

# A Parser as an Epistemic Artifact: A Material View on Models

Tarja Knuuttila and Atro Voutilainen<sup>†</sup>

---

The purpose of this paper is to suggest that models in scientific practice can be conceived of as *epistemic artifacts*. Approaching models this way accommodates many such things that working scientists themselves call models but that the semantic conception of models does not duly recognize as such. That models are epistemic artifacts implies, firstly, that they cannot be understood apart from purposeful human activity; secondly, that they are somehow materialized inhabitants of the intersubjective field of that activity; and thirdly, that they can function also as knowledge objects. We argue that models as epistemic artifacts provide knowledge in many other ways than just via direct representative links. To substantiate our view we use a language-technological artifact, a parser, as an example.

---

**1. Introduction.** In the philosophical literature on scientific models, two different currents can be discerned. On one hand there are efforts to establish, more or less formally, what scientific models are. The syntactic view on models, once the “received view,” and the semantic conception of models are both attempts of this kind. On the other hand, there are approaches concerned with different roles of models in scientific practices. From the very outset, these two views seem to have had competing goals (Bailer-Jones 1999). Whereas the first approach tries to give a unitary account of what models are, the other approach is interested in models in their very diversity. These goals appear to be derived from fundamentally different ways of approaching models in science, which can be seen once we compare the claims made by the proponents of both views.

Recently, da Costa and French (2000) and French and Ladyman (1999) have argued that the semantic approach to models can accommodate their diversity and their uses in science by representing them set-theoretically in

<sup>†</sup>To contact the authors, please write: Tarja Knuuttila, Department of Philosophy, P.O. Box 47, FIN-00014, University of Helsinki, Finland; e-mail: tarja.knuuttila@helsinki.fi. Atro Voutilainen, Department of General Linguistics, P.O. Box 9, FIN-00014, University of Helsinki, Helsinki, Finland; e-mail: Atro.Voutilainen@helsinki.fi.

terms of partial structures (da Costa and French 2000, 123–125). Moreover, French and Ladyman claim that “the specific material of the models is irrelevant; rather it is the structural representation, in two or three dimensions, which is important” (1999, 109). This can be contrasted with the attempt of Morgan, Morrison, and others to study the process of constructing and manipulating models, which is, in their opinion “crucial in gaining information about the world, theories and the model itself” (Morrison and Morgan 1999a, 8). Thus, while the adherents of the semantic conception attempt to represent models in science as more or less steady and as ready-made entities, the proponents of the practice-oriented approach are interested in the modeling process and in explaining why and how models are used in scientific endeavor.

From the point of view of the modeling process one can hardly dismiss the materiality of models as irrelevant. In this paper we argue that models can be treated as *epistemic artifacts* from the scientific practice point of view. This approach stresses the importance of their materiality for scientific research. As epistemic artifacts, scientific models are open to different interpretations and uses, functioning as both tools and objects of inquiry. Last, but not least, the conception of models as epistemic artifacts can accommodate many actually fabricated things that scientists themselves call models, but which the prevailing semantic conception fails to recognize as such—due to its predisposition to conceptualize models as abstract, theoretical entities. Our example of a parser, with which we substantiate our claims, is a case in point.

**2. Models in Scientific Practice.** Of the recent accounts of models, the one by Morrison and Morgan (1999b) seems to us a fruitful attempt to approach them from a practical point of view. In their approach, the workability and manipulability of models occupy a central position. Somewhat in the fashion of science and technology studies, Morrison and Morgan’s practice-oriented approach to models advocates the suspension of theorizing about models in favor of looking at how they are actually constructed, used, and conceived of in diverse scientific activities. Once we take a practical approach to models, the astonishing diversity of different kinds of things called models by scientists themselves becomes apparent. They can be physical (three-dimensional) objects, various mathematical structures, diagrams, computer programs, and so on. It seems, indeed, to be extremely difficult to try and say something general about so heterogeneous a group of things. Yet, despite their explicit reluctance to propose a theory of models, Morgan and Morrison have chosen to approach models as *mediators*.

To call models “mediators” means that they can be treated as “autonomous agents” that mediate between the theory and the world. It is the

independence of the models that makes them able to mediate. In arguing for this view, Morrison and Morgan distinguish four different angles from which to approach models. The four basic elements in their account of models as mediators are *construction*, *functioning*, *representation*, and *learning*. Morrison and Morgan begin their argument by claiming that because of their construction models gain their independence (at least partly) from theory and data, since, besides being composed of both theory and data, models typically involve also “additional ‘outside’ elements” (1999b, 11). Once independent, models can mediate in different ways. They can function as tools, on account of their autonomous nature. But the models that are used in science are actually more than just instruments: They are “investigative instruments.” To be a tool of investigation, an investigative instrument, is to involve some form of representation. Learning, for its part, is dependent on representation. According to Morrison and Morgan we can learn from models because they represent. But Morrison and Morgan make it clear that they do not conceive of representation in a traditional way as “mirroring” or as a correspondence. For them representation is “a kind of rendering—a partial representation that either abstracts from, or translates into another form, the real nature of the system or a theory, or one that is capable of embodying only a portion of a system” (1999b, 27).

We find Morrison and Morgan’s approach to models insightful but somewhat vague. Specifically, the important idea of models as mediators is left in the air, as the authors do not specify what they mean by mediation. Moreover, their analysis of representation points in at least two different directions. On one hand, they claim that we do not learn very much by looking at models but rather by building and manipulating them. On the other hand, they claim that models mediate between theory and world, as if both were stable entities between which models somehow provided a link. Seeing the world and the theory as separate domains in need of connection—or mediation—leads us to look for structural features common for both the model and the theory, and for the model and the world.

It seems that Morrison and Morgan leave things hanging when talking about “additional outside elements.” It is somewhat paradoxical that in their attempt to grant models an independent status as mediators, they end up articulating their view rather traditionally, relying on the categories of theory and data. However, in the very same collection Boumans (1999) makes the radical implications of Morrison and Morgan’s view explicit. He loosens the model from the grip of theory and data, thus making a model a truly independent entity.

Having studied three different business-cycle models, Boumans argues that models integrate a far broader range of ingredients than just “theory” and “data.” In his study a model is constructed out of many different in-

ingredients: analogies, metaphors, theoretical notions, mathematical concepts, mathematical techniques, stylized facts, empirical data, and finally, relevant policy views. Striving to match such diverse elements to one another tells us something interesting about modeling. It hints at the skill, experience, and hard work that are needed in it. The image of a scientist as a modeler is very different from that of a theoretical thinker. Boumans, in fact, likens model construction to baking a cake without a recipe (1999, 67).

**3. Models as Epistemic Artifacts.** Boumans's work exemplifies what is to us one of the most important insights of the practice-oriented approach to models in science, namely, to treat models as more or less complicated *epistemic artifacts*. That a model is an epistemic artifact implies first, that *human agency*, or rather traces of it, are more or less manifestly present in it. Second, it implies that models are somehow *materialized* inhabitants of the intersubjective field of human activity. Third, it implies that models can function also as knowledge objects.<sup>1</sup>

With regard to the importance of human agency for modeling, the modeling relation may appear to be dyadic, but it is, in fact, triadic. This can be seen once we ask in what sense models can represent real systems. Giere suggests that this be analyzed as a similarity between two objects, one abstract and one real (1988, 80–81, 93). But then, anything can be a model of anything else, as Giere notes as well. Any two things can always be brought into some relationship of similarity with each other. But this points to a characteristic of models that appears to be pivotal for our understanding of what a model is. When we choose something as a model of something else, we do it with some end in view. Our aims and goals explain which features we judge relevant and, consequently, how one thing can be *used* as a model of another thing (e.g., Wartofsky 1979, 6). Thus models cannot be understood without taking into account human agency.

The materiality of models, in turn, means that they are things that have their own construction and thus their own ways of functioning. Consequently, they are not open to all possible interpretations and uses. In other words, models, like all tools, have their own constraints and affordances. Yet, whether a certain property of a model is a constraint or an affordance is something that is relative to the use the model is put to. One might say that the constrained construction of a model promotes thinking, thus furthering research.

As we have seen, those writers taking the practice-oriented approach to models pay specific attention to the way we learn from building and

1. In developing our view we have benefited from the notion of epistemic thing presented by Knorr Cetina (1999) and Rheinberger (1997).

manipulating them. This is made possible by their materiality. Admittedly, it is somewhat against the philosophical tradition to treat models as materialized things inasmuch as philosophical literature has been predominantly interested in theoretical models, which, in turn, have been conceived of as abstract things. Let us, then, examine this purported abstractness of theoretical models.

Black characterizes theoretical models in the following way: We endeavor to understand some facts and regularities of an original domain with the help of “entities (objects, materials, mechanisms, systems, structures) belonging to a relatively unproblematic, more familiar, or better-organized secondary domain” (1962, 230). Theoretical models need not be built in the same sense as physical ones, yet they have to be described. The description of the model is not just any redundant realization of it made for the purposes of communication. In describing a model scientists try to deal with their problem and thus the theoretical model does not really come into being until it is actually described. Thus, we contest the idea that theoretical models dwell primarily as ready-made abstractions in the mind of an individual scientist—as if the act of describing the model were of secondary importance.

One reason why we are predisposed to thinking of models as abstractions is that many models appear as mathematical representations. But actually, mathematical models are often complicatedly constructed, even though they do not necessarily seem so. Take the business-cycle models Boumans studies, for example. Boumans calls these models “first-generation” mathematical models, and the mathematical representations that were finally arrived at do look neat, condensed, and even relatively simple from the mathematical point of view. These formalisms, in and of themselves, cannot tell how much work was involved in reaching them, how many individual ingredients they were built of, and the numerous translations from one representative “language” to another that were required—not to mention the tradition and accumulated knowledge inherent in such processes. Thus, scientific models are by their construction linked in complex and intricate ways to many kinds of knowledge originating from different fields and materialized in the machinery and modeling procedures. Because of that, models and simulations are not justified merely by what they produce; rather, part of their justification is “built-in” or internal to them (Boumans 1999, Winsberg 1999).

It seems that the ability of models to mediate is based on their materiality, but not in any straightforward way. In the absence of any extrasensory mediation, any mediation or communication between humans has to happen through sensory sign-vehicles of some kind. Yet, as mediators, epistemic artifacts extend a link to other forms of knowledge and artifacts that have made their compilation work possible, and they

mediate information only to those who have some relevant background knowledge.

Comprehending models as artifacts helps us also see that models provide an *artifactual* sphere of research of their own. One can, for instance, do experimental work with models. This applies especially to mathematical models and their computer simulations. In her study on simulation, Dowling (1999) found that the scientists used simulation also as “virtual laboratory” by black-boxing the internal structure of the program in order to interact with the computer in a more “experimental manner.”

Finally, it seems to us that the epistemic potentiality of models is closely related to the characteristics described above that make them constrained yet open. Since a model is a purposefully fabricated, constrained structure and, consequently, we know its construction, we expect it to help us to proceed rigorously in solving our problem. Or we try out different things with it, in an orderly fashion. Once working with it, we also expect it to astonish us, to produce something unexpected. For instance, McMullin claims that “a good model has a *surplus content* which enables the theory based on it to survive challenge and extend in all sorts of unexpected ways” (1968, 395). He explains this surplus content by claiming that it shows that the “model-structure has some sort of basis in the ‘real world’.” Indeed, but this basis is rather the reality of the model itself. As a thing, a model has an existence of its own. For this reason we cannot be totally in charge of it, however purposefully fabricated it may be. And yet, we can always try to move it into a new context and ask it new questions.

**4. A Parser as an Epistemic Artifact.** A parser is a language-technological artifact that assigns morphological and syntactic markup to written input texts and in this way provides a partial interpretation of the text. It is embedded technology used in many language-technological applications. To call a parser a model seems problematic: It is rather a program, an instrument, or a thing, than a model. If it is a model, then what is a model of? What does it represent? Using the parser as our practical example, we attempt to show how treating models as epistemic artifacts discloses the affinity of the parser to various other things scientists call models. In addition, we argue that the relation between modeling and representation is not as straightforward as it has been often thought to be. The representational vagueness, or rather openness, of a model does not prevent it from playing several important epistemic roles in scientific endeavor.

Linguists model language (structures consisting of sounds, words and sentences) with grammars. Pedagogically or scientifically oriented, a grammar is ideally comprehensive (all structures of the object language are described), constrained (only the correct structures are described), and

consistent (no contradictions). The scientific evaluation of grammars poses the following problem: How can we determine the relation between a grammar and its object? The traditional way is to subject the grammar to a readership and hope for a conventional academic discussion. However, humans are too unsystematic to interpret or apply complex rule systems consistently. Although people can be taught to apply very simple systems, such as multiplication tables, most likely their interpretations of a complex system such as an extensive grammar are based on too many uncontrolled sources for predictable results.<sup>2</sup> Maybe for this reason, grammars are often viewed as rather subjective things of little scientific interest.

Thus, to evaluate a grammar we need an interpreter with well-known knowledge sources and well-known operation: presently, a computer. Unfortunately, such interpreters cannot, for the moment, interpret traditional grammars as such; here grammars should be written in a more constrained manner for automatic interpretation. A natural language parser is a well-defined interpreter whose inputs are (i) natural language sentences and (ii) a formal language model (e.g., syntax), and whose output is sentences with a grammatical analysis. An extensive language model can be developed, and its relation to sizable instances of the object language can be mechanically evaluated in terms of correctness.

Below we present a case of Finite-State Constraint Grammar (FSCG) that draws on basic research on language engineering, carried out at the University of Helsinki, and continued in some language technology companies.<sup>3</sup>

FSCG is a mathematically well-understood method of assigning a grammatical analysis to sentences. A FSCG compiler/interpreter takes two inputs:

(1) sentences enriched with possible analyses listed as alternatives. An example (@@ = sentence boundary, @ = word boundary, @/ = boundary between two clauses, @< and @> delimit a center-embedded clause):

((@@ the Article  
 (OR: @ @/ @< @>)  
 man (OR: Noun Verb)  
 (OR: @ @/ @< @>)

2. However, what linguists can do rather consistently is analyze linguistic utterances when provided with a set of grammatical descriptors (e.g., word class symbols: noun, verb etc; functional categories: subject, object etc.) and guidelines for applying them to text (for a double-blind experiment, see Voutilainen 1999). We will return to this below.

3. This work is extensively documented in Karlsson et al. 1995 and three chapters by Koskenniemi, Tapanainen and Voutilainen in Roche and Schabes 1997, as well as in proceedings of numerous meetings of the ACL (Association for Computational Linguistics) since the early 1990s.

walks (OR: Noun Verb)

(OR: @ @/ @< @>)

away Adverb

(OR: @ @/ @< @>)

. @(@)

Alternative analyses are here listed with the “OR” operator. A sentence reading is a traversal from one sentence delimiter to another, for instance, the following correct one (from a total of 1024):

@@ the Article @

man Noun @

walks Verb @

away Adverb @

. @(@)

(2) formal grammar rules (constraints) compilable into finite-state automata. Two very simple constraints about legitimate part-of-speech sequences look like this:<sup>4</sup>

# “An article is followed by a noun or an adjective.”

Article → \_ . @ . [Noun|Adjective] ;

# “A modal auxiliary, e.g. ‘shall,’ is followed by an infinitive.”

ModalAuxiliary → \_ .. InfinitiveVerb ;

Parsing means intersecting all grammar automata with the sentence automaton. Sentence readings accepted by all grammar automata are proposed as analyses of the sentence. Multiple analyses usually means that the grammar is still underspecific and needs a more constrained reformulation. No analyses means that the grammar does not accept any of the available analyses; consequently the sentence is ungrammatical, at least with regard to the formal grammar.

The writing and testing of a comprehensive grammar is a process with six phases:

1. When writing a large-scale formal grammar, the linguist benefits from concrete and ready-to-use specifications and resources: descriptive symbols, their application principles (“manual”), and utterances representing the object language (e.g., an inventory of sentences representing different grammatical phenomena). These three can be combined into one resource: a manually annotated model corpus. To prevent misannotations in the corpus and to locate cases where the correct analysis is debatable, the annotation should be done, for the

4. More abstract structures have also been described in FSCG, e.g., the functional dependency structure of clauses and their interrelations.

- most part independently, by two or more linguists, using the double-blind method (see Voutilainen 1999).
2. Development and use of the lookup module, a system that enriches words and sentences with possible alternative analyses (as shown above).
  3. With the lookup module, utterances are enriched with possible linguistic analyses (= ambiguity lookup) and then disambiguated with the compiled finite-state grammar automata written so far. When utterances receive multiple analyses, the grammarian examines these analyses and usually realizes what kind of modifications should be made to the grammar (e.g., the addition of new constraints) to resolve more ambiguity.
  4. As new versions of the grammar emerge, they are tested against the development part of the model corpus. Here the compiler identifies cases where the parser's analysis disagrees with the benchmark analysis and reports relevant conflicts to the developer for diagnosis and correction.
  5. The test-operate-test cycle is continued until a sufficient accuracy has been reached.
  6. Finally, an evaluation is carried out against a held-out evaluation corpus.

The relation between a language and its model can be studied as the relation between the language instance enriched by linguists and the language instance enriched by the parser. After a successful modeling cycle, we have a well-defined interpreter that, using solely the formal language model as its knowledge base, produces analyses almost similar to analyses produced by expert linguists under optimal circumstances. A language model that may pass this kind of Turing test should be an interesting starting-point for further study. One can study, for instance, what properties distinguish the successful model from less successful models, what assumptions can be made about the "human language faculty," and so on.

In analyzing a parser as an epistemic artifact we should first note that the primary reason for constructing a parser is what it produces, a parsed (or partially interpreted) text that is needed for different language-technological applications. However, as soon as we start building a parser it becomes an interesting research problem in its own right. What kind of knowledge, know-how, and machinery is required to build such an artifact as a parser? The answers to these and many other questions become known in the process of fabricating the parser, and part of the resulting knowledge and know-how can be transferred to other areas of research as well. In addition, a substantial amount of knowledge about language can be obtained as a result of fabricating a parser and trying to get it work—for in-

stance, knowledge about the regularity of different (object) languages and even uncovering some new linguistic “facts.” The development of language technology has also made us aware how ambiguous our natural languages actually are.

Much of the epistemic value of a parser follows from its instrumental success. It can be used as a tool or module in investigating and making other language-technological appliances. On the other hand, language technology has given new tools and subsequently new possibilities for traditional linguistic research, since it has made it possible to investigate large linguistic data corpora. Moreover, it is also a tool for its own construction. For example, annotated corpora are used in parser fabrication, but parsers, for their part, are used to make those corpora. Finally, only when a parser functions well, when it produces reliably what we expect from it, can it become an interesting object in our effort to understand language and cognition. Most of the evidence and insight it gives about language and cognition is indirect, however, providing links to other bodies of knowledge.

In conclusion, as an epistemic artifact, a parser plays the roles of both *tool* and *object* of language-technological and linguistic research. It has both scientific and mundane instrumental uses in which it is used largely as a tool—as an embedded technology. As an object of research it is open to different interpretations and connections depending on which questions we want to ask of it.<sup>5</sup> These interpretations and connections are afforded by it because of its “nature” as a locus where different kinds of knowledge and instrumentality meet.

**5. Models and Representation.** Representation is an issue that constantly comes up in connection with models. Writers supporting different views on models have been relatively unanimous in agreeing that to be a model is to involve a representation of some kind. Interestingly, we use the concepts of model and representation in a similar fashion. We talk about models and representations when referring to material entities, processes, or (sign) vehicles, all of which can act as representatives of something else. But for a model to be a model of something else—that is, to represent something else than itself—a relationship is required. In this sense, a model—or representation—is a relation. One important reason to approach models as epistemic artifacts is to try and express this multifacetedness of what being a model encompasses. On the one hand, models are independent things that can be considered as having a certain constitution. On the other, since they are artifacts constructed and used in scientific activity, they are always

5. For the distinction between tool and object, see Engeström and Escalante (1995) and Miettinen (1998).

involved in epistemic relations of some kinds—and new relations are bound to accrue in continuous scientific use. It seems to us that this is what Suarez (1999) is after when he claims that scientific models are not merely objects (or structures) but heterogeneous combinations of mathematical or physical objects (or structures) and intended uses.

However, we would like to emphasize that although models bear traces of their intended use in their construction, they can be also used in many other ways. As epistemic artifacts, models are open-ended things that have their own history and dwell in our research practices in manifold ways as both tools and objects of inquiry. Often these different forms of existence of the models are closely related, as shown by the case of the parser, whose instrumental fitness is closely related to its epistemic value. Moreover, before we have built and worked with the model, there is often no way of knowing whether it represents something and how. To conceive of representation as afforded by models as a static relation between two structures, the real system and its representation, is to approach science from the point of view of “finished” science. Nonetheless, most science is preliminary, as Hartmann (1995) says. Treating models as epistemic artifacts challenges us to think about representation from a new angle. Models are often constructed by representing some aspects of some real systems or their functioning, but most scientific insight they give us is indirect, a result of working with a model in specific contexts.

The tendency of philosophers to treat models as abstract, theoretical and ready-made representative entities leaves much scientific work unrecognized and its epistemic value unexplained, and makes the work of experimentalists and applied scientists seem atheoretical (see Fox Keller 2000). We urge philosophers to investigate what is happening in the so-called applied sciences. They might find out that applied science is not just theory applied, but something more subtle and epistemologically interesting. And in that work models as epistemic artifacts play a pivotal role.

#### REFERENCES

- Bailer-Jones, Daniela (1999), “Tracing the Development of Models in the Philosophy of Science”, in Lorenzo Magnani, Nancy Nersessian, and Paul Thagard (eds.), *Model-Based Reasoning in Scientific Discovery*. New York: Kluwer, 23–40.
- Black, Max (1962), *Models and Metaphors*. Ithaca, NY: Cornell University Press.
- Boumans, Marcel (1999), “Built-in Justification”, in Mary Morgan and Margaret Morrison (eds.), *Models as Mediators*. Cambridge: Cambridge University Press, 66–96.
- Da Costa, Newton, and Steven French (2000), “Models, Theories and Structures: Thirty Years On”, *Philosophy of Science* 67 (Proceedings): S116–S127.
- Dowling, Deborah (1999), “Experimenting on Theories”, *Science in Context* 12(2): 261–273.
- Engeström, Yrjö, and Victoria Escalante (1995), “Mundane Tool or Object of Affection? The Rise and Fall of Postal Buddy”, in Bonnie Nardi (ed.), *Context and Consciousness: Activity Theory and Human-Computer Interaction*. Cambridge: MIT Press, 265–290.

- Fox Keller, Evelyn (2000), "Models of and Models for: Theory and Practice in Contemporary Biology", *Philosophy of Science* 67 (Proceedings): S72–S86.
- French, Steven, and James Ladyman (1999), "Reinflating the Semantic Approach", *International Studies in the Philosophy of Science*, 13(2): 103–121.
- Giere, Ronald (1988), *Explaining Science: A Cognitive Approach*. Chicago and London: University of Chicago Press.
- Hartmann, Stephan (1995), "Models as a Tool for Theory Construction: Some Strategies of Preliminary Physics", in William E. Herfel et al. (eds.), *Theories and Models in Scientific Processes*. Amsterdam: Rodolpi, 49–67.
- Karlsson, Fred, Atro Voutilainen, Juha Heikkilä, and Arto Anttila (eds.) (1995), *Constraint Grammar: A Language-Independent System for Parsing Unrestricted Text*. Berlin and New York: Mouton de Gruyter.
- Knorr Cetina, Karin (1999), *Epistemic Cultures: How the Sciences Make Knowledge*. Cambridge: Harvard University Press.
- McMullin, Ernan (1968), "What Do Physical Models Tell Us?", in Bob van Rootselaar and J. F. Staal (eds.), *Logic, Methodology, and Philosophy of Science III*. Amsterdam: North-Holland, 385–396.
- Miettinen, Reijo (1998), "Object Construction and Networks in Research Work: The Case of Research on Cellulose-Degrading Enzymes", *Social Studies of Science* 28(3): 423–463.
- Morrison, Margaret, and Mary Morgan (1999a), "Introduction", in Mary Morgan and Margaret Morrison (eds.), *Models as Mediators*. Cambridge: Cambridge University Press, 1–9.
- (1999b), "Models as Mediating Instruments", in Mary Morgan and Margaret Morrison (eds.), *Models as Mediators*. Cambridge: Cambridge University Press, 10–37.
- Rheinberger, Hans-Jörg (1997), *Toward a History of Epistemic Things*. Stanford: Stanford University Press.
- Roche, Emanuel, and Yves Schabes (eds.) (1997), *Finite-State Language Processing*. Cambridge: MIT Press.
- Suarez, Maurizio (1999), "Theories, Models, and Representations", in Lorenzo Magnani, Nancy Nersessian, and Paul Thagard (eds.), *Model-Based Reasoning in Scientific Discovery*. New York: Kluwer, 75–83.
- Voutilainen, Atro (1999), "An Experiment on the Upper Bound of Interjudge Agreement: The Case of Tagging", *Proceedings of the Ninth Conference of the European Chapter of the Association for Computational Linguistics, 8–12 June 1999*, Bergen, Norway. Association for Computational Linguistics.
- Wartofsky, Marx (1979), *Models: Representation and the Scientific Understanding*. Dordrecht: Reidel.
- Winsberg, Eric (1999), "Sanctioning Models: The Epistemology of Simulation", *Science in Context* 12(2): 275–292.