Computational Modeling: Opportunities for the Information and Management Sciences[*]

Steven Orla Kimbrough
kimbrough@wharton.upenn.edu

August 12, 2002

Contents

---

[*] File: comexp-ics-hemant.tex.

Abstract

Computational explanations appeal to computational models, in contrast to equations or axioms, to explain their target systems. These models are typically inspired by natural phenomena and the term *natural computation* has been used in the literature. Darwinian or evolutionary models and explanations are a prominent form. This paper presents and reviews the concept of computational explanation, and its uses in the biological and social sciences. Emphasis is placed on recent innovations in algorithms for computational modeling. Further, the paper briefly describes application areas that are exploiting the modeling resources lately developed.

. . . what counts as an explanation has become more and more difficult to distinguish from what counts as a recipe for construction. [33, page 203]

1 Kinds of Explanation

Scientific research reliably and routinely produces substantial progress within its several disciplines. Physicists increase our knowledge of physics, biologists advance biology, social scientists advance social science, and management scientists advance management science. Much more unusually, scientific research produces—invents or discovers—new ways of doing science. My topic in this paper is one such development. I aim to describe and characterize it in an introductory fashion and then to comment on what it means—on the opportunities it presents—for the management sciences.

The innovation I have in mind has been noticed by others, and indeed as we shall see has developed at least since the 19th century. It has not, as far as I know, been recognized with a commonly-accepted name, even though the idea has been accepted in at least parts of several disciplines. So to begin, I shall give it a name: computational explanation.

The core thought is that a computational explanation is an explanation that appeals to a computation to do its explaining. Phenomena are explained as resulting from a computational process. If the phenomena in question are produced by a computer program, then a computational explanation is surely one we would think appropriate, even mandatory.[1] What is new, and perhaps surprising, is the applicability of computational explanations in the biological, social, and management sciences. That is what I want to discuss.

There are other kinds of explanation and they have predominated in scientific study. Without pretending to cover the topic with any justice, it will suffice for present purposes to focus on two contrasting types of explanation: covering law, and axiomatic. Covering law explanations are most familiar from physics. A general rule or law is presented, typically in the form of a differential or other sort of equation, and given initial conditions the state of a system can be predicted

---

[1] Recognizing, of course, the legitimacy of asking for explanations of non-computational features, e.g., the purpose of the program.

and explained subsequently. Think of Newtonian physics. Such explanations are something of a “gold standard” in science; the problem outside of physics has been the scope of their application. If more could be found, that would be welcome; in the meantime progress accumulates under other standards.

In what I’m calling an axiomatic explanation, a system of axioms replaces the covering law. Often, the axioms are motivated as much or more by normative considerations as by empirical findings. Think of utility theory. An ideal kind of preference is axiomatized. With initial conditions we may then deduce—and explain and predict—behavior of an ideally rational agent. Much of economics and game theory follows this basic pattern of reasoning. The challenge is to abstract a reasonable representation, formalize a description of it axiomatically, derive properties of the formal system, and test the conformance of the formal system with particular concrete instances.

All of this is well and good. Why should anyone be interested in any other form of explanation? The general complaint has been that, at least outside of physics, the techniques of covering law and axiomatic explanation have been insufficiently productive. There is, many have thought, much more to be achieved with scientific thought than can be achieved with these two methods. In particular, the successes of computational explanation in its various forms are, on the positive side, what has so encouraged its pursuit.

How does computational explanation differ from covering law and axiomatic explanations? To a first approximation, a computational explanation is:

- Procedural

  In contract to declarative (both covering law and axiomatic explanations tend to be declarative).

- Constructive.

  In some sense it answers to “If you can’t make one, you don’t know how it works.” Each step in the procedure is doable by the entity in question; a mechanism is available for performing each of the steps in the procedure.

- Representational.

  Semantic, or meaningful. Basically, the procedure has an interpretation; its states carry information or meaning about the world outside of the system. Purely random arrangements do not, but once we have selection, there is (some) such information and meaning present.

  Computational explanations add to the more basic procedural or constructive explanations that they are about something. Computations aren’t just procedures. They are (at least) procedures that are the way they are because of what they are.

These requirements are in order of increasing stringency. I’ll try to make these points clearer in the sequel. We’ll proceed by example.

2 Modeling the Firm: Cyert & March

Perhaps the earliest monograph—still read today—emphasizing a procedural view of management studies is Cyert and March's 1963 *A Behavioral Theory of the Firm* [13]. They open by announcing a perspective that has come to be known as, or associated with, the Carnegie School of thought on economics, organizations, and decision making.

> This book is about the business firm and the way it makes economic decisions. *We propose to make detailed observations of the procedures by which firms make decisions and to use these observations as a basis for a theory of decision making within business organizations.* Our articles of faith are simple. We believe that, in order to understand contemporary economic decision making, we need to supplement the study of market factors with an examination of the internal operation of the firm – to study the effects of organizational structure and conventional practice on the development of goals, the formation of expectations, and the execution of choices. [13, page 1, emphasis added]

The views they oppose and wish to offer alternative to are, or are associated with, neo-classical economics and its assumptions of ubiquitous knowledge and heroic rationality, and hence of the irrelevance to economics of decision processes by finite beings.

> The assumptions of rationality in the [neo-classical] theory of the firm can be reduced to two propositions: (1) firms seek to maximize profits; (2) firms operate with perfect knowledge. [13, page 8]

Cyert and March politely suggest that such views are unrealistic. Even more politely, they assert their interest in questions not addressed in—and they argue not addressable by—the neo-classical theory of the firm.

> Our conception of the task we face is that of constructing a theory that takes (1) the firm as its basic unit, (2) the prediction of firm behavior with respect to such decisions as price, output, and resource allocation as its objective, and (3) an explicit emphasis on the actual process of organizational decision making as its basic research commitment. [13, page 19]

Here are examples of questions Cyert and March propose to investigate with a behavioral theory, in contrast to the neo-classical theory.

> What happens to information as it is processed through the organization? What predictable screening Biases are there in an organization? What is the effect of conflict of interest on communication? What difference does time pressure make?
>
> ...

> In general, there are many questions about the behavior of business firms but only a few answers. Existing theory is not equipped to answer most of the questions we have raised. Where an answer can be derived by brute force, it tends to be ambiguous or conspicuously inadequate. [13, pages 21–2]

In consequence

> . . . we need to consider in somewhat more detail the actual procedures used by business firms to make economic decisions. . . . [B]usiness firms adapt over time by learning a number of simple decision rules and procedures and . . . a behavioral theory of the firm should deal both with that adaptive process and with the procedural implications of long-run adaptation. [13, page 117]

Without saying these views of the Carnegie School have triumphed since 1963, we can say they are now widely accepted if not prevalent. The debate continues and is surely relevant to the subject of this paper, but it is not central for present purposes. What is central is their views on methodology and the structure of their results. First, a word on their methodology.

A Behavioral Theory of the Firm is quite self-conscious about its methods and approaches. Nearly as much space is given over to explaining what the results mean and why they are legitimate as is given to presenting the results themselves. A summary will have to suffice. Happily, Cyert and March are obligingly articulate.

> [This study] involved four major research commitments. They are commitments that evolved during the course of the research, but they constitute a general retrospective characterization of our research strategy:
>
> 1. *Focus on a small number of key economic decisions make by the firm.* In the first instance, there were price and output decisions; subsequently they included internal allocation and market strategy decisions.
> 2. *Develop process-oriented models of the firm.* That is, we viewed decisions of the firm as the result of a well-defined sequence of behaviors in that firm; we wished to study the decisions by studying the process.
> 3. *Link models of the firm as closely as possible to empirical observations* of both the decision output and the process structure of actual business organizations. The models were to be both explicitly based on observations of firms and subject to empirical test against the actual behavior of identifiable firms.
> 4. *Develop a theory with generality beyond the specific firms studied.* We wanted a set of summary concepts and relations that could be used to understand the behavior of a variety of organizations in a variety of decision situations.

> [13, pages 1–2; emphasis in the original]

Point 2 is central to their approach: actual processes matter. Point 1, that there should be "Focus on a small number of key economic decisions make by the firm," is a modeler's point. Simplifications must, can and will be made. Points 3, that the models should be more or less directly testable and tested empirically, and 4, that the behavioral models will generalize, can be seen as a "reputation bet" against the neo-classical tradition. That tradition is noted, at least in some quarters, for models of great generality and attenuated empirical validity. In any event, parsimony of focus, empirical validation, and generalization have to be seen on their own as desirable goals for this kind of modeling. Cyert and March are assiduous in making their case in this regard throughout the book.

Now to the structure of their results. *A Behavioral Theory of the Firm* reports on a series of detailed field studies undertaken with American firms in the 1950s. Cyert and March describe unusually intimate access to their subjects. Interviews were done, first-hand participant observations were conducted, and corporate archives were thoroughly read. As a result, a number of decisions were closely observed.

The resulting models typically have the following form:

1. A flowchart describing a decision process at a high level
2. Detailed, behavioral models of the (key) processes within the flowchart

Figure 1, above, restates one of their flowchart models, in this case a model of how a department store responds to sales information.[2] For key steps in the process, e.g., revising and using the reorder rules, Cyert and March offer rather simple formulas or if-then rules (depending on the case), which they arrive at by observation. They then apply the resulting detailed model to an independent data set and test the predictions. Here is a typical summary of results obtained:

> In order to test the ability of the model to predict the price decisions that will be made by the buyer on new merchandise, an unrestricted random sample of 197 invoices was drawn. The cost data and classification of the item were given as inputs to the computer model. The output was in the form of a predicted price. Since the sample consisted of items that had already been priced, it was possible to make a comparison of the predicted price with the actual price.
>
> The definition of a correct prediction was made as stringent as possible. Unless the predicted priced matched the actual price to the exact penny, the prediction was classified as incorrect. The results of the test were encouraging; of the 197 predicted prices, 188 were correct and 9 were incorrect. Thus 95 percent of the predictions were correct. An investigation of the incorrect predictions showed that with minor modifications themodel could be made to handle the

---

[2] This example was chosen for its simplicity. Most of the models are considerably more complex.

1. Form sales estimates
2. Advance orders
3. Observe sales feedback
4. Is sales goal being achieved?
    - (a) If yes, go to 7
    - (b) Else, continue
5. Search
    - (a) Renegotiate constraints
    - (b) Mark-down
    - (c) New items
6. Revise reorder rules
7. Use rules to reorder
8. Re-evaluate sales goals and return to 1

Figure 1: General form of reaction to sales goal indicators. After Figure 6.1, [13, page 138]

> deviant cases. However, at this point it was felt that the predictive power was good enough so that further expenditure of resources in this direction was not justified. [13, pages 158–9]

In retrospect, *A Behavioral Theory of the Firm* offered not only a *procedural* theory, but a *computational* one as well. The behavioral models for the elements of the flowcharts appeal only to simple rules and formulas, e.g.,

> RULE 1 The estimate for the next six months is equal to the total of the corresponding six months of the previous year minus one-half of the sales achieved during the last month of the previous six-month period. [13, page 143]

The upshot is that Cyert and March were able to offer a theory of the firm, at least a sketch of one, that relied on sometimes complex arrangements (flowcharts) of typically rather simple computations (viz., RULE 1 above), and that came with impressive predictive power and generalization. An impressive achievement of computational modeling.

3 Molecular Genetics

Historians and philosophers of science have lately noted and much commented upon how biologists, especially in developmental biology and in molecular genetics, have increasingly resorted to computational models of biological systems.[3] In 1998, Yuh, Bolouri, and Davidson published an account of "an experimental analysis of the multiple functions of a well defined cis-regulatory element [the promoter region] that controls the expression of a gene [*Endo16*] during the development of the sea urchin embryo" [62]. Remarkably, "The outcome is a computational model of the element, in which the logical functions mediated through its DNA target site sequences are explicitly represented. The regulatory DNA sequences of the genome may specify thousands of such information-processing devices" [62]. The authors focus on module A of the promoter (controlling) region for Endo16. Here is where they find a molecular computer. A sense of how they think it works can be had by examining Figure 2 (published later).

---

[3] I am much indebted to Robin Fox Keller [33, cf. pages 239ff] for her discussion in this regard of several especially striking examples. The Yuh, Bolouri, and Davidson, and the Wray papers constitute one of her examples. The interested reader should consult her excellent monograph. My treatment of the examples here relies on the original sources and uses more recent informtion. For an early analytical treatment of genetic regulation (the lactose operon) see [38].

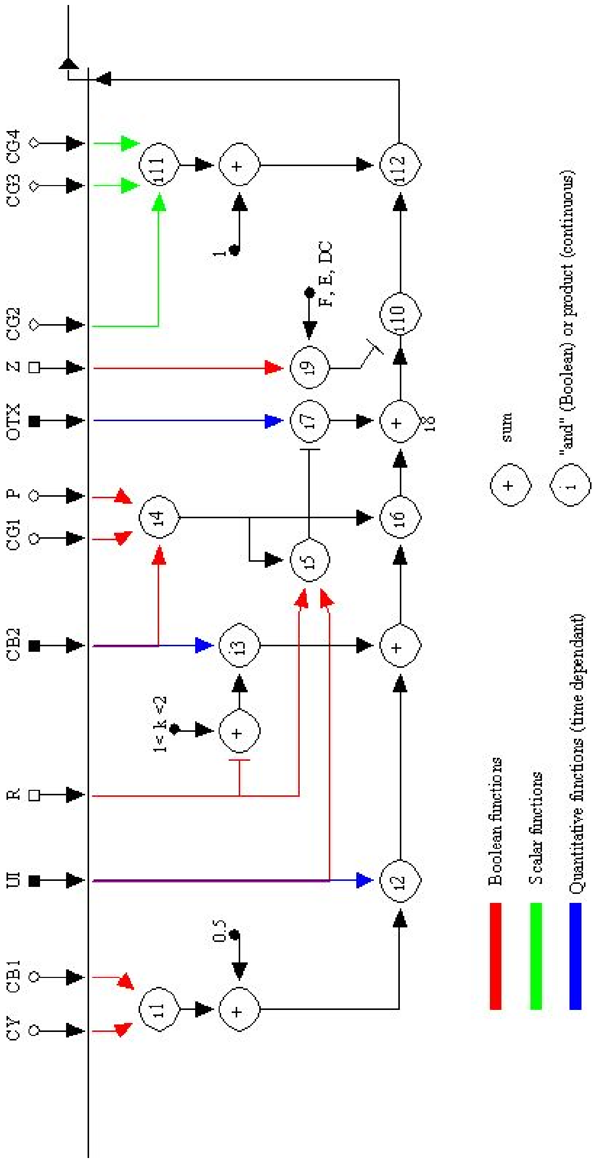

Figure 2: The Endo16 model of cis-regulatory control; from [63, 27 July 2002].

The paper finds that

> [T]he DNA sequence of module A specifies what is essentially a hard-wired, analog computational device. The requirement for this logic device is that there are many different inputs to the regulatory system that must be sorted appropriately. It is to us a remarkable thought that every developmentally active gene in the organism may be equipped with devices of this nature. ... The various functions mediated by module A are precisely encoded in the DNA sequence. Each target site sequence has a specific, dedicated function: Take the site away and the function is abolished; put it back in the form of a synthetic oligonucleotide and it reappears. [62]

Summing up, the authors write that

> The properties of the module A regulatory apparatus enable it to process complex informational inputs and to support the modular,polyfunctional organization of the *Endo16* cis-regulatory system. Perhaps the main insight from this experimental exploration is that these system properties are all explicitly specified in the genomic DNA sequence. [62]

That these are remarkable findings is confirmed by the presence of an editorial review of the piece appearing in the same issue in which it was published. The editorial, by Gregory A. Wray [58], serves to make the results of the research article more broadly accessible and to emphasize their significance. The following passage from Wray's editorial is methodologically significant for our purposes.

> To test their understanding of the *Endo16* promoter, Yuh and colleagues wrote a computer model that simulates these regulatory interactions. With the model, they made predictions about the consequences of specific promoter manipulations on transcription levels that were then tested experimentally. That these predictions were largely confirmed demonstrates not only an unusally complete understanding of how a particular promoter functions, but also the degree to which the *Endo16* promoter operates as an analog device. The “program that runs this tiny computer is directly in DNA as regulatory elements; its inputs are single molecules whose composition varies in time and among various cells of the embryo, and its output is a precise level of transcription.

The methodological commonality with Cyert and March is truly striking. Nor is it unusual in the biological literature. Many examples could be given. I'll close this section with an illustrative passage from a recent review article on morphogenesis. Notice how naturally the language of computation slips in, not as a metaphor, but as a description of how things are working.

> A major challenge now is to explore the possibility that there is also a conserved “morphogenetic code—a set of rules common to processes that are used repeatedly in different combinations to make

> functional organs. These instructions fall into two categories. First, there are basic subroutines that define essentially mechanical operations such as the packaging of cells into segments, the folding of epithelial sheets into tubes or cups, and the outgrowth of buds. Each of these modules utilizes sets of genes controlling properties such as differential cell adhesion, cell motility, cell-matrix interactions, and cytoskeletal organization.
>
> The second category determines how these subroutines are coordinated with cell proliferation and cell fate determination. This “project management depends on signaling centers that arise in the organ primordia or progenitor fields at positions initially determined by the primary embryonic axes. Each center is a group of cells that regulates the behavior of surrounding cells by producing positive and negative intercellular signaling molecules. [30]

4 Beginning with Darwin

Charles Darwin’s scientific and conceptual achievements continue to be celebrated and probed in search of deeper understanding. It is by now standard to note that evolutionary theories of adaptation (e.g., Lamarck’s and that of Charles’s grandfather Erasmus Darwin) predated *The Origin of Species* (1859). Darwin’s achievement was not to originate a theory of evolution. Rather it was (with Alfred Russell Wallace) to propose a workable (and largely correct) mechanism or *procedure* by which evolution comes about. In fact, Darwin did not use the term evolution in the first edition of the *Origin.* Instead, he consistently wrote of his theory of “descent with modification by natural selection.” Most fundamentally (and sufficient for present purposes) Darwin (and Wallace) put together three ideas or observations:

1. *Profusion* of individuals with *variance* of traits

   Every species has the capability to, and tends to, produce more offspring each generation than can possibly survive and reproduce. The individuals so produced are not all identical.

2. *Selection* among the variants

   Natural selection operates on populations of varying individuals, selecting for properties favorable to survival and reproduction.

   The individuals in a species vary in many different ways, including their capacity, at least in expectation, for survival and reproduction.

3. *Reproduction* of variants favored by selection, with *inheritance* of favorable traits.

   Inheritance may be approximate and far from perfect. What is required is a heritable association between the traits favorable to the parents and the traits of their offspring.

Given such a regime—of profusion with variance of traits, selection by traits, and reproduction with inheritance—it is nearly inevitable that evolution or "descent with modification by natural selection" will occur. Darwin and Wallace claimed that in fact it did routinely and that this process has in the main been responsible for adaptation and speciation. Such, boiled down for present purposes, is the (biological) theory of evolution by natural selection.[4] So successful has this theory become that now when we say something is an evolutionary theory or account, we *mean* it appeals to, or posits, a profusion-selection-reproduction process.[5] More accurately, we describe such accounts as *selectionist.* I shall employ the two terms as equivalents.

Once you have an evolutionary or selectionist perspective (theory), scientific progress largely consists in applying the theory in order to create particular explanations and predictions. In this way the theory both becomes useful and receives support. Darwin excelled in using his theory to explain and account for biological phenomena. Among other things, he provided evolutionary explanations for the structure of the phylogenetic tree, for sexual selection, for the fossil record, for the expression of emotions, and much else. And generally his accounts continue to be seen as correct in the main.

For present purposes, however, we are less interested in the biology and more interested in scientific methodology. From this perspective, Darwin's great achievement was to find a simple process (evolution by natural selection) that could be harnessed to model (for purposes, e.g., of explanation, prediction, and intervention) extremely complex and otherwise refractory phenomena (biology). Darwin offered a procedural model of adaptation and speciation. He did not offer a equational model or an axiomatic model (or indeed any other kind of model). Although the details of Darwin's procedure have been challenged and refined, serious models of adaptation and speciation since Darwin have not abandoned the core procedurality of the theory. (Gould's recent book [29] is an example of a serious effort to revise and refine evolutionary theory in biology.) Further, Darwin provided excellent examples of how to exploit a procedural theory. One observes the "pattern in evolution" (Eldredge's felicitous phrase, [21]), finds or posits special conditions and uses them in conjunction with the theory to account for the patterns.

Darwin (and Wallace) offered a procedural theory or model of biological adaptation and speciation. Lacking at the time a workable theory of how variation is produced and maintained, it might be questioned whether Darwin's was a constructive theory. (In his defense, variation, reproduction, etc. were clearly observable, even if their underlying mechanisms and causes were not well understood.) Nor could have Darwin conceived of the evolutionary process as computational, since that latter concept was unavailable in his time. The concept of computation is, however, available to us and we now turn to it.

---

[4] For a more detailed analysis see [35].

[5] With this qualification: reproduction may be weakened to mere survival, and profusion may be weakened to growth. These should be seen as limiting cases.

5 Church's Thesis and Computation, Broad and Narrow

The fundamental concepts and properties of computation were developed and discovered before the Second World War and before anything like a modern computing machine was built. It is easy to agree that determining the output of a function on the natural numbers should count as a computation. Given a well-defined function, f, taking one or more natural numbers as arguments, the process of applying the function to a particular configuration of argument values and determining the resulting value of the function is by all accounts a computation. It is sufficient to be a computation. And it is also assumed to be necessary, since no one has come up with a broader notion of computation. I shall call this the *broad concept of computation.*

Given this notion of computation, we can ask after the design of machines or collections of primitive functions that will perform computations. Alan Turing famously described an abstract machine, now called the Turing maching, which is able to perform some computations. That is, it can be set up so as to determine the outputs for some functions on the natural numbers. Turing's simple machine is surprisingly general and powerful (neglecting speed and other details of implementation). Yet it is known, originally by what is called the *Church-Turing theorem,* that there are some functions (on the natural numbers) the Turing machine cannot compute. Moreover, we have *Church's thesis* which says, essentially, that the Turing machine is maximally powerful in the sense that no other machine (or collection of primitive functions) can compute a function that the Turing machine cannot.

The Church-Turing theorem has been proved and is not in doubt. Church's thesis is not specific enough to be amenable to proof, but it can be disproved by a single counterexample. Because no one has come up with one since it was proposed in the 1930s, the thesis is generally judged to be true, or at least to be a warranted working assumption.

It is important to appreciate how truly broad the broad notion of computation is. Take any discrete procedure at all. We may think of it as transforming the (discrete) states of a given system from one set of values to another. We can always encode the values of the relevant states as natural numbers. Given this, the procedure can be viewed as a function on the natural numbers. Thus, in a certain sense, everything that happens (discretely) is a computation. Conversely, we can see that there is nothing special about the natural numbers other than that they are convenient entities for being the subject of functions. We could have used jelly beans instead, except that there aren't enough of them.

6 Evolutionary Mechanisms

Evolution in Darwin's sense of variation and selective retention has come to loom larger in importance than even Darwin might have imagined. There are broadly two ways in which evolution or selectionism has expanded its scope beyond (merely!) explaining adaptation and speciation in biology.

First, evolutionary mechanisms have been found to underlie other biological processes, notably the immune response, and learning and cognition. A longstanding puzzle in immunology has been to explain how an immune system can react differentially to about 4×10^9 antigens (foreign bodies provoking a reaction), when there are only on the order of 10^5 different genes in the human genome. The answer, it was discovered in the 1970s, is (roughly) that immune cells individually reshuffle their DNA, which allows production of new antibodies. Selection occurs by reinforcement of those cells producing successful antibody. Thus, the Darwinian process operates within the body of an organism, over a period of days.

Can Darwinian processes underlie behavior and cognition? As Richards describes in his masterful study, *Darwin and the Emergence of Evolutionary Theories of Mind and Behavior* [47], the positing and search for evolutionary mechanisms underlying behavior and cognition began with Darwin himself and were expanded upon dramatically by, among others, William James. Behaviorism in psychology succeeded in discouraging such investigations from the beginning of the twentieth century until the rise of cognitive psychology, roughly in the 1960s. Since then, evolutionary or Darwinian theories of mind have been espoused with increasing enthusiasm by scientists from a range of fields, including pyschology (e.g., Donald Campbell) and neurophysiology (e.g., William Calvin). Gerald Edelman, who won a Nobel prize for his contributions to immunology, has also been active.

The second way in which a Darwinian, selectionist perspective has proved valuable is in computational learning. As Holland (e.g., [31]) and others have shown, the simple genetic algorithm (GA) and its evolution programming variants are general-purpose computational algorithms. Evolution is a form of computation. Not only can we understand natural evolutionary processes as computations, but we can abstract and emulate them for practical purposes. A small industry—academic and commercial—has sprung up to do just that (see GECCO and other such conferences).

I assume that a manage science audience is familiar with GAs and evolution programs, and so I shall continue without discussing them in any detail. I do want to note, however, that not all *computational* theories of behavior and cognition are Darwinian. The following passage provides a convenient example.

> The idea that the representation of objects is essentially hierarchical and combinatorial was developed by the late David Marr in his influential 1982 book, *Vision.* Marr gave an account of how our brains process visual information, from the point at which light enters the eye to that at which we recognize the array of objects in the world around us; . . . This account was *computational* in the sense that the processes of vision were conceived as a series of computations of the sort that might be carried out on a digital computer. It was also based on what was known of the neurophysiology of vision. The

> way in which we perceive and represent the visual world is far from fully understood, however, and Marr's theory is in many respects incomplete. Nevertheless it is a useful starting point for a discussion of how objects are represented in the brain. [12, page 219]

The goal of completing a computational understanding of behavior and cognition remains one of the grand challenges of science, and one of the great opportunities of applications of computing.

7 Evolutionary Economics and Game Theory

We have touched on computational modeling in biology and in psychology. A computational, and especially evolutionary, perspective has also been taken and proved useful in the social sciences. I will focus on economics and game theory, where most of the work has been done and is being undertaken.[6]

Nelson and Winter's important work, *An Evolutionary Theory of Economic Change* [45] is an appropriate focus for our discussion, in addition to their recent review article [46]. Along with many others, Nelson and Winter have argued that the preoccupation in neoclassical economics with equilibria gives short shrift to explanation of many economic phenomena. General equilibrium theory and neoclassical economics, it is claimed, simply are not up to accounting for dynamics and bounded rationality, and do not actually give accurate descriptions of economic systems. Nelson and Winter describe their alternative, evolutionary approach as follows.

> [T]he modeling approach that we employ does not use the familiar maximization calculus to derive equations characterizing the behavior of firms. Rather, our firms are modeled as simply having, at any given time, certain capabilities and decision rules. Over time these capabilities and rules are modified as a result of both deliberate problem-solving efforts and random events. And over time, the economic analogue of natural selection operates as the market determines which firms are profitable and which are unprofitable, and tends to winnow out the latter. [45, page 4]

I refer the reader to their review article [46] for an overview of the extensive literature on evolutionary economics appearing since their book in 1982.

Epstein and Axtell's 1996 book *Growing Artificial Societies* [23] is an apt exemplar for a now active body of work that nicely complements evolutionary economics, although its origins and motivations are somewhat different. Epstein and Axtell (and the related work) take bounded rationality for granted. Methodologically, they are committed constructivists and avidly subscribe to the constructivist's motto: "If you can't build one, then you don't understand

---

[6] I note only in passing the popular concept of social evolution, based on what Dawkins called *memes*. This suggestive concept remains under development. I note, however, the related notion of "routines as genes" in Nelson and Winter [46].

it." Or, from the title of their book: "Social Science from the Bottom Up." (See also [22] for further reflections on constructivism.) This outlook coheres with, and lends support to, the general theme of emergence of surprising properties from (computationally) simple substrata that has captured the attention of so many on the leading edges of computational modeling and explanation.[7] Like Nelson and Winter (and the evolutionary economists generally), Epstein and Axtell (and that related literature) dispute many of the underlying assumptions of classical economics and present results (contrary or extending) that rely only on more realistic assumptions of what can be attributed to a finite agent. This is a rich, fruitful, and rapidly growing mode of research.

Finally, for our all-too-brief survey, evolutionary game theory has become a robust part of the study of games, along with classical game theory and behavioral game theory. (Samuelson provides a recent review from a rather classical or mainstream perspective see [50].) For our purposes, evolutionary game theory is concerned with

- The *behaviors* and especially *dynamics* generated by realistic strategies, and
- *Design principles* to support the fielding of artificial agents in strategic contexts

and assumes a sceptical attitude towards the adequacy of the Nash equilibrium as a solution concept for games played by actual agents.

Axelrod's work on iterated prisoner's dilemma (e.g., [2, 5]) is a paradigmatic example, although the relevant active literature is very large. It is known that there is exactly one Nash equilibrium in single-shot prisoner's dilemma and in prisoner's dilemma repeated a fixed (and known) number of times: mutual defection, to the detriment (hence 'dilemma') of both players. Also, if the game is to be repeated an indefinite number of times, there are a large number of Nash equilibria. In the definite and indefinite cases, human subjects tend to obtain rewards much higher than mutual defection. Thus, classical game theory either predicts poorly (the definite case) or not at all (the indefinite case). And classical game theory has little to offer on how to play the game, in either case. Axelrod held a series of contests in which entrant's computer programs played iterated prisoner's dilemma against each other. The data for these tournaments, quite surprisingly, revealed properties of generally successful strategies and of the remarkable (but situationally conditioned) power of TIT FOR TAT. This is but one example of surprising and useful results obtained from evolutionary game theory.

As noted, classical game theory has little to offer regarding how to design and field artificial agents in strategic contexts (see [19] for elaboration of the point). In response, most researchers with an interest in the question have looked to learning computations—usually either evolutionary computation or

---

[7] To mention just a few: Resnick and the StarLogo folks, SWARM, swarm intelligence, ant colony optimization, Wolfram and the cellular automata crowd, artificial life, natural computation [7].

some form of reinforcement learning, including learning classifier systems—to instruct their agents. If we don't know what to tell the agents, at least we can help them learn how to behave in strategic situations. Results to date are promising and manifold, but confined to the laboratory. Many of the studies have looked at standard game theory examples, such as prisoner's dilemma, stag hunt, or the ultimatum game (e.g., [64]). Results are beginning to appear for more realistic applications, such as supply chain management (e.g., [37]), and bidding at auction. The prospects for effective applications are genuinely encouraging.

8 Principles and Prospects

What does the foregoing intellectual history indicate for INFORMS researchers, researchers in Information Systems and/or Management Science? Where are relevant and promising applications likely to be found? Are there opportunities to contribute to science? In short, what work is there for us to do?

So far, I have introduced the concept of computational explanation, sought to clarify it enough for present purposes, and have proceeded by recounting examples. I now want to offer some organizing remarks in preparation for addressing the questions of the previous paragraph.

It is worth noting, in summary form, the academic disciplines that have been involved with computational modeling. What follows is merely a convenience sample, and an idiosyncratic one at that, meant only to indicate something of the breadth of the substantive contributions of computational modeling. (I apologize to all of those who should have been mentioned and aren't. Convenience is arbitrary.) Note that disciplinary distinctions here are especially problematic. To cite but one example, norms and conventions, are discussed by economists (e.g., [22, 60, 61, 59]), political scientists (e.g., [2, 3, 4, 5]), and philosophers (e.g., [8, 14, 52]). Also, connections among the disciplines are rampant, e.g., "The Biological Basis of Economic Behavior" [49].

1. Biology

   In addition to the papers cited above, Evelyn Fox Keller [33] has written an insightful and accessible monograph on computational models in biology, emphasizing embryology. Neurophysiologists, such as [20] and [10] have offered selectionist theories brain development and cognition. And most fundamentally, biology is the origin and home of selectionist theories of evolution.

2. Psychology and cognitive science

   Already mentioned is that selectionist accounts of cognition arose beginning in the nineteenth century and that detailed computational versions are in play today in biology (e.g., [10, 20]). Indeed, most of the field of cognitive science, and very much of artificial intelligence, can be understood as seeking computational models of cognition.

3. Economics (and commerce)

   The work of Nelson and Winter was discussed at length above [45, 46]. It has produced, or is associated with, a now large related literature. More recently, but also very productively, economists associated with the Brookings Institute and with the Santa Fe Institute have been active in exploring computational models of economic phenomena. See, e.g., [6, 23, 22], and [60, 61, 59]. Others have been active as well, e.g., [26].

4. Political science, organization science

   Robert Axelrod, political scientist at the University of Michican, has been a seminal contributor to computional modeling in the social sciences, cf., [2, 3, 4, 5]. Ian Lustick's work, e.g., [42], is an example of subsequent computational modeling in political science. In organizational science, I discussed Cyert and March's work at length, cf., [13]. A recent example in this tradition, drawing explicitly on developments in evolutionary explanation, is [41].

5. Philosophy

   A number of philosophers have been active in reflecting upon, and indeed producing, computational explanations. Much of the work is concerned with providing a naturalistic account of norms and conventions, e.g., Skyrms [52], Danielson [14, 15], and [8]. Paul Thagard, who holds a faculty position in both philosophy and computer science, has been active as a philosophically-oriented computational cognitive scientist, cf., Thagard [54, 55].

6. Computer and computational science

   Just as economists and game theorists have turned their attention to computational issues, so computer scientists and others with a computational science bent have begun to focus on game theory and economics, e.g., [51]. The recent DIMACS Workshop on Computational Issues in Game Theory and Mechanism Design,

   > http://dimacs.rutgers.edu/Workshops/gametheory/

   featured a representative selection of this work.

   Complementing the applications and explanations in the various fields, there has been a flowering of invention and investigation of algorithms for computational modeling. Genetic algorithms are perhaps the best known, but they are only a special case of a larger class that includes evolution programming, genetic programming, learning classifier systems, and evolutionary computation generally. See the GECCO conference for recent work, e.g.,

   > http://gal4.ge.uiuc.edu:8080/GECCO-2002/.

   Yet more generally, ant colony optimization, particle swarm optimization (PSO) [34], and various forms of "natural computation," algorithms inspired by naturally-encoun-tered (normally, biological) processes are very much in play [7].

Focusing now on the concerns of an INFORMS audience, what, we must ask, are we to make of all of this in the context of the applications-oriented world of management science and information systems? A small point is that this world has always recognized the distinction between, and the validity of both, theory and applications. Theory in this context has to include substantive scientific theory in biology, economics, political science, and so on. Our community should treat this as an opportunity to make contributions.

The larger point is that application opportunities are manifold. The basic applications methodology largely consists of exploring and modifying algorithms developed originally as computational explanations or abstractions of phenomena in the various sciences. Genetic algorithms are a clear example. I hope it is clear by now that they are just one of a rich family of algorithms originally developed as computational explanations or inspired by natural phenomena. There is much work to do to investigate (and even define) this larger class and to discover how it can best be used for practical purposes. Here, then, is a convenient classification of application areas pertaining to the management sciences. The comments are brief and meant to be only indicative of the attendant richness. Each taxon in the classification scheme—prediction, optimization, etc.—merits and receives the attention of entire careers. My aim here is to suggest and stimulate, not to pronounce and entomb.

1. Prediction

   It is surprising but true that despite many excellent techniques (e.g., statistical time series models), the prediction problem has not be definitively solved nor does progress on it appear to be exhausted. Koza [39] produced an early demonstration that evolutionary computation could be effective in finding good prediction models. Among others, Allen and Karjalainen [1] have applied evolutionary computation with success to stock market prices.

2. Optimization

   Genetic algorithms and evolutionary computing have long been applied to practical optimization problems. Goldberg's Ph.D. thesis under Holland [27] recorded an early and conspicuous success. Good reviews can be found in [11, 28, 43, 44]. Innovation and progress continues, cf., [36] and much else at the various GECCO conferences.

   Remarkably, during the last ten years several intriguing computational techniques, modeled on natural phenomena, have appeared, or been revived, and have shown much promise. These include ant colony optimization (e.g., [16, 18, 17]), particle swarm optimization (e.g., [34]), memetic algorithms (see [34, pages 245–255] for a not entirely favorable review; and [48] for a commerce-oriented application), reinforcement learning (often applied to game-playing agents; [53]), and various kinds of "natural computation" (e.g., [7]). These *families* of algorithms are only beginning to be explored.

3. Encoding

   Rule-based expert systems were an early and moderately useful way to encode knowledge for subsequent use. Their range of application has been limited by the cost and trouble of "knowledge acquisition," the process of obtaining information in a form that fits the encoding scheme. Holland's "classifier systems" [9, 28, 31], a software architecture that facilitates automated knowledge acquisition, were conceived, in part, as a response to this problem. There is now an active community working on generalizations of classifier systems, the class usually referred to as Learning Classifier Systems (LCS), [56, 57, 40]. See the annual International Workshop on Learning Classifier Systems.

4. Design

   Computational modeling is likely, I believe, to prove useful in design of supply chains, logistics systems, auctions, and generally in systems in which agents must confront strategic behavior. See [37] for a preliminary study on artificial agents and supply chains. The governing thought is that, in designing automated systems with strategic interactions, adaptive, learning artificial agents can be used to discover effective strategies and to test proposed system designs. See also [48] for an application in strategy formation. And there is much literature on design of game-playing agents, e.g., [24, 51, 25, 64]

5. Extraction

   Just as expert systems have long been used to encode information, neural nets have long been used to extract information from, and to obtain classifications of complex phenomena. Such applications continue, along with underlying innovations that make the use of neural nets more effective (e.g., by using evolutionary computation to adjust connection weights). In addition, the broader classes of algorithms, already discussed, show much promise in this task. John Holmes, for example, reports good success with using Learning Classifier Systems to perform data mining [32]. Indeed, by whatever algorithmic method, data mining and text data mining are very promising application areas for the methods here discussed.

9 Summary and Conclusion

The computational modeling enterprise seeks to model naturally-occur-ring processes as computations, either broadly or narrowly defined. Model and, of course, support explanations, predictions, interventions, and all the purposes of science generally. It is not (primarily) the investigation of algorithms as abstract entities (as in computer science); rather it is the attempt to discover and understand computational, or algorithmic, processes in nature and practice. Computational modeling may be contrasted with equational and axiomatic modeling, and as such it presents a genuinely distinct mode of doing science and engineering, i.e., of seeking explanations and predictions, and of supporting interventions and innovations.

The successes of computational modeling reach back at least to Charles Darwin, who (along with Alfred Russell Wallace) originated a theory of biological adaptation and speciation as due to "descent with modification by natural selection," of evolution as fundamentally proceeding by a process of variation and selective retention. Evolutionary theories—involving essentially processes of variation and selective retention—are characteristically procedural and (usually) constructive. They come with explanatory mechanisms that instantiate the variants on which selection acts. They are a special category of computational explanation: evolutionary computation. They are not, however, the whole story. Remarkably examples of a more general class have been discovered or invented of late and are showing much promise both for support basic science and for applications. Indeed, the two purposes are complementary.

It is an exciting time to be a computational modeler.

Acknowledgement

This material is based upon work supported by, or in part by, NSF grant number SES-9709548.

References

[1] Franklin Allen and Risto Karjalainen. Using genetic algorithms to find technical trading rules. *Journal of Financial Economics*, 51:245–271, 1999.

[2] Robert Axelrod. *The Evolution of Cooperation.* Basic Books, Inc., New York, NY, 1984.

[3] Robert Axelrod. The evolution of strategies in interated prisoner's dilemma. In Lawrence Davis, editor, *Genetic Algorithms and Simulated Annealing*, pages 32–41. Morgan Kaufman, Los Altos, CA, 1987.

[4] Robert Axelrod. An evolutionary approach to norms. *American Political Science Review*, 80(4):1095–1111, 1996.

[5] Robert Axelrod. *The Complexity of Cooperation.* Princeton University Press, Princeton, NJ, 1997.

[6] Robert Axtell. The emergence of firms in a population of agents. Working paper 99-03-019, Santa Fe Institute, Santa Fe, NM, 1999.

[7] Dana H. Ballard. *An Introduction to Natural Computation.* MIT Press, Cambridge, MA, 1997.

[8] Christina Bicchieri, Richard Jeffrey, and Brian Skyrms, editors. *The Dynamics of Norms.* Cambridge University Press, New York, NY, 1997.

[9] L. B. Booker, D. E. Goldberg, and J. H. Holland. Classifier systems and genetic algorithms. In J. G. Carbonell, editor, *Machine Learning: Paradigms and Methods*, pages 235–282. The MIT Press, Cambridge, MA, 1990. Reprinted from *Artificial Intelligence: An International Journal,* vlume 40, number 1–3 (1989).

[10] William H. Calvin. *How Brains Think: Evolving Intelligence, Then and Now.* BasicBooks, New York, NY, 1996.

[11] Carlos Artemio Coello Coello. A survey of constraint handling techniques used with evolutionary algorithms. technical report Lania-RI-99-04, Laboratorio Nacional de Informática Avanzada, Veracruz, México, 1999. http://www.lania.mx/˜ccoello/constraint.html.

[12] Michael C. Corballis. *The Lopsided Ape.* Oxford University Press, New York, NY, 1991.

[13] Richard M. Cyert and James G. March. *A Behavioral Theory of the Firm.* Blackwell, Cambridge, Massachusetts, second edition, 1992. Originally published in 1963 by Printice-Hall, Inc.

[14] Peter Danielson. *Artificial Morality: Virtuous Robots for Virtual Games.* Routledge, New York, NY, 1992.

[15] Peter Danielson, editor. *Modeling Rationality, Morality, and Evolution*, volume 7 of *Vancouver Studies in Cognitive Science.* Oxford University Press, New York, NY, 1998.

[16] Marco Dorigo and Vittorio Maniezzo andAlberto Colorni. The ant system: Optimization by a cology of cooperating agents. *IEEE Transactions on Systems, Man, and Cybernetics–Part B*, 26(1):1–13, 1996.

[17] Marco Dorigo and G. Di Caro. The ant colony meta-heuristic. In D. Corne, M. Ddorigo, and F. Glover, editors, *New Ideas in Optimization*, pages 11–32. McGraw Hill, London, UK, 1999.

[18] Marco Dorigo, G. Di Caro, and L. M. Gambardella. Ant algorithms for discrete optimization. *Artificial Life*, 5(2):137–172, 1999.

[19] Garett O. Dworman, Steven O. Kimbrough, and James D. Laing. Bargaining by artificial agents in two coalition games: A study in genetic programming for electronic commerce. In John R. Koza, David E. Goldberg, David B. Fogel, and Rick L. Riolo, editors, *Genetic Programming 1996: Proceedings of the First Annual Genetic Programming Conference, July 28-31, 1996, Stanford University*, pages 54–62. The MIT Press, 1996. File: GP96_k08.doc.

[20] Gerald M. Edelman. *Neural Darwinism: The Theory of Neuronal Group Selection.* BasicBooks, New York, NY, 1987.

[21] Niles Eldredge. *The Pattern of Evolution.* W H Freeman & Co, New York, 2000.

[22] Joshua M. Epstein. Learning to be thoughtless: Social norms and individual computation. *Computational Economics*, 2000. Also, The Brookings Institution, Center on Socail and Economic Dynamics Working Paper No. 6, January 2000.

[23] Joshua M. Epstein and Robert Axtell. *Growing Artificial Societies: Social Science from the Bottom Up.* The MIT Press, Cambridge, MA, 1996.

[24] Ido Erev and Alvin E. Roth. Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria. *The American Economic Review*, 88(4):848–881, 1998.

[25] Christina Fang, Steven O. Kimbrough, Stefano Pace, Annapurna Valluri, and Zhiqiang Zheng. On adaptive emergence of trust behavior in the game of stag hunt. *Group Decision and Negotiation*, forthcoming.

[26] Dhananjay K. Gode and Shyam Sunder. Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *Journal of Political Economy*, 101(1):119–137, 1993.

[27] David E. Goldberg. *Computer-Aided Gas Pipeline Operation Using Genetic Algorithms and Rule Learning.* PhD thesis, University of Michigan, Ann Arbor, MI, 1983.

[28] David E. Goldberg. *Genetic Algorithms in Search, Optimization & Machine Learning.* Addison-Wesley Publishing Company, Inc., Reading, MA, 1989.

[29] Stephen Jay Gould. *The Structure of Evolutionary Theory.* Harvard University Press, Cambridge, MA, 2002.

[30] Brigid L. M. Hogan. Morphogenesis. *Cell*, 96:225233, January 22, 1999.

[31] John H. Holland. Escaping brittleness: The possibilities of general-purpose learning algorithms applied to parallel rule-based systems. In R. S. Michalski, J. G. Carbonell, and T. M. Mitchell, editors, *Machine Learning II.* Morgan Kaufmann, 1986.

[32] John H. Holmes. Applying a learning classifier system to mining explanatory and predictive models from a large clinical database. In Pier Luca Lanzi, Wolfgang Stolzmann, and Stewart W. Wilson, editors, *Advances in Learning Classifier Systems*, volume 1996 of *Lecture Notes in Artificial Intelligence*, pages 103–113. Springer, Berlin, Germany, 2001.

[33] Evelyn Fox Keller. *Making Sense of Life: Explaining Biological Development with Models, Metaphors, and Machines.* Harvard University Press, Cambridge, MA, 2002.

[34] James Kennedy and Russell C. Eberhart. *Swarm Intelligence.* Morgan Kaufmann Publishers, San Francisco, CA, 2001.

[35] Steven O. Kimbrough. The concepts of fitness and selection in evolutionary biology. *Journal of Social and Biological Structures*, 3:149–170, 1980.

[36] Steven O. Kimbrough, Ming Lu, David Harlan Wood, and D. J. Wu. Exploring a two-market genetic algorithm. In *Proceedings of GECCO 2002 (Genetic and Evolutionary Computation COnference).* AAAI Press, 2002. Available at http://grace.wharton.upenn.edu/˜sok/sokpapers/2001-2/twomarketgas/ga202.pdf.

[37] Steven O. Kimbrough, D.J. Wu, and Fang Zhong. Computers play the beer game: Can artificial agents manage supply chains? *Decision Support Systems*, 33:323–333, 2002. File: DSS_KWZ_J25.doc.

[38] Steven Orla Kimbrough. On the reduction of genetics to molecular biology. *Philosophy of Science*, 46(3):389–406, 1979.

[39] John R. Koza. *Genetic Programming: On the Programming of Computers by Natural Selection.* MIT Press, Cambridge, MA, 1992.

[40] Pier Luca Lanzi, Wolfgang Stolzmann, and Stewart W. Wilson, editors. *Advances in Learning Classifier Systems*, volume 1996 of *Lecture Notes in artificial Intelligence: Subseries of Lecture Notes in Computer Science.* Springer, Berlin, Germany, 2001. Revised papers from the Third International Workshop, IWLCS 2000.

[41] Daniel A. Levinthal. Adaptation on rugged landscapes. *Management Science*, 43(7):934–950, 1997.

[42] Ian S. Lustick. Agent-based modeling of collective identity: Testing constructivist theory. *Journal of Artificial Societies and Social Simulations*, 3(1), January 2000. http://jasss.soc.surrey.ac.uk/3/1/1.html.

[43] Zbigniew Michalewicz. A survey of constraint handling techniques in evolutionary computation methods. In *Proceedings of the 4th Annual Conference on Evolutionary Programming*, pages 135–155, Cambridge, MA, 1995. MIT Press. http://www.coe.uncc.edu/˜zbyszek/papers.html.

[44] Zbigniew Michalewicz. *Genetic Algorithms + Data Structures = Evolution Programs.* Springer, Berlin, Germany, third edition, 1996.

[45] Richard R. Nelson and Sidney G. Winter. *An Evolutinary Theory of Economic Change.* The Belknap Press of Harvard University Press, Cambridge, MA, 1982.

[46] Richard R. Nelson and Sidney G. Winter. Evolutionary theorizing in economics. *Journal of Economic Perspectives*, 16(2):23–46, Spring 2002.

[47] Robert J. Richards. *Darwin and the Emergence of Evolutionary Theories of Mind and Behavior.* University of Chicago Press, Chicago, IL, 1989.

[48] Jan W. Rivkin. Imitation of complex strategies. *Management Science*, 46(6):824–844, June 2000.

[49] Arthur J. Robson. The biological basis of economic behavior. *Journal of Economic Literature*, XXXIX:11–33, March 2001.

[50] Larry Samuelson. Evolution and game theory. *Journal of Economic Perspectives*, 16(2):47–66, Spring 2002.

[51] T. Sandholm and R. Crites. Multiagent reinforcement learning in iterated prisoner's dilemma. *Biosystems*, 37:147–166, 1995. Special Issue on the Prisoner's Dilemma.

[52] Brian Skyrms. *Evolution of the Social Contract.* Cambridge University Press, New York NY, 1996.

[53] Richar S. Sutton and Andrew G. Barto. *Reinforcement Learning: An Introduction.* The MIT Press, Cambridge, MA, 1998.

[54] Paul R. Thagard. *Computational Philosophy of Science.* MIT Press, Cambridge, MA, 1993.

[55] Paul R. Thagard. *Coherence in Thought and Action.* MIT Press, Cambridge, MA, 2000.

[56] Stewart W. Wilson. Classifier fitness based on accuracy. *Evolutionary Computation*, 3(2), 1995.

[57] Stewart W. Wilson. State of xcs classifier system research. Technical Report 99.1.1, Prediction Dynamics, Concord, MA, 18 March 1999.

[58] Gregory A. Wray. Promoter logic. *Science Magazine*, 279(5358):1871–1872, 1998.

[59] H. Peyton Young. *Individual Strategy and Social Structure: An Evolutionary Theory of Institutions.* Princeton University Press, Princeton, NJ, 1998.

[60] H. Peyton Young. The evolution of conventions. *Ecoometrica*, 61:57–84, 1993.

[61] H. Peyton Young. The economics of convention. *Journal of Economic Perspectives*, 10:105–22, 1995.

[62] Chiou-Hwa Yuh, Hamid Bolouri, and Eric H. Davidson. Genomic cis-regulatory logic: Experimental and computational analysis of a sea urchin gene. *Science Magazine*, 279(5358):1896–1902, 1998.

[63] Chiou-Hwa Yuh, Hamid Bolouri, and Eric H. Davidson. The Endo16 model of *cis*-regulatory control. Pages on the World Wide Web, Accessed 2002-07-27. URL: http://strc.herts.ac.uk/bio/maria/NetBuilder/-Examples/Endo16/Endo16sim01.htm.

[64] F. Zhong, Steven O. Kimbrough, and D.J. Wu. Cooperative agent systems: Artificial agents play the ultimatum game. *Group Decision and Negotiation*, forthcoming.

`File: compexp-ics-hemant.tex`
`$Id: compexp-ics-hemant.tex,v 1.4 2002/08/13 03:37:18 sok Exp $`