory. Recognition memory develops earlier than recall memory both ontogenetically and evolutionarily, and the nesting of heuristics can similarly be seen as analogous to the addition of a new adaptation on top of an existing one.

3.2. FILLING THE MIND'S ADAPTIVE TOOLBOX

The various simple heuristics that are built up from building blocks and other nested heuristics can all be thought of as making up part of the *adaptive toolbox*: the collection of specialized cognitive mechanisms that evolution has built into the human mind for specific domains of inference and reasoning (Gigerenzer and Todd, 1999; see also Cosmides and Tooby, 1992; Payne et al., 1993). The adaptive toolbox contains all manner of psychological (as opposed to morphological or physiological) adaptations. These include so-called 'lower-order' perceptual and memory processes which can be fairly automatic, such as depth perception, auditory scene analysis, and face recognition, as well as 'higher-order' processes that are based on the 'lower' processes and can be at least partly accessible to consciousness. Within the class of higher-order mental processes fall fast and frugal heuristics for decision making, which themselves often call upon lower-order processes of cue perception and memory.

Lower-order perceptual and memory processes such as face and voice recognition are complex and difficult to unravel, in part because they make use of massively parallel computations. No one has yet managed to build a machine that recognizes faces as well as a 2-year-old child. Now consider a higher-order decision mechanism that makes inferences based on these processes, the recognition heuristic mentioned earlier (see also Section 4.1). This fast and frugal heuristic uses recognition to make rapid inferences about unknown aspects of the world: For instance, food whose taste one recognizes is probably safer than unrecognized food, and a university whose name one has heard of probably provides a more prestigious education than one whose name is unfamiliar. Although the mechanisms of recognition memory may be intricate and complex, the recognition heuristic can be described as an algorithm just a few steps long. We do not need to know precisely how recognition memory works to describe a heuristic that relies on recognition. This example illustrates an apparently paradoxical thesis: *higher-order cognitive mechanisms can often be modeled by simpler algorithms than can lower-order mechanisms.*

This thesis is not new, having been proposed in various forms over the past century (e.g., by proponents of the Würzburg school of psychology in the early 1900s — see Kusch, 1999). But it is central to the discussion of when we should postulate simple versus complex decision mechanisms in the adaptive toolbox. In particular, it offers us a way to evade the withering firepower of Morgan's Canon, the dictum that we must always prefer lower psychological processes as explanations over higher psychological processes. Morgan's Canon would have us explain any given observed choice behavior as the outcome of the lowest possible psychological mechanism, no matter how complex that mechanism might be. But armed only with Occam's Razor (see Forster, this issue, and Martignon and Schmitt, this issue), a preference for simplicity can overcome the barrage of Morgan's Canon and allow us to propose instead higher-order decision processes achieved via simple heuristics. (Fodor's Pop Gun, insisting that we always prefer a *higher*-order mechanism as an explanation over a lower-order one, can bring us to the same outcome — see Fodor, 1999; but as Fodor himself says, there is no principled reason to prefer his Pop Gun over Morgan's Canon, or vice versa, whereas there are many arguments for using Occam's Razor, as discussed in some of the papers in this issue.) We now tum to an overview of some of the kinds of high-order simple heuristics that can fill the mind's adaptive toolbox.

4. Four Families of Simple Heuristics

The decision-making building blocks just described can be put together to form classes or families of heuristics whose members are related by the particular search, stop, or decision rules they use. In this section I briefly introduce four such families of heuristics (out of many possible) covering decision situations that vary in the amount of information available, the number of options to choose between, and the distribution of options in time or space. These heuristics can be seen as models of the behavior of both living organisms and artificial systems. From a descriptive standpoint, they are intended to capture how real minds make decisions under constraints of limited time and knowledge. From an engineering standpoint, these heuristics suggest ways to build artificially intelligent systems — artificial decision-makers that are not paralyzed by the need for vast amounts of knowledge or for extensive computational power.

4.1. Ignorance-based decision mechanisms

One of the simplest forms of decision that can be made is to select one option from two possibilities, according to some criterion on which the two can be compared. What simple cognitive mechanisms can be used to make this type of decision? This will depend on the amount and type of information that is available in the environment. If the only information available is whether or not each possibility has ever been encountered before, then the decision maker can do little better than rely on his or her own partial ignorance, choosing either recognized options or unrecognized ones. For heuristics applicable to such situations, their information search building block merely specifies that recognition should be assessed for the alternatives being compared; the search stopping building block limits consideration to only this recognition information; and the decision building block indicates exactly how recognition information determines the final choice. This 'ignorance-based reasoning' is embodied in the recognition heuristic (Gigerenzer and Goldstein, 1996a; Goldstein and Gigerenzer, 1999), which uses the following decision rule: when choosing between two objects (according to some criterion), if one is recognized and the other is not, then select the former. For instance, Norway rats have evolved to behave according to a rule of this type, preferring to eat things they recognize through past experience with other rats (e.g., items they have smelled on the breath of others) over novel items (Galef, 1987).

Following the recognition heuristic will yield correct responses more often than would random choice in those decision environments in which exposure to different possibilities is positively correlated with their ranking along the decision criterion being used. Thus, the rats' food preference copying presumably evolved because the things that other rats have eaten (i.e., recognized items) are more often palatable than are random (unrecognized) items sampled from the environment. Such useable correlations are likely to be present for species with social information exchange where important environmental objects are communicated and unimportant ones are ignored, as well as for species in environments where important environmental objects are simply encountered more often or earlier in life.

4.2. ONE-REASON DECISION MECHANISMS

When multiple pieces of information are available (beyond only recognition), a second class of simple heuristics becomes applicable: those that rely on just a single cue to make a decision. Imagine that we again have two objects to compare on some criterion, and several cues that could be used to assess each object on the criterion. A one-reason decision heuristic could then work as follows: (1) select a cue dimension using some search building block and look for the corresponding cue values of each option; (2) compare the two options on their values for that cue dimension; (3) if they differ, then stop (this is the stop-search building block), and choose the option with the cue value indicating a greater value on the choice criterion (the decision building block); (4) if the options do not differ, then return to the beginning of this loop (step 1) to look for another cue dimension. Such a heuristic will often have to look up more than one cue before making a decision, but the simple stopping rule (in step 3) ensures that as few cues as possible will be sought, minimizing the time needed for information search. Furthermore, ultimately only a single cue will be used to determine the choice, minimizing the amount of computation that must be done.

To finish specifying a particular simple heuristic of this type, we must also determine exactly how cue dimensions are 'looked for' in step 1 — that is, we must pick a specific information search building block. For instance, the *Take The Best* heuristic searches for cues in the order of their ecological validity — that is, their correlation with the decision criterion, while the *Minimalist* heuristic selects cues in a random order (Gigerenzer and Goldstein, 1996a, 1999; see Martignon and Schmitt, this issue, and Bullock and Todd, this issue, for discussions and comparison with other decision mechanisms). Again, both stop their information search as soon as a cue is found that allows a decision to be made between the two options.

Despite (or often because of) their simplicity and disregard for most of the available information, these two fast and frugal heuristics can make very accurate choices. A set of 20 environments was collected to test the performance of these heuristics, varying in number of objects and number of available cues, and ranging in content from the population sizes of German cities to fish fertility to high-school dropout rates (Czerlinski et al., 1999). The decision accuracies of Take The Best and Minimalist were compared with those of two more traditional decision mechanisms that use all available information and combine it in more or less sophisticated ways: multiple regression, which weights and sums all cues in an optimal linear fashion, and Dawes's Rule, which counts up the positive and negative cues and subtracts the latter from the former. The two fast and frugal heuristics always came close to, and often exceeded, the performance of the traditional algorithms when all were tested on the data they were trained on (data fitting). This surprising performance on the part of Take The Best and Minimalist was achieved even though they only looked through a third of the cues on average (and only decided using one of them), while multiple regression and Dawes's Rule used them all. The advantages of simplicity grew in the more important test of generalization performance, where the decision mechanisms were assessed on a portion of each data set that they had not seen during training; in that case, Take The Best outperformed all three other algorithms by a clear margin. Thus, making good decisions need not rely on the standard rational approach of collecting all available information and combining it according to the relative importance of each cue — simply betting on one good reason, even one selected at random, can provide a competitive level of accuracy in a variety of environments.

4.3. Elimination heuristics for multiple-option choices

These results on the efficacy of simple heuristics in structured environments are not restricted to decisions made between two objects. More generally, fast and frugal heuristics can also be found that use as few cues as possible to categorize objects. *Categorization by Elimination* (Berretty et al., 1997), similar to Tversky's (1972) *Elimination by Aspects* model of preference-based choices upon which it is based, uses one cue after another in a particular order to narrow down the set of remaining possible categories until only a single one remains. When cues are ordered in terms of their usefulness in predicting the environment, accurate categorization can be achieved using only the first few of the available cues. Even more accurate performance can arise when categorization is based on only a single cue through the use of a fine-grained cue-value-to-category map (Holte, 1993), provided the decision maker is able to make the tradeoff required between extra memory and increased accuracy.

Estimation can also be performed accurately by a simple algorithm that exploits environments with a particular structure. The *QuickEst* heuristic (Hertwig et al., 1999) is designed to estimate the values of objects along some criterion while using as little information as possible. To estimate the criterion value of a particular object, the heuristic looks through the available cues or features in a criterion-determined order, until it comes to the first one that the object does not possess. At this point QuickEst stops searching for any further information and produces an estimate based on criterion values associated with the absence of the last cue. QuickEst proves to be fast and frugal, as well as accurate, in environments characterized by a distribution of criterion values in which small values are common and big values are rare (a so-called 'J-shaped' distribution). Such distributions characterize a variety of naturally occurring phenomena including many formed by accretionary growth (e.g. cities, some businesses, etc.).

4.4. Satisficing heuristics for sequential choices

The heuristics presented so far assume that all of the possible options to be chosen between are presently available to the decision maker. But a different strategy is called for when alternatives (as opposed to information about the alternatives) take time to find, appearing sequentially over an extended period or spatial region. In this type of choice task, a fast and frugal reasoner need not (only) limit information search, but (also) must have a stopping rule for ending the search for alternatives themselves. One instance of this type of problem is the challenge that faces individuals searching for a mate from a stream of potential candidates met at different points in time; another instance is the search for amenable habitats in which to settle. Here, Herbert Simon's (1955, 1990) notion of a *satisficing* heuristic is applicable: An aspiration level is set for the selection criterion being used, and the search for alternatives is stopped as soon as the aspiration is met. Simple mechanisms can be used to set the aspiration level in the first place, such as checking the first few alternatives and taking the best value seen in that set as the level to beat in further search (Todd, 1997; Todd and Miller, 1999; see Goodrich, Stirling and Boer, this issue, for more on satisficing approaches). The trick here is to balance the desire for a short, fast and frugal search on the one hand (achieved by checking as few initial alternatives as possible), against the need for enough information about the potential alternatives to set an appropriate aspiration level on the other hand (achieved by checking as many initial alternatives as possible).

5. Extending the Study of Simple Heuristics

5.1. Extensions presented in this issue

The classes of simple heuristics just described represent some of the initial landmarks that have been explored in the study of simple inference heuristics in comparison with complex decision machines. The papers in this special issue of *Minds and Machines* extend this exploration, focusing on the importance of environment structure, the way that simplicity can lead to robustness, and how simple strategies can do useful work. First, Terry Connolly discusses the kinds of environments and tasks in which simple 'hedge-trimming' decision strategies that act quickly on the basis of little information will be more appropriate than complex 'tree-felling' strategies that gather masses of evidence and ponder it thoroughly before deciding. When information is costly or difficult to obtain and incremental actions can make useful progress, the former approach is called for; when decisions and actions have an 'all-or-nothing' character and their outcomes are crucial and irreversible, it makes sense to fell the tree carefully.

Seth Bullock and Peter Todd continue this investigation of the impact of decision importance or significance, as well as of decision prevalence or frequency, on the kinds of choice mechanisms that fare better or worse in different environments. When the frequency structure of an environment is skewed — making some choices confront a decision-maker much more often than others — certain heuristics show a particular sensitivity to this structure, and hence can exploit the environment for better performance. By skewing an environment's significance structure — making some decisions much more important than others, for example giving some wrong choices lethal consequences — more complex decision strategies, akin to Connolly's tree-felling approaches, begin to have an advantage.

In his paper on how simple mechanisms could ever fit the environment structures of a complex world, Malcolm Forster argues that this is only possible because a degree of complexity is hidden 'under the hood' of the simple mechanisms. The structure of any given simple mechanism must have built into it some assumptions about the structure of the problems to which that mechanism will be applied, making it domain-specific to some extent. Making the wrong assumptions can seriously hinder the performance of an algorithm; but building in the right structure can yield simple mechanisms that generalize robustly to decisions across their particular domain.

This robustness from simplicity theme is taken up further by Laura Martignon and Michael Schmitt. They show how even a simple task — comparing two objects on the basis of a set of binary cues — can be tackled by a variety of decision mechanisms, including the one-reason heuristics described earlier, decision trees, Bayesian networks, and estimation algorithms. Within each of these approaches, a complex or optimal strategy can be compared with a simple 'fast and frugal' mechanism. Martignon and Schmitt demonstrate that the performance of the simple mechanisms typically approaches, and can often exceed, that of the more complex strategies, particularly when robustly generalizing to new data.

Kathryn Laskey, Bruce D'Ambrosio, Tod Levitt, and Suzanne Mahoney present the flip side of Forster's position on simplicity built on complexity, arguing that complex decision mechanisms can only work if grounded in simple reactive inference processes. They present a complex decision machine that they have built to help military commanders assess the overwhelming masses of information that flood in from battlegrounds; this system relies crucially on the initial use of simple quick inferences to distill that flood. Thus, simple heuristics enable complex information processing machines to produce overall assessments that are once again simple, in a form most useful for the minds of the people making the ultimate decisions.

Finally, Michael Goodrich, Wynn Stirling, and Erwin Boer consider how a decision-making agent can, at a higher meta-level, choose between available simple and complex decision strategies. The idea of satisficing described earlier can be applied to such situations, making meta-choices that resolve the tradeoffs between costs and benefits of different mechanisms without the meta-level processing itself becoming too complex. This approach is demonstrated in a setting in which human drivers must decide when to take over control of their moving vehicle from an autonomous speed-regulating device. Thus, Goodrich, Stirling, and Boer show another way in which simple heuristics and complex strategies, as well as the human minds and artificial machines in which they are usually embodied, can work together effectively.

5.2. Further Advances

Beyond the initial landmarks and the regions around them explored in this issue, there are of course many areas still wide open for exploration in the territorial dispute between simple inference heuristics and complex decision machines. After all, much of this territory has only recently been opened up. Each of the papers in this issue indicates profitable directions for advancing on current work; here I summarize a few of the most pressing topics, clustered into three main questions.

5.2.1. *When are Simple Heuristics or Complex Mechanisms Appropriate?*

Researchers have begun to identify particular simple heuristics for a variety of restricted decision tasks, but how far ‘up the cognitive ladder’ of higher-order processes can this search be extended? For example, recent work in behavioral ecology and behavioral robotics points toward simple possibilities for tasks that are extended in time, such as planning or navigation (Goodie et al., 1999). What other forms of reasoning can heuristics be applied to? Are there features of environment structure that determine when simple or complex mechanisms are needed and which particular mechanism is best suited? We do not have yet a well-developed language for describing those aspects of environment structure, whether physical or social, that determine the usefulness and shape the design of decision heuristics. In studying this aspect of ecological rationality we should turn for inspiration to other fields, including ecology and statistics, that have analyzed environment structure from different perspectives.

5.2.2. *How is the Appropriate Simple Heuristic or Complex Algorithm Selected?*

That is, when more than one decision mechanism is applicable in a given situation, how can a choice between them be made? Must the adaptive toolbox be equipped with an adaptive mechanic who performs this selection? Will the choice be made on the basis of the performance of the various candidate mechanisms and the demands of the current situation (as proposed by Payne et al., 1993; see also Goodrich et al., this issue), and if so, what performance criteria (accuracy, speed, frugality, robustness, logical coherence, etc.) are to be used? Or can heuristics and other mechanisms be invoked in a more automatic, simple fashion? The appeal of a fast and frugal approach to cognition will certainly be decreased if this is not the case.

5.2.3. *Where do Simple Heuristics (and Other Decision Mechanisms) Come From?*

How can the processes of learning and evolution interact to create the particular heuristics that humans and animals employ? Are the outlines of specific heuristics evolved, and the particular cues used in particular environments then subject to learning? When is individual learning used, and when are new decision mechanisms (or new cues) picked up through social information exchange, including culture? How can we as researchers come up with appropriate heuristics to investigate in humans and animals (the first step in the four-step research process described in Section 3)? Finally, when we are designing artificial decision-making systems, what engineering processes can we use to find the most useful heuristics or complex mechanisms — artificial evolution, machine learning, Bayesian approaches, or some other methods?

Tackling these questions will give us further insights into the relationships between simple inference heuristics and complex decision mechanisms. In the meantime, as the papers in this issue attest, we are already converging on an image of simplicity built atop complex underpinnings, complexity emerging from the interactions of simple mechanisms, and virtuous deep thought working hand in hand with fast and frugal snap judgments to yield the decisions that shape our paths.

References

Anderson, J.R. (1990), *The Adaptive Character of Thought*, Hillsdale, NJ: Erlbaum.

Berretty, P.M., Todd, P.M. and Blythe, P.W. (1997), ‘Categorization by Elimination: A Fast and Frugal Approach to Categorization’, in M.G. Shafto and P. Langley, eds., *Proceedings of the Nineteenth Annual Conference of the Cognitive Science Society*, Mahwah, NJ: Lawrence Erlbaum Associates, pp. 43–48.

Braitenberg, V. (1984), *Vehicles: Experiments in Synthetic Psychology*, Cambridge, MA: MIT Press.

Brooks, R.A. (1991), ‘New Approaches to Robotics’. *Science* 253, pp. 1227–1232.

Brunswik, E. (1943), ‘Organismic Achievement and Environmental Probability’, *Psychological Review* 50, pp. 255–272.

Bullock, S. and Todd, P.M. (1999), ‘Made to measure: Ecological Rationality in Structured Environments’, *Minds and Machines* 9, pp. 497–542.

Connolly, T. (1999), ‘Action as a fast and frugal heuristic’, *Minds and Machines* 9, pp. 479–496.

Cosmides, L. and Tooby, J. (1992), ‘Cognitive adaptations for social exchange’, in J. Barkow, L. Cosmides and J. Tooby, eds., *The Adapted Mind: Evolutionary Psychology and the Generation of Culture*, New York: Oxford University Press, pp. 163–228.

Czerlinski, J., Gigerenzer, G. and Goldstein, D.G. (1999), 'How Good are Simple Heuristics?' in G. Gigerenzer, P.M. Todd and the ABC Research Group, *Simple Heuristics that make Us Smart*, New York: Oxford University Press, pp. 97–118.

de Garis, H. (1999), 'The Artilect War', Unpublished manuscript. Available at http://www.hip.atr.co.jp/∼degaris/artilectwar.html.

Fodor, J. (1999), 'Not So Clever Hans (review of *If a Lion could Talk: How Animals Think*, by Stephen Budiansky)', *London Review of Books*, 4 February 1999, pp. 12–13.

Forster, M.R. (1999), 'How do Simple Rules "Fit to Reality" in a Complex World?', *Minds and Machines* 9, pp. 543–564.

Galef, B.G., Jr. (1987), 'Social Influences on the Identification of Toxic Foods by Norway Rats', *Animal Learning & Behavior* 18, pp. 199–205.

Gibson, J.J. (1979), *The Ecological Approach to Visual Perception*. Boston: Houghton Mifflin.

Gigerenzer, G. and Goldstein, D.G. (1996a), 'Reasoning the Fast and Frugal Way: Models of Bounded Rationality'. *Psychological Review* 103, pp. 650–669.

Gigerenzer, G. and Goldstein, D.G. (1996b), 'Mind as Computer: The Birth of a Metaphor', *Creativity Research Journal* 9, pp. 131–144.

Gigerenzer, G. and Goldstein, D.G. (1999), 'Betting on One Good Reason: The Take The Best Heuristic', in G. Gigerenzer, P.M. Todd and the ABC Research Group, *Simple Heuristics that make Us Smart*, New York: Oxford University Press, pp. 75–95.

Gigerenzer, G. and Selten, R. (in press), *Bounded Rationality: The Adaptive Toolbox* (Dahlem Workshop Report), Cambridge, MA: MIT Press.

Gigerenzer, G. and Todd, P.M. (1999), 'Fast and Frugal Heuristics: The Adaptive Toolbox', in G. Gigerenzer, P.M. Todd and the ABC Research Group, *Simple Heuristics that make Us Smart*, New York: Oxford University Press, pp. 3–34.

Gigerenzer, G., Todd, P.M. and the ABC Research Group (1999), *Simple Heuristics that make Us Smart*, New York: Oxford University Press.

Goldstein, D.G. and Gigerenzer, G. (1999), 'The Recognition Heuristic: How Ignorance makes Us Smart', in G. Gigerenzer, P.M. Todd and the ABC Research Group, *Simple Heuristics that make Us Smart*, New York: Oxford University Press, pp. 37–58.

Goodie, A.S., Ortmann, A., Davis, J.N., Bullock, S. and Werner, G.M. (1999), 'Demons versus Heuristics in Artificial Intelligence, Behavioral Ecology, and Economics', in G. Gigerenzer, P.M. Todd and the ABC Research Group, *Simple Heuristics that make Us Smart*, New York: Oxford University Press, pp. 327–355.

Goodrich, M.A., Stirling, W.C. and Boer, E.R. (2000), 'Satisficing Revisited', *Minds and Machines* 10, pp. 27–57.

Hertwig, R., Hoffrage, U. and Martignon, L. (1999), 'Quick Estimation: Letting the Environment do the Work', in G. Gigerenzer, P.M. Todd and the ABC Research Group, *Simple Heuristics that make Us Smart*, New York: Oxford University Press, pp. 209–234.

Heyes, C.M. (1998), 'Theory of Mind in Nonhuman Primates', *Behavioral and Brain Sciences* 21, pp. 101–148.

Holden, C. (1998), 'The Nuns' Story', *Science* 280, p. 1199.

Holte, R.C. (1993), 'Very Simple Classification Rules Perform Well on Most Commonly Used Datasets', *Machine Learning* 3, (11), pp. 63–91.

Kusch, M. (1999), *Psychological Knowledge: A Social History and Philosophy*, London: Routledge.

Laplace, P.S. (1814/1951), *A Philosophical Essay on Probabilities* (F.W. Truscott and F.L. Emory, Trans.), New York: Dover. (Original work published 1814).

Laskey, K.B., D'Ambrosio, B., Levitt, T.S. and Mahoney, S. (2000), 'Limited Rationality in Action: Decision Support for Military Situation Assessment', *Minds and Machines* 10, pp. 1–25.

Martignon, L. and Schmitt, M. (1999), 'Simplicity and Robustness of Fast and Frugal Heuristics', *Minds and Machines* 9, pp. 565–593.

Noble, J. and Todd, P.M. (in press), ‘Imitation or Something Simpler? Modelling Simple Mechanisms for Social Information Processing’, in K. Dautenhahn and C. Nehaniv, eds., *Imitation in Animals and Artifacts*, Cambridge, MA: MIT Press.

Payne, J.W., Bettman, J.R. and Johnson, E.J. (1993), *The Adaptive Decision Maker*, New York: Cambridge University Press.

Pinker, S. (1998), *How the Mind Works*, New York: Norton.

Sahakian, W.S. (1968), *History of Philosophy*, New York: Harper Collins.

Sargent, T.J. (1993), *Bounded Rationality in Macroeconomics*, Oxford: Oxford University Press.

Shanteau, J. (1992), ‘How Much Information does an Expert Use? Is it Relevant?’ *Acta Psychologica* 81, pp. 75–86.

Simon, H.A. (1955), ‘A Behavioral Model of Rational Choice’, *Quarterly Journal of Economics* 69, pp. 99–118.

Simon, H.A. (1956), ‘Rational Choice and the Structure of Environments’, *Psychological Review* 63, pp. 129–138.

Simon, H.A. (1987), ‘Rational Decision Making in Business Organizations’, in L. Green and J.H. Kagel, eds., *Advances in Behavioral Economics. Vol. 1*, Norwood, N.J.: Ablex, pp. 18–47.

Simon, H.A. (1990), ‘Invariants of Human Behavior’, *Annual Review of Psychology* 41, pp. 1–19.

Stigler, G.J. (1961), ‘The Economics of Information’, *Journal of Political Economy* 69, pp. 213–225.

Todd, P.M. (1997), ‘Searching for the Next Best Mate’, in R. Conte, R. Hegselmann and P. Terna, eds., *Simulating Social Phenomena*, Berlin: Springer-Verlag, pp. 419–436.

Todd, P.M. and Miller, G.F. (1999), ‘From Pride and Prejudice to Persuasion: Satisficing in Mate Search’, in G. Gigerenzer, P.M. Todd and the ABC Research Group, *Simple Heuristics that make Us Smart*, New York: Oxford University Press, pp. 287–308.

Tversky, A. (1972), ‘Elimination by Aspects: A Theory of Choice’, *Psychological Review* 79(4), pp. 281–299.

Tversky, A. & Kahneman, D. (1974), ‘Judgment Under Uncertainty: Heuristics and Biases’, *Science* 185, pp. 1124–1131.

Vriend, N.J. (1996), ‘Rational Behavior and Economic Theory’, *Journal of Economic Behavior & Organization* 29, pp. 263–285.

Whiten, A. and Byrne, R.W., eds., (1998), *Machiavellian Intelligence II: Extensions and Evaluations*. Cambridge, UK: Cambridge University Press.