

Can Agent-Based Models Assist Decisions on Large-Scale Practical Problems? A Philosophical Analysis

Use of Microstructurally Complex Models

The use of predictive agent-based models as decision assisting tools in practical problems has been proposed. This article aims at a theoretical clarification of the conditions for such use under what has been called post-normal problems, characterized by high stakes, high and possibly irreducible uncertainties, and high systemic complexity. Our argument suggests that model validation is often impossible under post-normal conditions; however, predictive models can still be useful as learning devices (heuristic purposes, formal *Gedanken* experiments). In this case, micro-structurally complex models are to be preferred to micro-structurally simple ones; this is illustrated by means of two examples.

DOMINIQUE GROSS AND ROGER STRAND

Dominique Gross and Roger Strand are researchers at the Centre for the Study of the Sciences and the Humanities, University of Bergen, Norway. One of Gross' research interests is the methodology of complex systems science. Strand is currently working on philosophical and policy-oriented aspects of biotechnology and biological complexity. E-mail: dominique.gross@sut.uib.no roger.strand@sut.uib.no

INTRODUCTION

Society faces today new large-scale social and ecological practical problems of human civilization. It has been argued that these problems involve system uncertainties and a richness of causal connection to a degree that impairs top-down modeling or in general, traditional "reductionist" science (analyzing problems by their parts). With reference to the philosopher Thomas Kuhn's concept of normal science [1], such problems have been called "postnormal," defined as urgent practical problems with high stakes and large and possibly irreducible uncertainties and complexities involved [2,3]. There have been numerous efforts to improve on this state of methodological inadequacy, ranging from statistical technique to new forms of interdisciplinary research (notably what is sometimes called the environmental sciences) [4–6].

There is currently some enthusiasm that the growing science of complexity can help to manage postnormal problems. Thus, Casti [7] refers to risks involved in nuclear power production, AIDS research, the global climate, the global economy, and genetic engineering, and writes [7, p. 35]:

For the first time in history we are in a position to do bona fide laboratory experiments on these kind of complex systems. [...] But now, thanks to the availability of affordable, high-quality computing abilities, we can actually construct silicon surrogates for these complex, real-world processes. We can use these surrogates as laboratories for carrying out the experiments needed to be able to construct viable theories of complex physical, social, biological, and behavioral processes.

Casti calls such models *would-be worlds*. He especially refers to agent-based modeling, that is, the set of techniques that above all are known through the work of the Santa Fe Institute. (Sometimes this modeling technique is also called “bottom-up modeling” or “individual-based modeling” [8].) In agent-based models, relations and descriptions of global variables of traditional (top-down) models are replaced by an explicit representation of the microscopic features of the system, typically in the form of microscopic entities (“agents”) that interact with each other and their environment according to (often very simple) rules in a discrete space-time. Agent-based modeling is becoming known to a broader public both within and outside science, and its explanatory successes have led to expectations that such models might be applied in a more direct sense, notably in industrial, economical, and political decision processes [7,9,10].

It has been demonstrated that a number of interesting qualitative features of rather complex natural systems can be reproduced and explained by means of relatively simple agent-based models. The complexity of these systems can thus be thought of as emerging from simple rules. One example is Per Bak’s sand-pile model [11–13], which, with its almost maximally simple microstructure, exhibits phenomena such as the $1/f$ -noise and punctuated

equilibrium [14], which are supposed to be features of natural systems such as biological evolution [15], earthquakes, or wood fires.

Turning to the history of science, we are reminded of Galilei: Once the right level of abstraction is found, models even of complex systems can be both simple and complete at the same time. And in fact it has been demonstrated that many patterns and structures observed in natural systems can be thought of as the result of rather simple generating mechanisms. Examples include the so-called L-Systems, allowing the generation of very realistic pictures of plants by following some simple algorithms [16], and the study of artificial life, where it has been demonstrated that very effective collective behavior of, for example, termites can be generated by strikingly simple rules for individual behavior [17].

In our opinion, the science of complexity has a large potential for the future advancement of knowledge. However, considering the urgency and importance of problems of the postnormal type, claims like Casti’s should be analyzed with care. If “would-be worlds” can be built and successfully used, they will be of immense practical significance, and the construction of “silicon surrogates” of the climate, the world market and the like, should be given top priority. On the other hand, if unvalidated models are used in decision-making processes, the outcome may be fatal, especially in decisions regarding irreversible problems. Global climate may be a paradigm case. Accordingly, there is a need to analyze under which conditions such models can be considered valid.

In this article we examine how computer-based models in general and agent-based models of a manageable size in particular might function as tools to assist decisions under postnormal conditions. Our main question will be whether these postnormal problems can and should be attacked by means of

simple models. When are we really “in a position to do bona fide laboratory experiments on these kind of complex systems?” Do, under postnormal circumstances, agent-based models provide the possibility to perform experiments on the digital “surrogate” to test the consequences of decisions before actually implementing them?

We first review various usages of models, before showing that the validation of an agent-based model with the purpose of prediction poses certain challenges that in many cases hardly can be met. In particular, we discuss models of slow, long-term processes in global systems (such as the climate) and human impact on such processes. We conclude with unpromising prospects of validating *predictive* models of such systems before the relevant practical decisions are to be made. Instead, one will probably have to be content with the clarifying, explanatory, and heuristic qualities of such models. In this regard the methodological norm of simplicity is seen to be ambiguous. Although it can be wise and even necessary to limit the level of detail of a model, we argue that simplistic features as linearization, stratification of causal relations, and overstated determinacy prohibit learning of certain complex systems. Accordingly, we encourage the construction of models with a rich microstructure.

METHOD AND TERMINOLOGY

The method of this article is that of philosophical analysis, performed in natural language. It does not review or describe the range of existing models and their actual state of validation, nor does it make generalizations based on examples of such models. Instead, it addresses normative questions such as “what is required for a proper model validation?” There is, of course, no one general answer to that question to be found in philosophy (or in science). Rather, individual modeling practices should be discussed scientifically in their particular contexts. However, ad-

equate concepts and figures of thought are required for scientific discourse, notwithstanding the responsible use of scientific results in matters of policy. The aim of this article is to provide effective concepts and lines of reasoning. Accordingly, our contribution should not be seen as an attempt of “proving” a position but as a preliminary result of a priori deliberations on the uses and scope of agent-based modeling.

A few key terms need precise definition because their use is rather diverse in the literature:

A *natural system* is a real system in the world, not excluding social or cultural phenomena.

A *model* is a formal, theoretical, and/or physical system intended to bear specified similarities with a given natural system.

Macro- signifies the phenomena of primary interest in the natural system and the model, the behavior of which is to be predicted, explained, or otherwise studied. Collectively these phenomena constitute the *macrolevel* of the system, depicted through *macrovariables* of the model.

Qualitative behavior of system and model is the flow of phenomena on the macrolevel or the state of the macrovariables through time, respectively. The qualitative behavior of a system or model may display regularities that can be described in terms of *qualitative features* such as statistical distributions, fractal dimensions, periodicity, dynamical regimes, and similar things.

Micro- signifies every part of the system and every specification of the model that is not macro. The microstructure of a model is the set of all such specifications. *Microvariables* and the *parameters* are a part of it, where the latter specify the conditions for the behavior of, or the interactions between, the microvariables of the model. If human intervention in the system can be encoded into a change of a set of parameters, this set is called *free parameters*. The microstructure also includes constraints in terms of the over-all design of the model, that is, what kind of microevents and entities (agents) that are programmed into the code and thus

can appear in the model. We call the set of these constraints the *microstructural frame*.

Variables is taken to be a primitive concept. In the case of agent-based models, they refer to descriptions of individual or collective states of agents and the environment as a function of time.

VARIOUS USES OF MODELS

When discussing the use of a model, it may be convenient to distinguish between the various usages of models, for instance by applying the following typology: (1) *predictive models*, aiming at a (quantitative or qualitative) prediction of future states of a specific, real system; (2) *explanatory models*, aiming at the elucidation of “essential” mechanisms, typically of a class of systems at a more general and/or idealized level; and (3) *heuristic models*, aiming more generally at the invention and discovery of unknown properties of some real or formal system through a learning process involving the manipulation of (“playing with”) the model. It is easily seen that this typology neither is exhaustive nor mutually exclusive. For instance, what sometimes are called “prescriptive models” [7] can, according to the modelers’ ambitions, be seen as predictive with free parameters, or as heuristic models, displaying formal consequences of choices of structure in a hypothetical system (formal *Gedanken experiments*).

An important insight from contemporary philosophy of science [18–21] is that the relations between a model of a given type and its practical use are diverse. Validation in the classical philosophical sense of logical empiricism would require verification of a one-to-one correspondence between the elements of the model and the selected set of observable entities of the natural system [22,23], but today many scientists and philosophers would argue that a model also can be very useful when such a strict procedure cannot be completed. First, the use of an explanatory or heuristic model may help to clarify assumptions and inferences in the decision-making process. Second, under

some circumstances it may be rational to take a chance and believe in a particular predictive model, even if it cannot be properly validated. Indeed, Popper’s philosophy of science [24] has been seen to show that complete verification is impossible and that we can approach the truth only negatively, by falsification and elimination of error. Finally, one can make the model first and then construct or change a natural system so that it fits the model and allows scientific or technological control.

Such ways of relating model to decision will happen, however, at the expense of the ideal of science as merely making neutral representations of a given system. Indeed, they are also easily interpreted as interventions in the natural and/or cultural world. This is shown in an illustrative manner in Casti [7], where he argues that perfect fidelity of a model is neither sufficient nor necessary for its successful use. The author invokes as a metaphor Picasso’s portrait of Gertrude Stein and writes [7, p. 22] :

Picasso replied: “Everybody thinks she is not at all like her portrait, but never mind. In the end she will manage to look just like it.” In fact, in later years the portrait was indeed acclaimed as being an admirable likeness of the writer. If we were to think of Picasso’s portrait as a model of Gertrude Stein, then what was seen as the reality in some sense was the model. [. . .] [T]he theory [read: model] tells you what you can observe.

Taking this metaphor seriously, the story can be interpreted in three plausible ways: (1) Picasso managed to predict Stein’s later appearance by grasping some essential principles of her personality. (2) Stein, living with the knowledge of this famous portrait, was the victim of self-fulfilling prophecy. (3) The existence of the portrait conditioned the audience’s perception of reality. The first interpretation is a case of making a neutral representation; the

latter two are also seen as acts that intervened in the life of others.

Accordingly, models may be successfully used not only because they can be seen to represent the natural system correctly but also because of luck; by changing the terrain according to the map; because of self-fulfilling prophecies (when modeling social phenomena); and so forth. In some contexts, such strategies may be perfectly rational because the model may be validated and improved along the way through trial and error. In the postnormal case, however, it may be that such correction amounts to detection of irreversible damage, for instance, when a species is extinct, the global temperature has risen or the accident has happened. This implies a need for *ante hoc* validation. In the next sections, we will analyze the possibilities of validation in such cases, also taking into account the possible characteristics of the natural systems encountered in postnormal problems.

VALIDATION OF AGENT-BASED MODELS OF COMPLEX SYSTEMS

Methodological Problems of Validation Procedures

Although there are many different ways of using a model, even a predictive one (see Various Uses of Models), the approaches to *validation* are fewer. In general, a predictive model can be scientifically validated by comparing its predictions with future or past observations or by verifying structural similarity between the model and the present empirical and/or theoretical knowledge of the system. As everybody would agree, the optimal case is to be able to combine these approaches. We now look into the methodological aspects of each of the three approaches in the context of models in the management of postnormal problems.

Validation through prediction:

When the scope and accuracy of the predictions of a model have proven satisfactory in repeated testing events, it is rational (though not infallible) to expect the model to stay trustworthy under similar conditions in the future. Under normal conditions, such procedures are

theoretically unproblematic. However, many important systems and problems, especially in postnormal situations, do not allow repeated testing events. The remainder of this article is devoted to these kinds of systems/problems because they are abundant in the realms of policy. Many natural and societal systems are too valuable to allow experimental intervention on a realistic scale for the sake of acquiring knowledge. Moreover, systems such as forests or climate indicators typically change too slowly with human intervention to allow a normal validation procedure with tests of predictions on a relevant time-scale. On the other hand, decisions based on false predictions of the model can have severe consequences. Basically, the validity of the model has to be clear before application, and validation through prediction is then of less relevance.

Validation through retrodiction:

The reasoning behind retrodiction, or tests against historical data, goes like this: Given the existence of a historical record of sufficient quality and a model that reproduces the record correctly (i.e., it gives the correct retrodictions for a part of the history), the model may be trusted also for the future. The reproduction of the historical record thus plays a role analogous to real-time prediction.

Typically, there will be a huge number of possible models that yield, within some error margin, a correct retrodiction of the historical record. These models can be very different from one another in terms of microstructure, from the very simple to the very complicated and from “realistic” representations to highly idealized abstractions. The modeling process can thus be thought of as the selection among these possible models. The selection will be seen to depend on methodological criteria, but also on the purpose of the model, the particular experience of the modelers, and other factors. Probably, this process neither can nor should be seen as a matter of following a set of fixed rules [25], but rather as a creative, open-ended process [26]. On the other hand, modelers do have a need for

regulative rules-of-thumb along the way to closure of the modeling process, and a much-cited rule is the norm of simplicity, often called Occam’s razor [27].

Occam’s razor is probably indispensable in the sense that simplification and idealization are necessary parts of making workable models in finite time. In light of the success of Per Bak’s sandpiles or the simple models of social insects, it might at first sight seem that Occam’s razor also helps solve the problem of selection among retrodictively correct models. One might be encouraged to search for the simplest model that makes correct retrodictions and at the same time satisfies standards of plausibility (being somewhat “realistic” and compatible with existing models and theories).

The rationality of this procedure depends on strong assumptions about the modeled system and the purpose of the model, namely that the available information in the historical record entails the future behavior of the system (at least inasmuch as the modeler is interested in it). This is a plausible assumption when dealing with ant colonies or, say, the typical systems in classical physics (and here we may invoke Rosen’s [28] notion of *mechanisms*). The assumption is not justified when dealing with systems (1) of which our historical record is limited and uncertain, (2) that we suspect to involve internal feedbacks and nonlinear interactions to a degree where even small changes of parameters may be important, and (3) on which there is evidence of recent large-scale human intervention.

Postnormal problems often involve the latter kind of systems. Indeed, the issue at stake typically is to evaluate the evidence for and against change of qualitative behavior of natural systems caused by recent human activities such as large-scale greenhouse emissions, industrial fisheries, and so on. For example, to ask if the Gulf Stream will disappear from Northern Europe is to ask if the system leaves the range of our historical record and thus the valid range of any model validated by retrodiction alone. There may, of course, in any case

exist retrodictively correct models that happen to predict the future correctly. The point is that there is no a priori reason to believe that Occam's razor (or similar rules) is a rational strategy toward the problem of selection among the model candidates. Accordingly, there are small prospects of a proper validation through retrodiction of models in postnormal problems.

Structural similarity: Validation through retrodiction and Occam's razor are rarely practiced alone. In most modeling practices, there is some explicit or implicit element of realism in the sense of striving for structural similarity between the model and the system as we know it, making it "plausible" or "credible." (For an exception, see Kvalheim [29] and Wold [30].) To *validate* a model on the grounds of structural similarity is, however, unrealistic in the case of complex systems because such systems typically include huge numbers and massive heterogeneity of parts, localities, forces, and causal relations. Trying to build a model with the purpose of true and valid representation of the system's microstructure may in many cases be likened to try to get to the moon by climbing a higher tree. In sum, there can probably be no proper validation of models in the management of postnormal problems.

Pseudo-validation: In our discussions with scientists, some have argued that our concept of validation is too rigid. First, it is argued, the lack of proper validation of individual models may be somewhat helped by collective validation. If a great number of models all yield the same (qualitative) prediction, it seems that the prediction was somewhat justified. The problem remains, however, to decide if the sameness could be due to bias in input data or to some unnoticed similarity of model design. Second, stressing the importance of skills and expertise in modeling, one may trust modelers known for successful models that allowed proper validation. In our view, this will not change the fact that postnormal problems differ from the normal situation exactly in the prospects of modeling success. Finally, it is sometimes ar-

gued that policy decisions anyhow have to be based on our present beliefs about the future, and therefore it is rational to replace unscientific beliefs with scientific predictions, even if they are unvalidated. However, this argument is flawed. It is rational to use all available sources of knowledge, but there are more ways to use a model than accepting its predictions. In the remainder of the article we discuss the rational use of unvalidated models.

Unvalidated Models

Assuming that postnormal problems frequently entail questions like "Will the system change its qualitative behavior?", we have seen why models to this purpose cannot be validated and why predictions from unvalidated models should not be trusted. Unvalidated models may be useful, though, to facilitate learning processes. First, we may discover new scenarios. Second, the model may have explanatory value in terms of the exploration of the implication of one's theoretical beliefs (formal *Gedanken experiments*). Other use is problematic (see Various Uses of Models).

Several well-known, agent-based models, although not postnormal, also are either unvalidated or not even validatable. One might think of Tierra [31], an attempt to model biological evolution by means of self-reproducing, machine-code programs that compete for CPU-time and memory. In a strong sense Tierra cannot be validated, but nevertheless it provides important insights into the nature of evolution. Similar ideas apply for many other of the best-known agent-based models.

Because under postnormal conditions most models are unvalidated anyway, one might, in the spirit of Casti's quote in the Introduction, apply the ideas of agent-based modeling and build simple models of the respective systems to assist our decision-making process. Hence, a relatively primitive microstructure (not necessarily agent-based, but explicit in the formulation of the microstructure) should suffice to come to, if not good predictions, at least some understanding of the system. If the models are kept simple, it would

even be possible to build and compare a relatively great number of them, which in turn would lead to more knowledge about the system and a better feeling for its intricacies.

However, we do not think so. Postnormal problems raise different questions and demand other answers than, for example, theoretical problems in evolutionary biology do. In practical decisions of the postnormal type, the impact of a rich microstructure cannot be disregarded as theoretically uninteresting deviations from general principles. The space of possible decisions is of very high dimension, there are nonscientific considerations to be taken into account, and not at least, the inherent nonlinearities of the system make idealizations suspicious. In-principle understanding of the system is of little interest in these situations; what is needed is rather an understanding of the specificity of problem and of its specific dangers. Thus, simple agent-based models in the style of Tierra, which might lead to important theoretical insights in some contexts, can be expected to be of little value under postnormal conditions.

The challenge is thus to design models of high quality for learning processes. It appears that microstructural complexity of models is one quality criterion. However, in order to reach this conclusion, we have to clear an ambiguity in the notion of model simplicity: We propose to distinguish between "simple" models, "simplistic" models, and "detailed" models.

Tierra, the sandpile model, or models of ant societies would then be typical examples of simple models. Although they all display considerable complexity on the macroscopic level, their microstructure typically relies only on a few rules. Thus, "simple" can entail "complex macrostructure." To make Tierra microstructurally complex, one could allow a wider range of dynamics by relaxing constraints imposed on it by its microstructural frame, for instance, the constant, nonadaptive size of the virtual CPU and the mechanism for "birth" and "death" of agents. In a microstructurally

complex system, one cannot rule out interlevel and time-variant interactions. The intrinsic dynamic of the system may change the “rules of the game” on any level: from anaerobic to aerobic life, from marine to terrestrial life, and from vegetative to sentient beings. In the postnormal context, one typically worries about change of qualitative behavior of the system. Accordingly, rule-changing dynamics cannot be excluded if the model is to match the possible flexibility and plasticity of the system.

This notion of microstructural complexity must not be confused with a high level of detail of a model, which we understand as comparable to toy models; instead of being limited to a few characteristics of the system, they would rather be attempts of an exact representation of every part of the system, regardless of its relevance for the problem at hand. A higher level of detail typically leads to higher hardware and computing-time demands. Importantly, too detailed models with many parameters can be less robust from both a practical and principle perspective; they might become brittle, which makes it extremely difficult to gauge the models against the system. In principle, if every “relevant” parameter is to be included, the model will be less robust against low quality in input information [27].

By “simplistic” we denote the lack of sufficient microstructural complexity of an *unvalidated* model in learning processes, with regard to problems involving complex natural systems. It appears that simple agent-based models may have features that can render them simplistic. We conjecture that under postnormal conditions, this will usually be the case. In a certain sense, computer models will always fail to match perfectly the microstructural complexity of the system because such models have to be explicitly specified by some language whereas the real world escapes that requirement. That does not impair the prospects of learning processes through such models. Indeed, the use of unvalidated models would be to learn what given explicit specifications entail in terms of behavior. However, practical interest in or scientific information

about dynamical features of the system may render whole sets of model design as too constrained to serve our learning process. For instance, we cannot learn much from a model that assumes that human beings always behave as rational actors if we are interested in human relations that are scientifically known to have a strong emotional component.

One may argue that scientific knowledge and thus the judgment of a model being simplistic is fallible and that the future will show that insufficient knowledge of the system caused our belief in its irreducible microstructural complexity. A strong historical case in point would be the replacement of the microstructurally complex Ptolemaic world view by the much simpler Copernican model. From a normative point of view, though, it is clear that the transition of belief (and use) from the Ptolemaic to the Copernican model was not rational until the latter was validated through prediction (which took quite some time).

The positive conclusion is thus to encourage models that take into consideration contributions to knowledge from a broad spectrum of the sciences. One crucial aspect is that of quality control of the input knowledge translated into design of microstructure. Any uncertainty in the input to the model will be likely to translate into an uncertainty of the output. In postnormal contexts, where uncertainties may not be reducible in practice or even in principle, it is therefore important to mark this uncertainty rather than hiding it in simplistic models. Another challenge is to learn more about the sources of complexity. In the next section we illustrate the discussion of simplistic models by two examples of natural systems: urban traffic and the climate.

NATURAL SYSTEMS: TWO EXAMPLES

There is no consensus in the literature on how to define or measure complexity [32]. In the previous section we argued for the relevance of a concept of microstructural complexity, understood as a richness, diversity, and dynamic of the parts of the system and their inter-

relations. We now illustrate our discussion by looking at two specific natural systems: the climate and a traffic system.

The following discussion of the global climate closely follows Bengtsson [33]. The problem of predicting the future climate is easily seen to be postnormal: If the suspicion of anthropogenic global warming and its dramatic consequences are justified, large-scale action ought to be taken rapidly. It is unclear how much perturbation the climate can tolerate without undergoing qualitative change. Furthermore, it is unclear what these changes would imply for mankind. Moreover, the question about the development of the global climate has a considerable impact on our lifestyles and habits and is thus closely connected to political and sociological questions, which have to be taken into account in the decision-making process.

One of the biggest problems in climate research is the difficulty of validation of models. Following Bengtsson, the three problems of validation from the section Methodological Problems of Validation Procedures are easily recognized: First, the climate system is slow, unique, and valuable. The decisions cannot await model validation through prediction; by then it would be too late. Second, there is the contingent, but still typical problem of incomplete historical records, decreasing the value of validation through retrodiction. The estimated temperature before 1450 has large error bars, and the observational evidence before the end of the 18th century covers only about 3% of the earth’s surface (Europe and central China).

Third, the issue at stake really is whether a qualitative change of the climate system is encountered or not. Indeed, the historical record indicates “a general ongoing cooling until 1900” [33, p. 4], although it is not clear what caused this trend. Although changes of longer periods seem to be due to solar irradiation changes, “the causes of shorter time scales are still rather mysterious” [33, p. 3]. From the turn of the 20th century there appears to be a global warming, and this development coincides with human emission of green-

house gases. Because these new causes and factors imply that the system is entering a new domain of its parameter space, one is forced into speculation on how the anthropogenic factors interact with other relevant factors and what the net result will be. Following Bengtsson, it seems that a warming process may result in different interplays between the feedbacks, for example, as caused by an increase in vapor and clouds. At present, it is generally believed that increased water vapor results in a positive feedback, but there are deviating results. Even subtler and more uncertain is the question about the cloud feedback.

The change in cloud forcing due to enhanced greenhouse forcing is strongly model dependent, with some models giving a positive feedback, others a negative [33, p. 12]. [...] [Accordingly] climate models have to be realistic and rather detailed, since any systematic model deficiency could create an erroneous response pattern. [...] Simple models could in this context be quite misleading. [33, p. 17]

The model thus has to predict the effects of more clouds in the atmosphere. Now, these effects are likely to vary considerably with the conditions in a way that is not fully known at present. Worse, the suspicion is, as mentioned, that the climate might enter into a new qualitative state where present feedbacks play a different role and maybe new ones have to be taken into account.

The constraints on the qualitative effects that can be observed in the model depend crucially on the microstructural frame. These constraints can be relieved by implementing into the model higher-order features of the system that are thought to have an impact on the macrobehavior. In the present example, one might think of cloud feedback, global currents in the oceans, or the abundance of greenhouse gases in the atmosphere. Furthermore, the microstructural frame constrains the possible interactions between conjectured relevant factors. For example, a model of the effects of global warming may

implement the impact of the melting of the polar caps on the currents (especially the Gulf Stream, see Bernes [34]). In the model, such interactions can be observed only if they are explicitly programmed, that is, their microstructure is accordingly enriched. Thus, a less constraining microstructural frame means higher microstructural complexity. This has to be contrasted with a mere augmentation of the level of detail, which might be achieved by a more realistic representation of topological features. Although it is unclear how much one actually can learn about the climate from models that are microstructurally complex in this sense, it seems clear that decision-relevant learning is not to be achieved when models are kept simple.

Our second example is Casti's [7] description of "Transims," an agent-based model of the road traffic system of the town Albuquerque, New Mexico. Transims represents the traffic flow down to the level of the single car and has a scaled-down internal representation of the road system, allowing accurate study of traffic flow patterns such as traffic jams. The model has also been used with considerable success to test the effects of "perturbations" of the system, such as new traffic light patterns or additional roads or bridges.

The success may at first sight be surprising because a road traffic system might appear just as complex as the climate (even involving human acts and thus intentionality), prohibiting meaningful predictive models. However, this is not so because the system is both fast and easily and accurately observed and allows testing on a nearly daily basis, enabling validation through prediction. The situation is, though, somewhat different when effects of new roads are to be predicted. In this case there is a considerable time delay between the prediction and the testing, and as in the case of the global climate, there is a need for *ante hoc* validation. The success of the model then depends on the amount of required knowledge of the microstructure: Will the behavior of the

agents still conform to the same rules? Interestingly, the microstructural complexity required for the model depends on the nature of the questions to be answered. In traffic planning, one would normally be content with traffic prognoses made under the assumption of normal human behavior. In that case, the individual agents act under so many constraints so as to rule out interesting effects of unbounded intentionality. If one wants to predict events more tightly connected to spontaneous manifestations of free will, such as blind violence or overthrow of government, one would have to try to model the system a lot more carefully with respect to microstructural complexity. For instance, expressive or desperate behavior may play a crucial role, and the emergence of such behavior may be the result of a highly contingent or opaque historical situation. One may think of the Los Angeles riots.

A normal traffic-planning model without the extreme ambitions of forecasting exceptional events such as riots is thus seen to have a moderately complex microstructure. There are the easily programmed traffic rules and the individual routes of the agents, derived from empirical data and presumably quite stable over long periods of time. The description can thus be reduced to a limited number of simple rules plus some statistics. This is different from the global climate, of which there is a lot more to know and understand on the level of microstructure.

CONCLUSIONS

There is a definite need for new scientific approaches to the emerging large-scale practical problems in complex environments, such as the climate. The question we pursued in this article was whether microstructurally simple, agent-based models could be useful in this respect. Because such models have been very successful in the study of the emergence of complexity out of simple rules, it seemed reasonable to expect that this modeling practice might also be applied for the purpose of prediction related to postnormal problems. However, we have argued that in certain contexts

typical of postnormal problems, proper validation of predictive models cannot be had.

We are then left with the possibility of using unvalidated models for explanatory and heuristic use. Even when it is unknown whether they yield correct

prediction, these unvalidated models can be useful if they are microstructurally complex enough to allow a learning process. Accordingly, we suggest that simple agent-based models are only of limited use in the realm of postnormal problems.

ACKNOWLEDGMENTS

We thank two anonymous referees and William Boos, Ragnar Fjelland, Eivind Hiis Hauge, Kjellrun Hiis Hauge, Judith Ann Larsen, Konrad Richter, Tuomo Saloranta, and Rune Vabø for valuable comments and inspiring discussions.

REFERENCES

1. Kuhn, T. *The Structure of Scientific Revolutions*. University of Chicago Press: Chicago, 1962.
2. Funtowicz, S.; Ravetz, J. *Futures* 1993, 25(7), 739.
3. Funtowicz, S.; Ravetz, V. *Environ Toxicol Chem* 1994, 13(12):1881.
4. Martinez-Alier, J.; Schlüpmann, K. *Ecological Economics: Energy, Environment and Society*. Basil Blackwell: London, 1987.
5. 1999. *Futures Special Issue. Post-normal Sciences* 1999, 31(7).
6. Young, P.; Parkinson, S.; Lees, M. *J Appl Stat* 1996, 23, 165.
7. Casti, J. *Would-Be Worlds*. John Wiley & Sons: New York, 1997.8. Judson, O. *Tree* 1994, 9(1), 9.
9. Sardar, Z.; Ravetz, J. *Futures* 1994, 26(6), 563.
10. Arthur, W. *Science* 1999, 248, 107.
11. Bak, P. *How Nature Works*. Oxford University Press, 1997.
12. Bak, P.; Tang, C.; Wiesenfeld, K. *Phys Rev A* 1987, 38, 364.
13. Bak, P.; Tang, C.; Wiesenfeld, K. *Phys Rev Lett* 1987, 59, 381.
14. Bak, P.; Sneppen, K. *Phys Rev Lett* 1993, 71, 4083.
15. Solé, R.; Manrubia, S.; Benton, M.; Bak, P. *Nature*, 1997, 388, 764.
16. Prusinkiewicz, P.; Lindenmayer, L. *The Algorithmic Beauty of Plants*. The Virtual Laboratory. Springer-Verlag: New York, 1990.
17. Bonabeau, E.; Theraulza, G.; Deneubourg, J.; Aron, S.; Camazine, S. *Self-Organization in Social Insects*, April 1997. Santa Fe Institute Working Paper 97-04-032.
18. Harré, R. *The Philosophies of Science*. Oxford University Press: Oxford, 1985.
19. Toulmin, S. *The Philosophy of Science*. Hutchinson: London, 1967.
20. Giere, R. *Understanding Scientific Reasoning*. Holt, Rinehart, and Winston: Chicago, 1984.
21. Funtowicz, S.; Ravetz, J.R. *Uncertainty and Quality in Science for Policy*. Kluwer: Dordrecht, 1990.
22. Nagel, E. *The Structure of Science. Problems in the Logic of Scientific Explanation*. Routledge & Kegan Paul: London, 1961.
23. Hempel, C. *Philosophy of Natural Science*. Prentice-Hall: Engelwood Cliffs, NJ, 1966.
24. Popper, K. *The Logic of Scientific Discovery*. Hutchinson, London, 1959.
25. Feyerabend, P. *Against Method: Outline of an Anarchist Theory of Knowledge*. NLB: London, 1975.
26. Pickering, A. *The Mangle of Practice*. The University of Chicago Press: Chicago, 1995.
27. Brooks, R.; Tobias, A. *Mathematical and Computer Modelling* 1996, 24(4), 1.
28. Rosen, R. *Life Itself*. Columbia University Press: New York, 1991.
29. Kvalheim, O. *Chemometrics and Intelligent Laboratory Systems*, 1992, 14, 1–3.
30. Wold, S.; Dunn, W., III. *J Chem Information Computer Sci* 1983, 23(1), 6.
31. Ray, T. In *The Philosophy of Artificial Life*; Boden, M., ed. 1996, 111–145.
32. Rescher, N. *Complexity—A Philosophical Overview* Transaction 1998.
33. Bengtsson, L. In *Proceedings of the Symposium on Global Biogeochemical Cycles in the Climate System*, Jena, September 22–27, 1999. In press.
34. Bernes, C. *Arktisk miljø i Norden—urørt, overutnyttet, forurenset?* [Arctic Environment in the Nordic Countries] Nordisk Ministerråd, København, 1996.
35. Boden, M., ed. *The Philosophy of Artificial Life*. Oxford Readings in Philosophy. Oxford University Press: Oxford, 1996.