

examples in the book are purely mathematical examples, with no interpretation. Bibliography, which consists of 105 entries, is useful, but by no means satisfactory: too many important references regarding evidence theory are not included in it.

On the whole, the book is a welcome addition to the rapidly growing literature on evidence theory, particularly for a mathematically oriented reader. As revealed by the authors in the Preface, the prospective second volume will deal with support functions.

REFERENCES

- Shafer, G. [1976], *A Mathematical Theory of Evidence*. Princeton Univ. Press, Princeton, N.J.
 Barnett, J. A. [1981], "Computational methods for a mathematical theory of evidence." *Proc. Seventh Int. Joint Conf. on AI*, pp. 868–875.
 Yen, J. [1986], "A reasoning model based on an extended Dempster-Shafer theory." *Proc. AAAI'86*, pp. 125–131.
 Yen, J. [1989], "GERTIS: A Dempster-Shafer approach to diagnosing hierarchical hypotheses." *ACM Communications*, 32, pp. 573–585.

GEORGE J. KLIR
Dept. of Systems Science
State University of New York
Binghamton, NY 13902, U.S.A.

LIFE ITSELF: A COMPREHENSIVE INQUIRY INTO THE NATURE, ORIGIN, AND FABRICATION OF LIFE, by Robert Rosen. Columbia University Press, 1991, xix + 284 pages.

The history of biology is marked by equivocation or outright failure to identify and define its fundamental categories: what is alive? what is an organism? While a small cadre of systems theorists still struggle with this core question of both systems science and biology, biologists tend to simply ignore this fundamental lack, occupied if not content with the empiricism of their work.

Therefore, before discussing some of the details of the book, I will satisfy the eager and impatient reader by giving away the elegant punch-line from this new, fundamental work in theoretical biology by a pre-eminent systems theorist, Robert Rosen: "A material system is an organism if and only if it is closed to efficient causation [p. 244]." The meaning and consequences of this statement are what this book is about.

CATEGORY THEORY

Life Itself deals freely with the mathematics of *elementary* category theory, abstract algebra, calculus, and systems theory. While the mathematics is at a high conceptual level, the mathematics is not deep, and can be followed easily by a reader with some mathematical patience and breadth of exposure.

Rosen looks to a mathematics based on category theory rather than the traditional Whitehead-Russell language of sets and functions to be the canonical language for systems theory and the "sciences of complexity". Invented to ease the fusion of

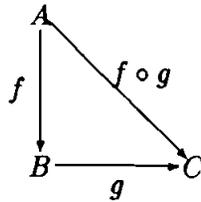


Figure 1 A Composite Relation.

abstract algebra and topology, as a language, category theory has some great advantages. It is hoped that these advantages will allow category theoretic expressions to more closely capture the *semantic* aspects of theories expressed in natural language without sacrificing any objectivity or rigor.

First, category theory is inherently *structural* and *graphical*, based as it is on the block-and-arrow diagram so common in the information sciences. In category theory a functional relation $f: A \mapsto B$ is denoted by a labeled, directed arrow: $A \xrightarrow{f} B$. Such a relation is called a *component*, and can be seen as processing inputs $a \in A$ into outputs $b \in B$. For a more complex example, the set-theoretic functional relations

$$f: A \mapsto B, \quad g: B \mapsto C, \quad f \circ g: A \mapsto C$$

are denoted in category theory as the group of related components in Figure 1. Notations built in this language are more than convenient. They allow expression of complex relations in a *topological* form which is both conducive to human understanding and more expressive than simple functional notation.

Also, category theory is inherently *reflexive*: the interrelations between sets, mappings, functionals, and other elements of mathematics, usually regarded as distinct types of entities, are easily expressed. Such concepts as functions determined by other functions and the self-application of relations are naturally expressed. In particular, we can propose a relation $B \xrightarrow{\Phi} H(A, B)$ where $H(A, B) = \{f: A \mapsto B\}$, so that we can say $\Phi(f(a)) = \Phi(b) \mapsto f$. The component notation $A \xrightarrow{f} B \xrightarrow{\theta} H(A, B)$ can also be interpreted as $A \xrightarrow{f} B \xrightarrow{\theta} f$, and has the category theoretic shorthand expression of Figure 2. The effect is that the function *f itself* is selected by Φ on the



Figure 2 Self-entailment of Relations.

basis of its own *products* b . These kinds of circular relations will prove crucial to the definition of the organism.

The search for a canonical language for cybernetics and systems is not unique to Rosen (see also Brown 1972, Heylighen 1990, Joslyn 1991, Klir 1985). His “relational” language of category theory is similar to some of these others, and seems ideal for cybernetics, concerned as it is with modeling and abstract structure in the most general sense. Category theory captures the abstract *structural* relations among components, and promises to serve as a generic modeling language for both simple and complex systems.

Relational diagrams reflect pure “organization”, measured by the density of entailments within them. This concept of organization is completely removed from those concerning disequilibrium, improbability, or entropy common in classical information systems theory. The elements of a relational diagram are devoid of any explicit referents (for example, an explicit representation of time). None of the baggage of dynamical systems need be brought to bear.

ENTAILMENT AND CAUSALITY

Thus, relations of the form $A \xrightarrow{f} B$ are the fundamental units of expression for Rosen. Each one is called a *component*, and expresses a form of general *entailment*—a relation of *necessity*—from A to B . As such they are uninstantiated, and can represent any kind of relation (for example, ontological or epistemic). When we are considering entailments in the world, then we regard of them as *causal* relations. Thus, through analyzing entailment structures we can profitably use the classical causal categories of Aristotle.

Consider a set of entailments $A \xrightarrow{f_i} A$, $1 \leq i \leq n - 1$, and let $A \xrightarrow{f_n} B$. We can consider $F = \{f_i\}$ as a set of *inference rules*, which can be denoted as $\vec{F} = \langle f_i \rangle$ (where $\langle \cdot \rangle$ is an ordered vector) when they are applied in a *sequence*. Then the domain $\text{dom}(f_1)$ of f_1 can be seen as a set of *axioms* which may act on axioms and *theorems* in $\text{dom}(f_i)$, $i > 1$ to produce other theorems. Finally consider a theorem $b \in B$ derived from an axiom $a \in A$ according to the rule sequence \vec{F} :

$$f_n(f_{n-1}(\dots f_1(a) \dots)) \mapsto b. \quad (1)$$

We can see b as an *effect* which naturally generates the question “why b ?” There are multiple answers depending on the Aristotelian modality of the question, and each answer maps to a classical category for both Aristotle, logical systems, and dynamical systems:

Because	Category	Logic	Dynamics
a	Material	Axioms	Initial conditions
F	Efficient	Inference rules	Dynamical equations
\vec{F}	Formal	Algorithm	Trajectory

We note that these causal categories are *independent* of each other, and are themselves not entailed: the same *b* could be reached with different axioms, different inference rules, and/or a different order of application of those rules.

The Aristotelian category of “final cause”, of course, requires more consideration, bringing up as it does issues of teleology, purpose, function, vitalism, and even *meaning*. But in equation 1 we note that the symbol *b* itself has yet to be mapped to an Aristotelean category. This is the way that final cause is introduced, by understanding *b* as both an effect and reflexively as *itself* a cause: a *final* cause of that which entails *it*. Thus final causes are *contingent*, dependent on the larger system of which they are a part.

We can examine the relation between the causal categories for a single entailment $A \xrightarrow{f} B$ such that $f(a) \mapsto b$ with $a \in A$ and $b \in B$:

Why?	Cause	Category
<i>b</i>	Because <i>a</i>	Material
<i>b</i>	Because <i>f</i>	Efficient
<i>f</i>	Because <i>b</i>	Final
<i>a</i>	Because <i>b</i>	Final

The symbols *f* and *a* themselves are not entailed *by* the system; they are *given* from outside the system. Therefore they are interpreted as *finally caused*, their “purpose” being to cause *b*.

Thus, when we attempt to involve final cause in our explanatory scheme, we come to recognize a number of serious weaknesses in classical formalisms. First, since a final cause appears to violate expected causal temporality, being subsequent to its effect, the classical linear flow of time from axioms to theorems is not observed. Second, in classical entailment schemes, entailments *themselves* are not subject to further entailment. They are always “given” from “above”, explainable only in terms of their *final* cause, or purpose, never in terms of their *efficient* cause, or explanation as to how they came about: “In short, the efficient cause of something *inside* the system is tied to final cause of something *outside* the system. [p. 246]”

MODELS AND SIMULATIONS

The movement toward organisms depends on a detailed argument that there is a fundamental, unbridgeable gap between general *models* and special case *simulations*. This distinction is used, in turn, to generate that between *organisms* and *machines*.

The basics of modeling theory are well established and easily stated (Cariani 1989, Klir 1991, Pattee 1988, Rosen 1985, Turchin 1977). Given two systems of “internal entailment” $W \xrightarrow{g} W$ and $K \xrightarrow{h} K$, and given an *encoding* relation $W \xrightarrow{E} K$ and complementary *decoding* relation $K \xrightarrow{D} W$, then if $\forall w \in W, D(h(E(w))) \mapsto g(w)$, then the system $\langle\langle W, g \rangle, \langle K, h \rangle, E, D\rangle$ is a *model*.

If *W* is a system of causal, ontological entailments (an aspect of the natural world), and *K* is a system of inferential, epistemic entailments (a knowledge system), then we can describe “natural law” as the establishment of a modeling relation between

the causal system and the inferential system which brings them into congruence through the encodings and decodings of physical measurements.

What's crucial in this definition is that while g and h are entailments *within* W and K , the relations E and D *between* W and K themselves are not entailed. Thus, while g and h can be regarded as "syntactic" aspects of the formalism, E and D are necessarily *semantic* in nature: there is no *necessary* relation between the entailments within W and K and the codings. Thus, the codings, and in particular the measurement procedures of natural science, represent the category of final cause, reflecting the *purpose* of that which posited them.

Models are distinguished between *analytic* and *synthetic* based on the mathematical properties of their state spaces. Given a set of measurements $W \xrightarrow{f_i} K$ and a representation of the domain W as the Cartesian product of the equivalence classes of the f_i , then the result is an analytical model. There might be multiple such analytical models, and the set of all such analytical models forms a category under the partial ordering of inclusion among the equivalence classes.

Since the measured values in an analytic model are elements of Cartesian products, they all "overlap" each other. Due to these "holistic" effects, they cannot be "reduced". On the other hand, in a "synthetic" model the equivalence classes of the measurements f_i establish a *partition* on the domain W into disjoint W_i . Thus, synthetic models *can* be successfully "decomposed" into reduced units. The partitioning of synthetic models generates a *finite* semi-lattice with a unique upper bound (a finest model) and multiple lower bounds (coarsest models) such that the upper bound is the global intersection of all the lower bounds. Thus, they reflect the essence of *reducibility*: all measurable properties can be expressed as a function of the observables in the maximal model, and thus reduced to them.

Rosen asserts that synthetic models form the core, indeed the *entirety*, of modern physical science. The largest synthetic model is the "atomic" decomposition of the system, the finest micro-model available, the canonical "Hamiltonian" from which all other properties can be derived by simply appending state spaces. However, synthetic models are in fact a very special case: while every synthetic model is analytic, the contrary is generally not true.

Simulations are kinds of synthetic models and are strongly contrasted with general models. Again, assume two systems of entailment $W \xrightarrow{g} W$ and $K \xrightarrow{h} K$. We recognize a distinction in K between its *hardware*, corresponding to the efficient causes in K (f in the relation $f(a) \mapsto b$), and its *software*, corresponding to the material causes (a). When there exists an encoding $W \xrightarrow{E} K$ such that each state in W results in the specification of some initial condition, or *program* for the software of K , then K simulates W .

The differences between models and simulations are reflected in the fact that, in a simulation, the efficient causes (hardware) of the simulated system are converted into material causes (software) for the simulator, while in a model, any hardware/software distinctions in the object are preserved. Because of this "corruption", this loss of information, nothing about the *actual* inferential structure of W can be learned from its simulation, unlike a proper model.

Instead of a *congruence* between two different inferential structures, themselves unentailed, simulation gives us an *entailment*, by the simulator, of the entailment it simulates. In other words, simulation turns [its object] into an effect, a consequence of the simulator, which may have no relation to it at all. [p. 200]

MECHANISMS AND ORGANISMS

Rosen defines a *mechanism* as any system all of whose models are in turn simulable,² and demonstrates that mechanisms must necessarily have only synthetic models. Mechanisms have a necessarily *linear* causal structure. This arises when, in an attempt to provide an explanation, an entailment, for the initial conditions *a* and inference rules *f* of a given system, a meta-state space is hypothesized. Then further meta-meta-state spaces are hypothesized to explain the initial conditions and inference rules of the meta-system, etc. As with Ptolemaic epicycles, an infinite regress is immediately recognized. We could consider the limit of this sequence, but the result would not itself be a mechanism, as it would not have a maximal model.

Thus in mechanisms there must always be a level which is itself unentailed, which never has an *efficient* cause, but can only *be a final* cause of all that follows from it. Thus Rosen's primary conclusion: mechanisms, the very stuff of existing science, necessarily have very "impoverished" entailment relations. The class of *simple* systems (those mechanistic, fractionable, reducible, simulable systems with decidedly *linear* entailment structures) is necessarily smaller than the class of *complex* systems with general analytic models. But whereas, for example, control theorists recognize their linear systems as very simple and special cases of the general non-linear cases, three hundred years of science has been dedicated to the idea that the special class of simulable systems is in fact a *universal* paradigm for explanation of natural phenomena.

As machines, being synthetic, are the special case, so we arrive at the concept of the *organism* as the proper general case. Organisms have properly analytic models, and their entailment structures can be very rich, containing many loops. An organism cannot therefore be constructed as a machine, but perhaps in the limit of a series of machines (cf. epicycles). But further, organisms contain many *parts* which *are* machines; indeed they admit to multiple, complementary, individually incomplete mechanistic models. This explains the tantalizing mirage of the machine metaphor in biology: in many ways organisms are *similar* to machines, and have many *parts* which act like machines, yet they stubbornly refuse to yield to mechanistic models.

With this explanation in mind, let us return to the punch-line, Rosen's definition of life.

A material system is an organism if and only if it is closed to efficient causation. That is, if *f* is any component of such a system, the question "why *f*?" has an answer within the system, which corresponds to the category of efficient cause of *f*. [p. 244]

Closure to efficient causation means that no "production rule" in an organism is "given": all must be generated from *within* the organism, thus stopping the infinite regress. Rosen compares this move to that of Newton's second law, which by making acceleration a function of phase, terminates an infinite Taylor's expansion for position to two places.

²He also defines a *machine* as a mechanism any of whose models is a "mathematical machine". This is a rather unfortunate use of terms, obviously circular. How we know that a model is a mathematical machine is not clear.

Rosen provides a specific example of a “candidate organism” which is closed to efficient causation and expresses the three biological functions of metabolism, repair, and replication. Assume three entailments:

$$A \xrightarrow{f} B, \quad B \xrightarrow{\Phi} f, \quad f \xrightarrow{B} \Phi$$

where we understand appropriately that these mappings select *other* mappings, as discussed in Figure 2. We note that *each* of the functions f , B , and Φ , is itself entailed by one of the others, and thus the system is an *organism*. We can describe f as a *metabolic* function, which is in turn selected, based in its outputs, by the *repair* function Φ . Finally, the *replication* function B selects the repair function based on *its* outputs.

CONSEQUENCES FOR BIOLOGY

We see that for Rosen the distinction between machines and organisms is the same as that between purely syntactic models and those which capture the semantic relations of measurement and natural language. This whole program, derived from the “relational biology” of Rashevsky, represents a distinctly heretical challenge to biology, dedicated as it is to the machine metaphor of organisms. This has many consequences which are just hinted at in the final chapter of *Life Itself*.

Rosen vehemently attacks some of the fundamental tenets of modern biochemistry. For example, he discusses the definition of phenotype as simply the immediate chemical results of genetic translation, thus reducing biological form and function to chemistry:

On the face of it, [this] is an astonishing claim. It is not adduced on the ground of great success in faithfully translating anatomical, embryological, or physiological processes into a syntactic chemical language. Quite the contrary, in fact, is true. It is adduced, rather, primarily on the grounds that otherwise, we simply could not answer “why?” questions about these processes with a “because . . .” framed exclusively in terms of *intrinsic*, fractionable (i.e., chemical) structure. That is, unless we *identify* phenotype with biochemistry, we can no longer claim that the functional genetic factors originally posited by Mendel can be identified with fractionable pieces of chromosomal structure. If [this identification] fails, then mechanism fails; but as we have stressed, the alternative is not vitalism, it is complexity. [pp. 261–262]

Rosen’s discussion of evolutionary biology is especially cogent. Traditional biologists are caught between the requirements of mechanism and anti-vitalism: organisms must be regarded as deterministic, reducible machines, while the abundance of evolutionary change must not be regarded as the result of extra-physical processes. Evolution serves to rescue biology from mechanism, but at the cost of itself being entirely beyond *any* explanation, physical or otherwise. Rather, life and evolution must be seen as a monumental chronicle of accidents.

If we *did* admit entailment into the evolutionary realm, then only two alternatives seem visible: (1) these entailments are themselves mechanistic, in which case biology disappears back into mechanism again, and loses forever its distinguished character, or (2) these entailments are not mechanistic, which seems to mean they must be Vitalistic again. Hence we are *driven* to expunge entailment from evolution entirely, not on any intrinsic scientific grounds, but because of the psychological requirements of biologists. [p. 257]

Instead, we must recognize that the complexity of non-simulable systems allows us to regain a biology of *soma* and in turn allows a description of evolutionary processes in terms of natural laws at the level of *morphological* phenotypes (in the spirit of Rene Thom and D'arcy Thompson).

Rosen reasserts the view from cybernetics about the orthogonality—the complete separability—of structure and function. While it is well understood that structure is not entailed by function (there are multiple structures capable of performing any function), Rosen makes the stronger claim that (in general) function is not entailed by structure in the sense that one can compute such functions. This is (at least) because of complexity issues, but also because individual components can participate in multiple functions.

This has significant implications in protein folding for example, where Rosen asserts that there can be neither an algorithm to calculate tertiary structure from primary structure nor the *biological* function of “active sites” from tertiary structure. Indeed, genetic sequences and active sites in protein structures are therefore primary (and common) examples of true “emergence”: entirely new entities (informational sequences and the chemical properties of the interaction of binding sites respectively) whose properties are completely non-inherent in the properties and entities of the underlying system.

In the end, *Life Itself* is a strong denial of the Church-Turing thesis, seen in this case as the statement that all models are mechanisms, and thus all systems are simulable. These arguments are very similar to those of Kampis (1991). Proteins are in the domain of “component systems”, which are inherently complex and unyielding to dynamical models. Rosen and Kampis also show a similar use of the concepts of complexity, chaos, and the immensity of computational models of truly complex systems.

Life then entails the presence of complexity and the absence of simplicity, while mechanisms entail simplicity, and thus the absence of life. However, complexity is not a *sufficient* condition for life. At the end of this stunning book, the reader is left wondering what *other* kinds of systems might exist between simple mechanisms and the most complex organisms, and what natural correlates might exist.

REFERENCES

- G. S. Brown, *Laws of Form*. Julian Press, New York, 1972.
- P. A. Cariani, *On the Design of Devices with Emergent Semantic Functions*. Ph.D. Dissertation, SUNY-Binghamton, Binghamton NY, 1989.
- F. Heylighen, “A Structural Language for the Foundations of Physics.” *Int. J. General Systems*, 18, 1990, pp. 93–112.
- C. Joslyn, “Software Support for Principia Cybernetica Development.” In: *Workbook of the First Principia Cybernetica Workshop*, edited by F. Heylighen, Principia Cybernetica, Brussels and New York, 1991, p. 49.
- G. Kampis, *Self-Modifying Systems*. Pergamon, Oxford, 1991.
- G. Klir, *Architecture of Systems Problem Solving*. Plenum, New York, 1985.
- G. Klir, *Facets of Systems Science*, Plenum, New York, 1991.
- H. H. Pattee, “Simulations Realizations, and Theories of Life.” In: *Artificial Life*, edited by C. Langton, Addison-Wesley, Redwood City CA, 1988, pp. 63–77.

R. Rosen, *Anticipatory Systems*, Pergamon, Oxford, 1985.

V. F. Turchin, *The Phenomenon of Science*, Columbia University, New York, 1977.

CLIFF JOSLYN
Systems Science, SUNY-Binghamton
327 Spring St. #2
Portland, ME 04102 USA
cjoslyn@bingsons.cc.binghamton.edu
joslyn@kong.gsfc.nasa.gov

EVIDENCE THEORY AND ITS APPLICATIONS, Volume 1, by Jiwen Guan and David A. Bell. North-Holland, Amsterdam and New York, 1991. viii + 351 pages.

Prior to the publication of this book, the only book fully devoted to evidence theory had been the classic book by Shafer [1976], in which evidence theory was introduced. Since 1976, interest in evidence theory, which is now often referred to as Dempster-Shafer theory, has been rapidly growing [Shafer, 1990].

Although the book by Guan and Bell is heavily built on the seminal book by Shafer, as acknowledged by the authors themselves, it also includes significant new developments in the theory. Well presented are important contributions by Barnett [1981] and Yen [1986, 1989], as well as some original contributions by the authors themselves.

The book consists of ten chapters. Chapter 1 contains required prerequisites from set theory and introduces mathematical notation employed in the book. A special emphasis is given to Mobius inversions, which are essential in investigating the relationship among different evidential functions.

Basic functions of evidence theory are introduced in Chapter 2: mass functions (basic probability assignments), belief functions, plausibility functions, and comonality functions. It is shown that there exists a unique transformation between any pair of these functions and, consequently, each of these functions alone is sufficient to characterize a particular body of evidence. It is also shown that evidence theory is a generalization of Bayesian theory of probability.

Chapter 3 is devoted to a thorough discussion of the principal operation for combining evidence in evidence theory, which is usually referred to as Dempster rule. The rule is also compared with the Bayesian rule of conditioning.

Chapters 4–7 deal with two special classes of mass functions, which are particularly suitable for representing evidence in many applications. These are called simple mass functions and separable mass functions. Included are simplification algorithms, canonical decompositions, and special rules of combination that are applicable to these classes of functions.

Issues of computational complexity in evidence theory are discussed in Chapter 8. The main purpose of the chapter is to elaborate a methodology proposed by Barnett [1981], by which evidence can be evaluated in linear time. The methodology is based on partitioning the problem space in several independent ways and clustering evidence in the partitions.

Chapters 9 and 10 deal with various issues of using evidence theory for representing uncertainty in expert systems. The material is heavily based on ideas developed by Yen [1986, 1989].