THE ONTOLOGY OF CYBERSPACE:
Questions and Comments

David R. Koepsell

Department of Philosophy
and School of Law
State University of New York at Buffalo
Buffalo, NY 14260

koepsell@acsu.buffalo.edu

William J. Rapaport

Department of Computer Science
and Center for Cognitive Science
State University of New York at Buffalo
Buffalo, NY 14260

rapaport@cs.buffalo.edu
http://www.cs.buffalo.edu

Abstract

This document consists of two papers: "The Ontology of Cyberspace: Preliminary Questions", by David R. Koepsell, and "Comments on 'The Ontology of Cyberspace'," by William J. Rapaport. They were originally presented at the Tri-State Philosophical Association Meeting, St. Bonaventure University, 22 April 1995. This document is SUNY Buffalo Department of Computer Science *Technical Report 95-25* and SUNY Buffalo Center for Cognitive Science *Technical Report 95-09*.

THE ONTOLOGY OF CYBERSPACE:
Preliminary Questions

David R. Koepsell

Department of Philosophy and School of Law
State University of New York at Buffalo
Buffalo, NY 14260

koepsell@acsu.buffalo.edu

Not since the first work of fiction was produced have philosophers been confronted with a totally unexplored realm of ontological inquiry. The development of the computer and the concept of cyberspace have ushered in such a realm. Cyberspace is something whose nature remains largely unexplored despite its ripeness for ontological investigation. This is problematic for a number of reasons as shall be seen.

The concept of cyberspace is a vague one. Some confuse it with the notion of virtual reality (VR). VR, however, is but a small aspect of cyberspace as shall be shown. Even without considering the special problems posed by VR, a number of ontological questions can be posed regarding the goings-on within computers. The following questions should be explored and the implications of their various answers likewise considered: What is cyberspace? Is it or does it have dimension? Are there things in cyberspace? Are things in cyberspace properly called objects? Are such objects or is cyberspace itself substance(s) or process(es)? Is cyberspace or the objects in it real or ideal? What is the categorical scheme of cyberspace? How should cyberspace fit into a broader categorical scheme? Addressing these questions should complete a comprehensive ontology of cyberspace.

A comprehensive ontology of cyberspace is important both philosophically and sociologically. Philosophers desiring to understand reality typically do not exclude whole categories of phenomena from their consideration. On the contrary, ontologists must take account of as many phenomena as possible (at least those which are open to ordinary experience) in formulating an ontology. For this reason alone the realm of cyberspace should be considered. But other benefits may arise from an ontological consideration of cyberspace. Examining cyberspace offers the rare opportunity of a relatively constrained world of phenomena. Simply put, there is only as much cyberspace as we create. Such a neatly delineated realm offers the possibility for in-depth and exhaustive study unavailable in the world of ordinary experience over which we have limited control. An exhaustive ontology of cyberspace may serve as a sort of manageable laboratory for ontological explorations. Ontological methodologies may perhaps be more easily tested in such a lab. Finally, inquiry into the ontology of cyberspace may shed new light on ontology in general.

There are a number of sociological reasons for a comprehensive ontology of cyberspace. Computers and computer networks are becoming ubiquitous. As these tools permeate our lives so do the questions they pose and the conflicts to which they give rise. How shall we treat cyber-objects philosophically and practically? This question is now a very real problem. For instance, the question of whether computer software is patentable is one of real social importance. The implications of our ontological decisions regarding cyberspace are too potentially far-reaching to be made without some deep reflection. Cyber-objects are proliferating rapidly and how we treat them must ultimately depend upon the ontology of cyberspace which we accept. Furthermore, should artificial intelligence research (AI) ever make some real in-roads, the ontology we accept will inform how we treat artificial intelligences.

1 A GRAPHIC EXAMPLE.

The problem of software and the attempts by the law to deal with it as intellectual property underscores the practical importance of working out the ontology of cyberspace.

What is software? Federal law defines software as a "set of statements or instructions to be used directly or indirectly in a computer to bring about a certain result."[1] Software consists of binary code which, when used by a computer, causes the computer to do certain things. The things that the computer does depend upon how the binary code is organized.

Software forms the basis for all computerized phenomena. A computer without software consists of nothing but switches which may be either on or off. Without some form of instruction for these switches, either etched into Read Only Memory (ROM) or loaded into Random Access Memory (RAM), the switches would sit idle. But once software is introduced to these switches by way of ROMs or loaded from such media as magnetic disks or tapes into RAM, the computer can perform any number of functions.

Software is both like a written work (such as a book) and like a machine. It is like a book in that it is composed of symbolic information. It is like a machine in that it can perform certain differing functions. It is this dual nature which has confounded the courts' treatment of software.

Copyright law has typically applied to written works and protects, for a limited period of time, an author's interest in his or her original works to the extent that they are in a substantially concrete form. Thus, the source code of a computer program, which is written in such languages as C, Pascal, Basic, etc., is copyrightable. However, the object code, which is what the source code is compiled into by a computer in order to be executed by it, is not so clearly copyrightable:

> It has been suggested even if the object code is analogized to a recording of a phonograph record or tape in a form that the machine can recognize and play back—such items being clearly copyrightable—the transportation of the binary code into a circuit design that replicates the on-off switching of the binary form object code using [various]...processes raises serious questions of copyright protectability since the work is not a "writing" but is, in microchip form, part of a machine, a utilitarian object, and hence not copyrightable.[2]

Machines and other such "utilitarian objects" are normally afforded the protection of patent law. "Whoever invents or discovers any new and useful process, machine, manufacture, or composition of matter, or any new and useful improvement thereof, may obtain a patent therefor."[3] A patent essentially grants one the right to exclude others from producing, using, or selling an invention during a specified time period. Historically software has not been afforded patent protection. This is because it is not clear that software is always either a "process, machine, manufacture, or composition of matter." But, some patent protection may be available to software. For example, "the method which instructs a computer to perform its operating functions as distinguished from the instructions themselves, is said not to be subject to copyright but is protected, if at all, by patent law."[4] This is because such methods are perhaps most akin to processes, which are generally patentable.

---

[1] 17 U.S. Code §101.

[2] 18 American Jurisprudence (Second), Copyright §51; see Ray A. Mantle, "Trade Secret and Copyright Protection of Computer Software," 4 *Computer Law Journal* 669 (1984).

[3] U.S. Patent Act §101.

[4] 18 American Jurisprudence (Second), Copyright §52; see *Apple Computer, Inc. v. Franklin Computer Corp.*, 714 Federal Rep. 2d 1240 (3d Circuit), *cert. dismissed*, 104 Supreme Ct. Reporter 690 (1983).

Is software an idea, a written work, a process, or a machine? It is not clear what, in fact, it is. It may, at different times, be any one or all of these, or it may be something altogether different. It is at least clear that categorizing software is not only problematic theoretically, but that this problem results in practical difficulties as well. The legal system is working out these problems in accord with its concerns for practicality and economic efficiency, but the philosophical underpinnings of the problem remain unresolved.

What follows is an introduction to the problem of developing a comprehensive ontology of cyberspace. A brief history of the development of cyberspace precedes the philosophical discussion. It is hoped that a historical context can help clarify the territory and aid in setting forth some key definitions. Next, a review of the scant literature which currently exists on the topic will be made. It will be shown that no serious philosophical works have been produced which satisfactorily ask or answer the questions set forth at the beginning of this paper.

To date, the literature fails to seriously address the problem of the ontology of cyberspace because, mostly, of a failure to understand ontology as a discipline. Thus, we shall look at what constitutes ontological investigation and ask how the literature on cyber-ontology fails to meet the criteria for such inquiry.

Finally, a framework for a proper (or at least careful) ontological investigation of cyberspace shall be proposed. This framework is one which has evolved already through the law of intellectual property and which is being applied by jurists to resolve the problems associated with software.

2 CYBERSPACE.

The term "cyberspace" is a new one, coined by the science- fiction novelist William Gibson. As used in his "cyberpunk" books (a genre more or less invented by him) the term refers to what he calls a "consensual hallucination" which futuristic computer network users experience when "jacking-in" to a world-wide network. However, the term has come into popular usage as a name for the Internet. This term need not be confined simply to networks. If cyberspace is considered to apply to all phenomena occurring electronically "within" and among computers, the term's usefulness for ontological inquiry is clear. The question remains open: is cyberspace an existence or an occurrence at all or something quite unique? For now the word "cyberspace" serves a valuable purpose as a term of entry into the topic at hand.

We may say that cyberspace developed along with computers. Any machine which has relied upon electronic switching to store and process operations can be said to be a computer. In *The Metaphysics of Virtual Reality*, Michael Heim offers the following explanation: "cyberspace suggests a computerized dimension where we move information about and where we find our way around data."[5] The very use of the word "dimension" shows that this definition is imprecise and begs the ontological questions set forth above, but it captures a sort of general understanding of the notion of cyberspace. Curiously, the term cyberspace is ordinarily applied only to digital-computerized phenomena. One gets the notion that had Babbage's Difference Engine formed the basis of all modern computerization, using mechanical rather than electronic switches, there would be no cyberspace. This problem may require consideration eventually as we try to unravel the ontological questions.

[5] New York: Oxford University Press, 1993, pp. 77–78.

Electronic switching, first through vacuum tubes, then transistors, and now through silicon chips, forms the ultimate basis for computerized phenomena. Computers are the medium for an increasing number of information-transactions. For the limited purpose of clarifying the term "cyberspace", these transactions shall be said to occupy, occur, or exist in cyberspace. Programming is such a transaction. Digital communication is another. The computations that computers carry out as a result of programming are such transactions. The term "cyberspace", for now, shall refer to the complex of the electronic switches and the information-transactions that occur by way of these switches within and among computers. E-mail exists and moves in cyberspace. Computer programs exist and function within cyberspace. Virtual Reality exists and occupies cyberspace. Financial transactions often occur in cyberspace.

3 THE ONTOLOGICAL PROBLEM.

No serious philosophical approaches to the ontology of cyberspace have been made to date. This is not to say that it is totally undiscovered country. In fact, a few have noted that there are philosophical implications of cyberspace open for exploration.

> Cyberspace is more than a breakthrough in electronic media or in computer interface design. With its virtual environments and simulated worlds, cyberspace is a metaphysical laboratory, a tool for examining our very sense of reality.[6]

Unfortunately, the most promising title in this study, Heim's *The Metaphysics of Virtual Reality*, deals only cursorily and superficially with the ontology of cyberspace. Chapter seven of this book, entitled "The Erotic Ontology of Cyberspace" sets forth the problem quite succinctly: "We need to give an account of (1) the way entities exist within cyberspace and (2) the ontological status of cyberspace—the construct, the phenomenon—itself."[7] Heim works hastily through a number of hypotheses including 1) that "Platonism provides the psychic makeup for cyberspace entities"[8] and 2) that Leibniz's "logic, metaphysics, and notion of representational symbols show us the hidden underpinnings of cyberspace."[9] But Heim fails to back up these already vague hypotheses with any sound reasoning or in-depth discussion. Rather, we are asked to accept such statements as:

> The Central System Monad is the only being that exists with absolute necessity. Without a sysop, no one could get on line to reality. Thanks to the Central System Monad, each individual monad lives out its separate life according to the dictates of its own willful nature while still harmonizing with all the other monads on line.[10]

Ultimately, Heim does not answer the ontological questions he sets forth. This lapse is forgivable in the context of the book as a whole. *The Metaphysics of Virtual Reality* is primarily a popularized account of various sociological and psychological concerns which arise with this new medium. The book focuses on virtual reality instead of cyberspace as a whole. In fact, it confuses these two terms throughout. But it is a work directed at, and emerging from the popular culture surrounding cyberspace. Its concerns are not rigorously philosophical.

---

[6] Heim, p. 83.
[7] *Ibid.*, p. 84.
[8] *Ibid.*, p. 91.
[9] *Ibid.*, p. 92.
[10] *Ibid.*, p. 99.

What then are the philosophical concerns that must be addressed in an ontological investigation of cyberspace? Answering this question requires an agreement on the role and methods of accepted ontology. In practice and over time, various ontologies and their methods have differed markedly. But some philosophers have attempted to generalize about the aims of all ontology as well as to formulate schemes for its accomplishment. There is at least a general consensus among philosophers that ontology is the study of, variously: existence, being, reality and/or the meanings of each of these words.[11] Thus, an ontology of cyberspace would consider these questions as applied to the phenomena which we agree to constitute or comprise cyberspace. Although Heim fails to adequately address these concerns, the problem is clear given an understanding of the general subject of ontology.

4 THE LEGAL FRAMEWORK FOR THE ONTOLOGY OF CYBERSPACE.

While philosophers have not yet adequately addressed the ontological problems presented by cyberspace, the legal system has been grappling with the practical problems raised by the emergence of computerized media. Intellectual property law has developed a useful categorical scheme which may be rather readily used as a point of departure for an ontological study of cyberspace. The usefulness of this categorical scheme is partly due to the fact that the law of intellectual property has been built around a crude ontology. What follows is the legal categorical scheme for intellectual property ranging from those objects which are afforded the least legal protection to those which are afforded the most:

NOT COPYRIGHTABLE OR PATENTABLE

1. *Ideas*, procedures, processes, systems, concepts, laws of nature, principles, information, and utilitarian works. All of these are generally outside the scope of legal protection.

COPYRIGHTABLE

2. *Expressions.*
   (a) Non-literal elements: The "fundamental essence or structure" such as unique settings, characters, plot, etc. These are afforded copyright protection to a certain degree. Such protection often excludes parodies.
   (b) Literal elements: The expression of original ideas as distinguished from the ideas themselves are protected inasmuch as they are reduced to a tangible medium of expression from which they can be perceived, reproduced, or otherwise communicated either directly or with the aid of a machine or device..

TRADE SECRETS

3. *Certain formulas*, patterns, devices or compilations of information used in a business which give a competitive advantage over competitors are afforded trade secret protection. Trade secrets are not protected against independent discovery. A trade secret need not achieve the same level of advancement needed for patent protection.

[11] See, e.g., Williams, C. J. F., *What is Existence?* (Oxford: Clarendon Press, 1981), p. 1; Johansson, Ingvar, *Ontological Investigations* (New York: Routledge, 1989), p. 1; Hartmann, Nicolai, *New Ways of Ontology* (Chicago: Henry Regnery Co., 1953), p. 11.

PATENTABLE

4. *Utilitarian objects*,[12] machines, processes, manufactures or compositions of matter. These are all afforded patent protection which gives the patent holder the exclusive right to the art disclosed whereas copyright protects only the expression of an idea. Patents protect the means of reducing an inventive idea to practice.

The categorical scheme above seems to reflect a certain ontology or at least a recognition of distinction between types and tokens. Thus, this scheme will serve adequately as a foundation for a foray into the ontological investigation of cyberspace. It is a ready-made bridge constructed by the legal system between the practical and theoretical problems posed by computerized media.

The courts have begun the theoretical task by trying to fit software into this scheme as noted above. The problems and reasoning applicable to software are applicable by extension to the phenomena associated with the Internet and virtual reality. The questions set forth at the beginning of this paper may all be addressed within the categorical scheme developed by the law to deal with intellectual property.

5 CONCLUSION.

What has been presented above is an attempt to outline a research program for the investigation of the ontology of cyberspace. The goal of this study is clear: a comprehensive ontology of cyberspace. Such an ontology should provide answers for the questions at the beginning of this paper.

By following the reasoning used by the courts and already applied to intellectual property disputes regarding software, and by analyzing the legal categorical scheme, a useful and comprehensive ontological approach to cyberspace may be developed.

[12] Utilitarian objects are tangible things. Utilitarian works are expressions of ideas with specific uses, such as recipes, instructions, etc.

COMMENTS ON "THE ONTOLOGY OF CYBERSPACE"

William J. Rapaport

Department of Computer Science and Center for Cognitive Science
State University of New York at Buffalo, Buffalo, NY 14260

rapaport@cs.buffalo.edu, http://www.cs.buffalo.edu

In "The Ontology of Cyberspace" (1995), David R. Koepsell raises some interesting, exciting, and important questions, and suggests directions for answering them. I will begin my comments with some observations on cyberspace as a domain of ontology, I will then offer some quibbles about Koepsell's two sample ontological investigations (on the ontological status of software and of hardware), and I will conclude with some remarks on methodology.

1 CYBERSPACE AS A DOMAIN OF ONTOLOGY.

Koepsell says that "Examining cyberspace offers the *rare* opportunity of a relatively constrained world of phenomena. ... [T]here is only as much cyberspace as we create" (p. 2; my emphasis). Now, I *agree* with the artifactual nature of cyberspace, although the limitation on its size or quantity to what "*we* create" might not be accurate: The apparent large size, complexity, and topological structure of the World Wide Web, for example, might be an emergent property of a small, in any case limited, number of computers linked in combinatorially explosive ways. Though it has its origins in us, it can take on a life of its own.

I *disagree*, however, with its rarity or uniqueness. For the same can be said of fiction, that is, that it is "a relatively constrained world of phenomena", with "only as much" of it "as we create". To this pair, we might add individual mental worlds (or mental models of the world) as another such category of ontological domain. Koepsell also cites questions about the ontological status of "artificial intelligences" (p. 3), by which I take it he means computational cognitive or rational agents. In addition, though not explicitly cited by him, there are interesting and important ontological (and moral) questions raised by the new discipline of Artificial Life (see, e.g., Langton 1989, 1994; Meyer & Wilson 1991; Langton et al. 1992; Varela & Bourgine 1992; Meyer et al. 1993; Brooks & Maes 1994; Cliff et al. 1994).

This suggests that, far from being "rare", there are quite a few "relatively constrained worlds of phenomena" that are open to ontological investigation: Besides the real world (which, of course, includes all the others), there are, in addition to cyberspace, individual mental worlds and fiction. And cyberspace itself seems open to at least three, *prima facie* distinct sorts of ontological questions: (1) questions about "artificial intelligences" and artificial life, (2) questions about the *legal* status of cyberspatial entities, and (3) questions about what I shall call their *logical* (or "purely" ontological) status, such as (among others) (a) the relation of software to hardware (and the more general question of the exact nature of the implementation relation), (b) the relation of (static, textual) program to (dynamic, causal) process, and (c) the (real or perceived) topology of the Web. (The questions Koepsell raises on page 1 fall in this category, too.) There is also a fourth "relatively constrained world of phenomena" open to ontological investigation, which I will mention a bit later.

2 SAMPLE ONTOLOGICAL INVESTIGATIONS.

By way of examples, Koepsell discusses the ontological status of software and has brief remarks on the ontological status of hardware. Although he raises interesting questions, I disagree with some of his claims.

2.1 Ontological Status of Software.

First, I think he is mistaken in his analysis of software. He says that "software consists of binary code ..." (p. 3). This, however, is not necessarily the case: Programs written in high-level programming languages need not be compiled into binary machine languages (witness the Lisp machines of a few years ago). In any event, surely a Pascal or a C program falls under copyright or patent laws just as much as any binary machine-language implementation of them. And sometimes it is the "look and feel" that's copyrightable, rather than either the high-level program or the low-level binary code *per se*.

Koepsell also likens software to hardware in a way that I take issue with. He says that software has a "dual nature", being "like a book [that is, a written work] in that it is composed of symbolic information" and being "like a machine in that it can perform certain differing functions" (p. 3). However, although it *is* like a book, it is *not* like a machine. Software can't perform anything, any more than a book can. Rather, it is the central processing unit of the computer, with its fetch-and-execute cycle, that can perform things, depending on the "switch settings" determined by the software (more precisely, the values stored in certain registers when the program is loaded).

Koepsell, citing a legal text, says that the software " 'in microchip form [is] part of a machine' " (p. 4). But being part of a machine doesn't mean that it *is*, or even is *like*, a machine. And, in any case, the "microchip form" is merely the settings of the switches—it is the machine in a certain state, a state in which it was put by the loading of the software.

2.2 Ontological Status of Hardware.

The second claim about which I think Koepsell is mistaken is his analysis of (or, more precisely, is a remark he makes about) hardware. He says that "Any machine which has relied upon *electronic* switching to store and process operations can be said to be a computer" (p. 6; my emphasis). By this criterion, a telephone network would be a computer. Now, although this claim is wrong, it does afford an exciting insight into the ontology of cyberspace. It is wrong, on the one hand, because the term 'computer' should be reserved for *anything*—human or machine; electronic, mechanical, or otherwise—that computes in the sense of Turing. Whether something is a computer or not should have nothing to do with its implementation in electronic switches or any other medium. But it is exciting, on the other hand, because the telephone network may provide an interesting source of analogies for cyberspace, such as the ability to access distant places virtually instantaneously. Upon reflection, this is not surprising, given electronic computers' dependence on phone lines as their communication channel.[13] So I suggest we add the telephone network as the fourth item on our list of "relatively constrained worlds of phenomena".

3 METHODOLOGY.

Let me conclude with two brief remarks on methodology. First, Koepsell says that "Cyberspace is something whose nature remains largely unexplored ..." (p. 1). While more exploration *is* needed, I think more has been done than Koepsell's adverb 'largely' suggests. There have been lots of explorations of the ontological status of computer software not only with respect to its legal status (e.g., there are several texts, such as Johnson 1985, Gould 1989, Ermann et al. 1990, Forester & Morrison 1990, Dunlop & Kling 1991; and there is a regular column, "Legally Speaking", by Pamela Samuelson, in *Communications of the ACM*) but also in connection with Searle's Chinese-Room Argument (e.g., Eric Dietrich's recent *Thinking Computers and Virtual Persons* (1995)). And the ontological status of artificial intelligences (and artificial life) has surely been explored (if not exhaustively) in science fiction (especially Isaac Asimov's writings on robotics and Stanislaw Lem's prescient moral fable about artificial life, "Non Serviam" (1971)), in the pages of, e.g., *AI Magazine* (not to mention other professional journals and conference proceedings), and in Joseph Weizenbaum's classic *Computer Power and Human Reason* (1976).

Second, a problem with Koepsell's methodology is that he *seems* to be using legal citations in an effort to understand the ontology of cyberspace. But it should go the other way round: Computer scientists and philosophers should provide guidance on the ontological issues, in order to aid legal interpretation. Now, to be fair, Koepsell recognizes this problem when he says, "The legal system is working out these problems in accord with its concerns for practicality and economic efficiency, but the philosophical underpinnings of the problem remain unresolved" (p. 5). However, he goes on to suggest (in the section on "The Legal Framework for the Ontology of Cyberspace") that "the legal categorial scheme" might help in the purely philosophical enquiry. I repeat that I don't think this is the right way to begin (though it's certainly convenient). Why constrain philosophical ontological enquiries to fit within a categorial scheme motivated by economics and pragmatics rather than truth? It would be far better to get philosophically-inclined computer scientists and computationally-trained philosophers to set the ground rules for the lawyers. I look forward to future endeavors—by Koepsell and others—in this regard.

---

[13] The matter is actually rather more complicated. Roughly, local computers are more-or-less directly connected to each other; computers communicating over longer distances use dedicated phone lines and satellite links (not necessarily "the" phone system) and something called "virtual circuits":

> Within the [Computer Science] department, everything uses ethernet. Ethernet is a protocol which runs on any of a variety of physical media—the ones in use include coaxial cable, and twisted pair cable (which is much like regular copper phone cable). However, at the local level, nothing travels over real phone company cables.
>
> The various buildings on campus are connected using a protocol called FDDI, running over fiber-optic lines.
>
> Campus is connected to the world by what is called a T1 link. This is a dedicated phone line, running at 14Mbps (which is basically the capacity of 24 regular dedicated phone lines). This is over NYNEX phone wire, but is dedicated, in that it is not dialled when needed, but is permanantly connected. This connects us to SprintLink. (There is talk of them upgrading this connection to a T3 link, which runs at about 43Mbps!)
>
> SprintLink is Sprint's Internet backbone. It probably uses a variety of media, and a variety of protocols, as it transfers data around the country (probably including FDDI, T3 and/or ATM over fiber and over copper). (Davin Milun, personal communication, 11 April 1995.)

REFERENCES

1. Brooks, Rodney A., & Maes, Pattie (eds.) (1994), *Artificial Life IV: Proceedings of the 4th International Workshop on the Synthesis and Simulation of Living Systems* (Cambridge, MA: MIT Press).

2. Cliff, Dave; Husbands, Philip; Meyer, Jean-Arcady; & Wilson, Stewart W. (eds.) (1994), *From Animals to Animats 3: Proceedings of the 3rd International Conference on Simulation of Adaptive Behavior* (Cambridge, MA: MIT Press).

3. Dietrich, Eric (ed.) (1994), *Thinking Computers and Virtual Persons: Essays on the Intentionality of Machines* (San Diego: Academic Press).

4. Dunlop, Charles, & Kling, Rob (eds.) (1991), *Computerization and Controversy: Value Conflicts and Social Choices* (Boston: Academic Press).

5. Ermann, David; Williams, Mary B.; & Gutierrez, Claudio (eds.) (1990), *Computers, Ethics, and Society* (New York: Oxford University Press).

6. Forester, Tom, & Morrison, Perry (1990), *Computer Ethics* (Cambridge, MA: MIT Press).

7. Gould, Carol C. (ed.) (1989), *The Information Web: Ethical and Social Implications of Computer Networking* (Boulder, CO: Westview Press).

8. Johnson, Deborah G. (1985), *Computer Ethics* (Englewood Cliffs, NJ: Prentice-Hall).

9. Koepsell, David R. (1995), "The Ontology of Cyberspace: Preliminary Questions," paper presented at the Tri-State Philosophical Association Meeting, St. Bonaventure University, 22 April 1995.

10. Langton, Christopher G. (ed.) (1989), *Artificial Life: The Proceedings of an Interdisciplinary Workshop on the Synthesis and Simulation of Living Systems (September 1987, Los Alamos)* (Reading, MA: Addison-Wesley).

11. Langton, Christopher G. (ed.) (1994), *Artificial Life III: Proceedings of the Workshop on Artificial Life (June 1992, Santa Fe)* (Reading, MA: Addison-Wesley).

12. Langton, Christopher G.; Taylor, Charles; Farmer, J. Doyne; & Rasmussen, Steen (eds.) (1992), *Artificial Life II: Proceedings of the Workshop on Artificial Life (February 1990, Santa Fe)* (Reading, MA: Addison-Wesley).

13. Lem, Stanislaw (1971), "Non Serviam," in S. Lem, *A Perfect Vacuum*, trans. by M. Kandel (New York: Harcourt Brace Jovanovich, 1979).

14. Meyer, Jean-Arcady, & Wilson, Stewart W. (eds.) (1991), *From Animals to Animats: Proceedings of the 1st International Conference on Simulation of Adaptive Behavior* (Cambridge, MA: MIT Press).

15. Meyer, Jean-Arcady; Roitblat, Herbert L.; & Wilson, Stewart W. (eds.) (1993), *From Animals to Animats: Proceedings of the 2nd International Conference on Simulation of Adaptive Behavior* (Cambridge, MA: MIT Press).

16. Varela, Francisco J. & Bourgine, Paul (eds.) (1992), *Toward a Practice of Autonomous Systems: Proceedings of the 1st European Conference on Artificial Life* (Cambridge, MA: MIT Press).

17. Weizenbaum, Joseph (1976), *Computer Power and Human Reason* (San Francisco: W. H. Freeman).