ETHICAL ASPECTS OF COMPUTATIONAL MODELLING IN SCIENCE, DECISION SUPPORT AND COMMUNICATION

Rudolf M. Füchslin[1,2,*], Jacques Ambühl[3], Alessia Faggian[4], Harold M. Fellermann[5], Dandolo Flumini[1], Armando Geller[6], Martin M. Hanczyc[4], Andreas Klinkert[7], Pius Krütli[8], H. G. Matuttis[9], Thomas Ott[10], Stephan Scheidegger[1], Gary Bruno Schmid[11], Timo Smieszek[12], Johannes J. Schneider[1], Albert Steiner[7], Mathias S. Weyland[1]

[1] Institute of Applied Mathematics and Physics, Zurich University of Applied Sciences Zurich, Technikumstrasse 9, Winterthur, 8400, Switzerland.
[2] European Centre for Living Technology, Ca' Bottacin, Dorsoduro 3911 Calle Crosera, 30123 Venice, Italy.
[3]Short Term Consultant. CAHMP Project. World Bank Group. Geneva.
[4]Laboratory for Artificial Biology, Department of Cellular, Computational and Integrative Biology (CIBIO), University of Trento, 38123, Trento, Italy.
[5]School of Computing, Newcastle University, Science Square 1, Newcastle upon Tyne NE4 5TG, United Kingdom.
[6]Scensei (Switzerland) GmbH, Fraumünsterstrasse 11, 8001 Zurich, Switzerland.
[7]Institute of Data Analysis and Process Design, Zurich University of Applied Sciences, Rosenstrasse 3, 8400 Winterthur, Switzerland.
[8]ETH Zurich, Transdisciplinarity Lab, 8092 Zurich, Switzerland.
[9] Department of Mechanical Engineering and Intelligent Systems, The University of Electro-communications, Tokyo 182-8585, Japan.
[10]Institute of Computational Life Sciences, Zurich University of Applied Sciences, Schloss 1, 8820 Wädenswil, Switzerland.
[11]Swiss NeuroCreations GmbH, Hambergersteig 25, 8008 Zurich, Switzerland.
[12]Muesmattweg 70, 4123 Allschwil, Switzerland

[*] Corresponding author e-mail: furu@zhaw.ch

Abstract. The development of data science, the increase of computational power, the availability of the internet infrastructure for data exchange and the urgency for an understanding of complex systems require a responsible and ethical use of computational models in science, communication and decision-making. Starting with a discussion of the width of different purposes of computational models, we first investigate the process of model construction as an interplay of theory and experimentation. We emphasise the different aspects of the tension between model variables and experimentally measurable observables. The resolution of this tension is a prerequisite for the responsible use of models and an instrumental part of using models in the scientific processes. We then discuss the impact of models and the responsibility that results from the fact that models support and may also guide experimentation. Further, we investigate the difference between computational modelling in an interdisciplinary science project and computational models as tools in transdisciplinary decision support.

We regard the communication of model structures and modelling results as essential; however, this communication cannot happen in a technical manner, but model structures and modelling results must be translated into a "narrative." We discuss the role of concepts from disciplines such as literary theory, communication science, and cultural studies and the potential gains that a broader approach can obtain. Considering concepts from the liberal arts, we conclude that there is, besides the responsibility of the model author, also a responsibility of the user/reader of the modelling results.

Keywords: Ethics, Computational Modelling, Transdisciplinarity.

1 Introduction

This article deals with the responsible and ethical use of computational models in science, communication and decision-making. We intend to report experiences we collected over the last couple of years and a focussed discussion during a workshop at the WIVACE 2021 conference, held in Winterthur, Switzerland from Sep 15 to Sep 17, 2021[1]. A systematic treatment is given e.g. in [1] or earlier [2]; for a recent discussion based on case studies see [3], and an overview with a focus on philosophy see [4]. The societal role of artificial life has been a central topic at the ALIFE 2019 conference in Newcastle, UK; various articles on the topic can be found in [5]. Computational modelling still increases its importance in science and, as we have seen during the COVID-19 pandemics, is recognised as a central supportive tool in political decision–making processes. We, as modellers, realised the necessity for embedding and relating our work into a framework that also includes ethical considerations. In this work, we report on our findings, mainly derived not from theoretical work but our daily practice.

This report has been made possible by the joint efforts of different research projects. The leading role is thereby with the EU-funded project ACDC (Artificial Cells with Distributed Cores to Decipher Protein Function, funded by the European Union's Horizon 2020 research and innovation programme, https://acdch2020.eu/). The workshop was organised in an open format but initiated and supported by ACDC. We deliberately have chosen to collect examples from different fields to achieve generality and not restrict ourselves to specific research areas. Together with the open nature of the workshop at WIVACE 2021, this fact justifies the rather long list of co-authors. The first author takes the principal responsibility for the paper, the contributing co – authors are listed in alphabetical order. An extended version of this paper will be the base of a report to the EU as a deliverable for the ACDC project.

This article focuses on computational models and not on theory in general. In our view, computational models are a subclass of general models. A general model may

[1] `https://www.wivace2021.org/`.

define the terms of discourse; a computational model adds quantifiable relations between these terms (or, as we will define it later, uses variables subject to computation). These quantifiable relations are the basis for the implementation of a simulation. A further note on computational models and machine learning: We wrote this report with a somewhat "physics-oriented" interpretation of the term "computational model." Roughly said, we discuss models in which the model variables have an interpretation/ a semantics right from the start and do not acquire this interpretation throughout a training phase. We compare this to machine learning, say, an artificial neural net. The weights attributed to connections of neurons in a neural net have no interpretation before the network's training. Furthermore, even after the training, relating these weights to observables after the training phase is a complicated and only partially understood task. However, we point out that almost everything we write in this article applies to models with variables with direct semantics as well as structures such as artificial neural networks. We do not discuss the issue of the semantics of internal variables further, but we emphasise its importance and point out that it has many layers. For an eloquently written discussion of the "symbol grounding" – problem, i.e. the question of how internal variables are related to objects in the internal world, see [6]. Note, for example, that symbol grounding can become very intricate if one deals with a model that works with variables with a probabilistic interpretation.

Computational models have always been of importance. However, confluent trends of the last couple of years have increased the role of models in science and the relation between science and society. It is a sign of hope and the reason for optimism that recently, powerful youth movements and responsible politicians recommended and even demanded to unite behind the sciences. As scientists, we should welcome this trust. However, we are obligated to reflect on how we can justify this trust and how we have to communicate the results of computational models to avoid misunderstandings and prohibit misuse.

The responsible and ethical use of models is the main topic of this report. We discuss different aspects of the question of responsibility and ethics. Responsible use is closely related to several recurring requests R1-R3:

- R1: A clarification of the differences between computational models as tools in science and computational models that are part of decision support processes that go far beyond the social context of science itself.
- R2: A better understanding of interface processes. Models and simulations are powerful instruments for linking different fields of human expertise and establishing a relation between the real world and different abstractions of it. First, this linking between reality and abstraction requires the construction of interfaces. The presence of interfaces usually implies some form of translation processes, which in general leads to systematic information loss (because a model only represents a part of reality) as well as different types of translation errors (which, for example, can be the result of the limited precision of measurements). Second, models can help to connect different (abstract) universes of discourse (to use the term from computer science) or languages specific to social groups. As scientists, we are used to working within some more or less well-defined area of discourse in which a

common language exists, and the discussion participants share an implicit understanding of boundary conditions, interpretation rules, and the like. As soon as one uses the results of scientific discourse as input for a more general and interdisciplinary decision support process, one must not take this implicit understanding for granted. Scientists have to prepare the output of a computational model in a way that is digestible by actors outside the field in which the model has been set up.

- R3: Not only those developing and using computational models in order to produce output have a responsibility. There is also a responsibility of the reader. By the term "reader," we mean those who take up the result of a computational model but may not have area-specific knowledge to interpret these results as it is standard by those who are within the area of expertise. The scientists have to require the reader to be aware of (probably area-specific) limitations of computational models but also have to explain these limitations in a form that the reader can understand.

In the communication with the public and members from scientific fields not directly connected to computational modelling or natural science, it became clear that (at least) two main issues need to be addressed: the purpose of models and the relation of theory and data science.

The first issue (purpose of models) is a very fundamental one. Many people believe that a model is always a tool for setting up predictive simulations. We summarise this idea of a computational model in Fig. 1. The figure illustrates the model-based control of a robot as a predictive tool for controlling the robot's dynamics. The critical aspect is that the simulation produces a sufficiently faithful analogue of the dynamics in the real world in an appropriately designed mathematical representation. The term "sufficiently" refers to the quality of the control of the robot. Somewhat loosely expressed, sufficient means that the control ensures that the robot reaches a given set of goals. One can draw a similar picture of a simulation for weather forecasts. Also, the simulation result should give a sufficiently accurate prediction of reality in that context.

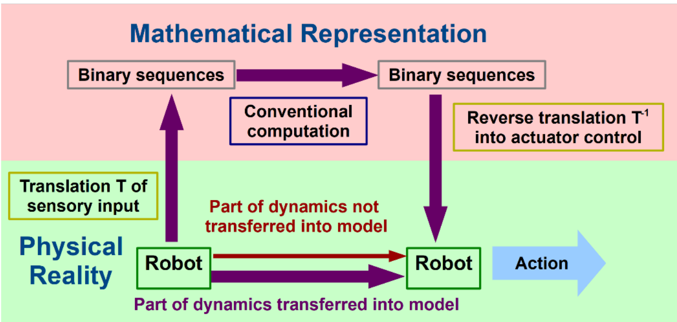

Fig. 1. Computational models as tool for the control of a robot. The computational model serves as a tool for predicting the dynamics of a technical system in a non-trivial environment.

Prediction is an obvious goal of a simulation but certainly not the only one. A non – exhaustive list of (in parts overlapping) purposes for models includes:

- Check the **understanding** of the past: By comparing model output and measured data, one can check whether the model explains what has happened or whether the model is insufficient and must be extended.
- **Optimise** the present: This applies to static or stationary situations with parameters. If one has an optimisation goal, *in silico* experiments can help find optimal values for these parameters.
- **Predict the future** if one knows the present and has a reliable model of the dynamics: Ideally, a model allows one to predict the future, given sufficient knowledge of the initial conditions. However, even if the initial conditions are known, the prediction of the future may still be imprecise, partly because models most often only approximate the actual dynamics, partly because many models contain stochastic components. We point out that even if the latter is not the case, i.e., the dynamics of the model is entirely deterministic, the future may still be difficult to predict, mainly if the dynamics of the model exhibits chaotic components.
- **Estimate the future** if one only guesses the present. Most often, one has only limited knowledge about initial conditions. In that case, either one estimates the initial conditions by statistical means (including the values of the parameters determining the dynamics) or constructs initial conditions based on plausible assumptions.
- **Explain what we see:** A model can help to give meaning to data[2] in the sense that the dynamics or a state can be explained by referring to the model's variables (see below the discussion of Fig. 2). We point out that the purpose of "giving meaning to data" is probably the one in which a restriction of the discussion to computational models is least necessary. Models as tools for explanation do not necessarily need to produce computational results but can serve as platforms for defining concepts and interactions qualitatively. Furthermore, note that an explanation of the behavior of variables (e.g. by showing correlations) differs from the first item in the list, which refers to the question whether the model can reproduce what has been observed.
- **Analyse the dynamics**: it is a fact that human beings are pretty good at understanding equilibrium states but are quite often surprised by the consequences of feedback processes, second-round effects, stochasticity (the consequences of fluctuations), or non-linear dynamics in general. This is an ideal application for models and simulations: with the help of simulations, one can get "a feeling" for the dynamics of a model and explore relevant settings by trying out the consequences of changes in inputs, variations in parameters and the like.
- Detect **emergent dynamics and structures:** Quite often, systems exhibit emergent phenomena. There are many examples in which a model sheds light on the underlying processes by showing that a specific mechanism leads to an observed emergent phenomenon. However, we point out that a model may well give a plausible explanation for certain emergent phenomena; this plausibility must not be confused with evidence.
- **Explain system behaviour from more fundamental dynamics**. A model may explain the behaviour of a specific system from the fundamental properties of its

[2] This has been pointed out by Marcello Pellilo from the University Ca'Foscari, Venice.

constituents. In engineering, a material model (i.e. a model taking material constants and fundamental physical laws as input) may explain a whole family of systems, such as FEM – models in civil engineering or the study of biological matter using molecular dynamics simulations.

- **Decision Support:** Even if, because of lack of data about initial conditions, a model cannot give sufficiently reliable information about the future development of the system under consideration, it may still help to evaluate the range of possibilities via scenario analysis. Such an analysis contains reasonable best and worst cases and a quantitative evaluation of the sensitivity of parameters. Especially studies of parameter sensitivity may help allocate potentially limited resources to determine those parameters that influence and output in a critical manner.
- **Models as platforms for discussion:** Models can serve as platforms for discussions (between modellers and experts or even between experts with domain knowledge with the modellers as moderators) of assumptions, parameter dependencies and qualitative aspects of system behaviour.
- **Virtualisation:** A sufficiently precise model of reality may enable "in silico" experimentation. Besides benefits concerning costs and speed, virtual experiments enable the study of seemingly unphysical situations, e.g., by turning off selected physical mechanisms. Such "knockout experiments" shed light on the actual physical circumstances' importance.
- **Produce data and train modern controllers** (e.g., deep neural nets): various types of artificial intelligence and, more generally, machine learning are now part of the modeler's toolbox. Many of these algorithms require vast amounts of data for training, often much more than experiments can provide. One can train neural nets with simulated data, comparable to the training of pilots in a flight simulator.
- **…**

The second main issue relates to the role of data. In 2008, Chris Anderson wrote an article in Wired titled "The End of Theory: The Data Deluge Makes the Scientific Method Obsolete.", [7]. Although the article's content was much more nuanced, the upcoming of data science and the broad availability of data brought some people to conclude that models and theories are unnecessary. We agree that the availability of cheap sensors, the possibility to transfer data from the sensors to some data processing unit without the need to install complicated hardware, and the ease with which the vast amounts of data can be analysed changes science in a deep sense. However, we still think that models are of relevance. We point out that the amount of data needed to replace a model or theory is often underestimated (particularly in the social sciences). The need for large amounts of data means that even if it were possible in a fundamental sense to dispense with models and only rely on data, it may still not be practical. In addition, one may discuss whether data can replace models in all circumstances. If one used computational approaches only for prediction or scenario building, something like the "Master Algorithm" envisioned, e.g., in [8], could, in principle, do the job. However, other model purposes, such as giving meaning to data, are hard to conceive without a model based on variables with semantics. We also point out that data-driven approaches usually perform poorly in generalisations, at least in the present situation. One may well predict a specific system's behaviour using a data-driven approach, but this ability for prediction is usually challenging to transfer to a different system. In contrast, a system modelling platform based on a fundamental material model relating variables with physical interpretation predicts the behaviour of large classes of structures and systems. The tension between purely data–driven approaches and semantics lies at the heart of the discussion about "explainable AI", see e.g. [9] and with emphasis on medicine, [10].

It goes without saying that in this article, we do not claim to give final answers to the request R1-R3; our experiences over the last couple of years in various disciplines enable us to report some of our findings and some generalisations distilled out of these findings over the last couple of months. Thereby we span modelling of evolutionary processes, statistical physics, modelling of cellular processes, model-based therapy optimisation, traffic simulations, applications of AI in medicine and control of large industrial entities, to give some examples.

We point out the importance of the last two years. All the contributors to this article have worked for decades in various applied and fundamental sciences fields, always using or developing computational models. In addition, some of us are now serving in the decision support for the Swiss government concerning the Covid 19 – pandemics. The change from publishing results in peer-reviewed journals to more direct decision support gave us (sometimes for the first time) the experience of science that has a short-term impact. However, it also made us reflect on the responsibility of developing and using computational models. In addition, it became clear that there is a considerable difference between scientific modelling for science itself and the use of scientific models as tools in a broader, nonscientific (or not exclusively scientific) but rational context.

This report does not focus on the necessity of quality control and standard operating procedures but more on aspects of communication. Also of crucial importance is that scientists being part of decision support processes should reflect on their role. Our experiences led to establishing a network of scientists and decision-makers discussing the role of computational models based on our recent research activities, ranging from purely scientific activities in EU-funded projects to decision-support. Our goal is to study and describe the difficulties of computational modelling in a broader context and in a permanent manner that includes publications and network activities.

The article is organised as follows: Sec. 2 discusses our perspective on models, model building, and implementation. The section defines a couple of terms and presents our view of the process of model building. Sec. 3 discusses the use of models in science, which exhibits relevant differences to the use of models in decision–making. The latter involves a full transdisciplinary mode of communication and is treated in Sec. 4. The interaction of science with stakeholders outside the scientific discourse requires the use of according means of communication (which we call narratives). In Sec. 5, we postulate a responsibility on the narrator's side as well as on the side of the reader. We emphasise the potential for science and science communication to learn from the vast body of literary theory, for the practice of communication, but also its conceptualisation. In the conclusions (Sec. 6), we relate our findings to the current status of modelling in society.

2 The Process Of Modelling

There is a broad discussion of what exactly one understands by the term “model.” This discussion ranges from literary theory over philosophy to the foundations of mathematics and logic (where mathematical structures such as the natural numbers serve as “models” for sets of axioms, e.g., the Peano axioms). In order to reduce misunderstandings that may occur if one speaks about a concept that appears in many different branches of science, we start with defining some basic terminology that we will use in this article:

- A **model** establishes relations between different classes of objects (which we will call fundamental terms). These relations can be static but also consist of rules and descriptions of the dynamics of these objects. Importantly, we require a model to be based on a rational description of the objects and their interactions. Thereby, we understand by the term “rational description” a language-based, sufficiently intersubjective formulation that a sufficiently knowledgeable group of experts can understand.
- A **theory** describes the fundamental terms and interactions of the model in a context that may well go beyond the scope of the model. To give an example here: relativistic quantum field theory is a theory in the sense of a framework, and the standard model of particle physics is a model that one builds within the framework of this theory.
- A **fundamental term** is an object of consideration in its broadest sense. In other words: the fundamental terms of the model are the objects the model is dealing with and talking about.
- A **variable** in a model is a numerical representation of such a fundamental term. We use the term variable in a loose sense as it can be either a single number, a list of numbers, or an instance of a more complicated class object.
- A **parameter** is a number that characterizes some aspects of the interaction or processing of the variables. The difference between parameters and variables may depend on context and is often somewhat arbitrary.
- A **computational model** is a model expressed in variables that can be subject to data processing.
- An **observable** in an experiment is a quantity for which a feasible measurement process exists.
- A **representation** of reality (or briefly representation) is a relation between observables and variables of the model. Ideally, there is a one-to-one correspondence between variables in a model and observables in an experiment. In any case, there should be some relation between observables and model variables in order to constitute a relationship between the model and the reality.
- An **implementation** of a model is some software that translates a computational model in operations on some hardware.
- If we execute an implementation of a model with some input variables, we call this run a **simulation**.

The focus on computational models has at least two direct consequences:

1. Since we usually work in an interdisciplinary context, we have to use software implementations and data formats that are widespread and easy to handle. In practice, we are restricted to data that can be expressed in the form of real or integer numbers. If the fundamental terms of the model are not generically numbers (e.g., in the case of sentiment analysis or image analysis), we must be aware that we need a mapping from the non-numerical terms to some sort of numerical representation. Such a mapping comes with its difficulties and is usually the source of various errors, some of which are systematic. Although the development of artificial intelligence probably will enable the classification of more complex data, e.g., the analysis of graphs, these methods are not yet widely available or easy to use outside of relatively narrow contexts. The range of problems one encounters if one deals with non-numerical data is broad: It starts with the fact that the translation of non-numerical into numerical data usually requires some classification. The classification criteria themselves are often chosen in an ad hoc manner and do not rely on a clear scientific strategy. The criteria reflect expert knowledge but lack proper rationalisation.
2. If the first restriction results from the request that the fundamental terms of the model are expressable as numbers (or sets of numbers), the second is about the simulations we perform with these numbers. Simulations are only helpful if one can efficiently do them. That means algorithms and hardware must allow performing the necessary computation within a timescale that is compatible with the needs of those who take the model outputs bases for decision making.

Based on these considerations, one must be aware that the construction and implementation of the model is a process that requires several tightly connected steps. Fig. 2 gives an illustration of the process as we see it.

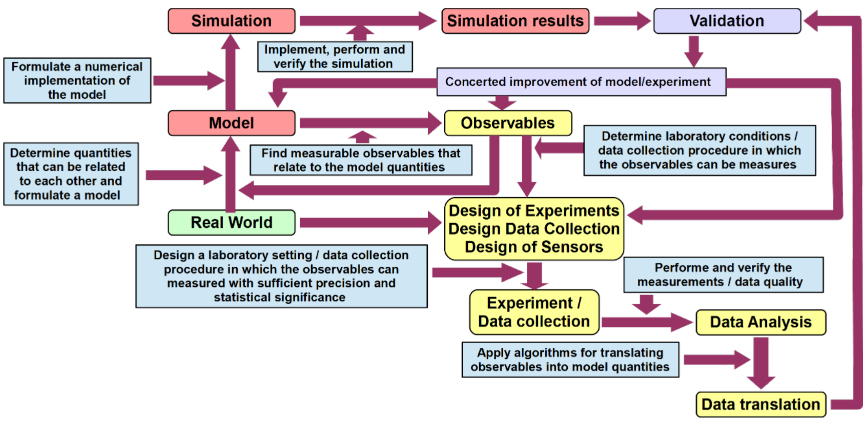

Fig. 2. The process of modelling.

Our analysis of the model building process establishes the first set of responsibilities: one has to perform all indicated steps according to the standards defined by the area of expertise and the methods one uses. These standards include scientific aspects and standard operating procedures in managing complex processes such as software development. We will not discuss this first set of standards in detail here. They are, again depending on the field, developed to a high level of sophistication, sometimes even formulated in terms of norms. Instead, we focus on two aspects that seem to be obvious but are nevertheless not intensively discussed and, in our experience, the source of problems in many projects that include computational models, namely modelling as a systemic task and the difference between implicit and explicit knowledge.

2.1 Modelling Is A Systemic Task

The most important lesson one has to learn in modelling is that a model is not an entity independent of the user and the experimental input (by the term “user,” we understand those taking the results from a simulation but not necessarily involved in the technical aspects of model building). From the project management point of view, it turns out to be a significant challenge to orchestrate the interaction between users, data providers/experimentalists, and those developing and implementing a model:

- The users must be well informed about what they can expect from a model and outside its scope. Whereas a model may be suitable for, say, the purpose of formulating scenarios or making the qualitative aspects of the system's dynamics transparent, it may not be able to produce reliable predictions. Different reasons may cause this: for example, the input data cannot be given with sufficiently high precision. Alternatively, the model itself bases on assumptions that may not be appropriate in the situation under consideration. Responsible project management must clarify what the different partners can expect from each other.
- The modellers often have a solution that seeks a problem (in practice: one aims to transfer software developed for the simulation of one type of system to another one). The project manager is obliged to raise sufficient awareness to ensure that a model matches a given problem without too much need for reinterpretation of the model’s initial semantics. We point out the responsibility of project management: One should not assume that the modellers have the necessary domain knowledge and experts in the domain are usually not familiar with the technical details of simulation software. The project probably is not either but has to initiate the necessary discussion process.
- Taking up the previous point, but in more generality: Project management must not expect the individual actors for the different tasks in Fig. 2 to act without guidance. This means that people in simulation and experimental science are certainly experts in their respective fields. Nevertheless, this does not imply that they see how their joint efforts result in a benefit. In other words: Bringing together competencies and orchestrating their results such that a benefit emerges is a competence on its own. Project management must recognise opportunities for the emergence and exploitation of synergies between modelling, simulation, and experiments.

A specific but frequent challenge emerges at the boundary of pure science and application. In pure science, one often wors with so-called “toy models.” Their purpose is to elucidate the mathematical structure of the model and to study qualitative features of the emerging dynamics of a given system. Sometimes, such models are surprisingly powerful, despite their deliberate simplicity (see, e.g., the Ising-model in the study of phase transitions of magnetic systems. There is a temptation to transfer such models from a qualitative into a quantitative, real-world context. A well-known example that illustrates the situation in our view quite well is Schelling's segregation model. The Schelling model shows how micro-motives can transparently lead to macro behaviour. Although one can learn an important lesson from this model and its simulations, it is certainly not a tool that one should use in urban planning, and it is certainly too simple to represent and explain actual social dynamics; for a very illustrative discussion see [11].

2.2 Modelling Requires Turning Implicit Into Explicit Knowledge

What one can model and simulate depends on the knowledge that can be injected into the model-building process in Fig. 2. In what follows, we distinguish implicit and explicit knowledge. By the term explicit knowledge, we understand the knowledge that

- is well-defined and can be expressed in some reasonably general language such that it can be communicated to sufficiently well-informed non-experts outside the domain of expertise on the consideration,
- is sufficiently formalised (quantitative and qualitative) that one can translate it into some algorithm.

In comparison, implicit knowledge consists of unwritten but (within the domain of expertise) generally accepted assumptions and standard operating procedures. In the context of science, this difference is quite often well understood, in the sense that we may know from experiments that a particular procedure works but not always why this is the case. However, even if we lack this knowledge, we are in science usually aware of this fact, and various practices have been established to deal with this situation (e.g., phenomenological models and explanations).

The situation is different if one uses models for decision support. Many historically grown social and economic structures are not well understood, even though they are working in a very stable manner (one may argue that it is precisely that stability that made a detailed analysis unnecessary). To give a famous example of a seemingly simple consumer good: shoelaces (we learned about this example in a televised interview with the German economist Hans Werner Sinn, who brought it up to illustrate a similar point). It is probably possible to get shoelaces in almost all locations worldwide, provided a certain standard of living has been achieved. It is reasonable to assume that there is probably no single human being who understands all aspects of the production and distribution of shoelaces in all detail. Nevertheless, the supply of shoelaces seems to be a self-organised process without any plan (or any underlying model). Shoelaces are a simple product; today's technology and politics give examples on all levels of sophistication (the information flow in a medium-sized enterprise usually differs from what one reads on the organigram, and the resilience of supply chains in a globalised economy is a very recent and intricate topic). In the situations described, one is not even capable of formulating a phenomenological model; not only does one not know how the system's components interact, but the components themselves also may not be known. Whereas experimentalists in a scientific project may not know everything about the system they analyse, they usually have a good understanding of what they do not know. In contrast, the actors in a social system may not even know that the system exists.

Self-organisation is ubiquitous in human societies, and deliberately so, at least in liberal ones. At first glance, it may look reasonable to analyse socioeconomic processes in sufficient detail such that a model of that process can be set up. In a very optimistic worldview, turning implicit into explicit knowledge is a process that is worth it on its own. We adhere to this position in the context of science but point out some issues that one needs to consider, particularly if one applies models in a broader context:

- A model may shift established balances of power. If a predictive model existed of a sufficiently large part of society, those who have access to this model would have considerable power. Modelling enables control, and control implies power, which can be misused.
- Applied to socio-economical and socio-technological systems, a model (particularly but not exclusively a successful one) seduces to centralise process control. This may be sensible, especially for cost-effectiveness, but may hinder the self-organisation of processes. Self-organised processes have specific advantages (e.g., usually they are resilient), which is not always the case for centralised processes. One has to evaluate carefully whether one wants to give up the benefits of self-organisation, and the first step on this path is the construction of a model.

The consequence of these considerations is by no means a request for less modelling. Nevertheless, if models result in power, we should enforce and guarantee transparency and equal or, at least, democratically controlled access to the models; for the importance of transparency, see, e.g. [1]. This holds for models guiding political processes and, as well, models in science (the term “democratic” then refers to the scientists involved in the project and the standards and practices of their respective fields).

2.3 Models Can Enable Thinking But May Also Set Up Limitations

In science, the problem of models that are (mis-)used for exercising power seems to be of minor relevance. This is because, in science, the role of models is usually not the control of processes but insight. Nevertheless, if a computational model turns out to be supportive of experiments, a sort of feedback may start to take effect. In practice, models are not complete. And even if they were, there would undoubtedly be settings that are easier to simulate than others for technical reasons. Then, there is a tendency to do what can be calculated and not necessarily what is most interesting from a pure domain-specific perspective. This may be perfectly reasonable, but one should bear in mind that the question of whether a process can be simulated is somewhat extrinsic to the process itself.

This aspect becomes apparent in teaching, where models play an important role. One often discusses simplifications and idealisations such as frictionless movements or perfect crystals. One tends to choose examples because they result in equations that can be solved with the mathematical means available to the students in the context of the applied model. But: Whether or not a process is easy to calculate does not necessarily say something about its relevance to nature; this is a deep issue, see [12]. Restricting the analysis of processes to that one can compute may be reasonable, especially in teaching. However, one should always keep in mind that there are phenomena outside the somewhat artificial boundaries imposed by requesting computable models.

The fact that models may impose limitations on scientific investigations is not restricted to teaching. As an example, we point out that the notion of an "integrable system" in the theory of dynamical systems is a fundamental concept, somewhat loosely defined as systems with conserved quantities and therefore restricted to submanifolds of the phase space under consideration. Chaotic behaviour is closely related to integrability, or better, its absence. Even though mathematics was already known in the 19th century that the behaviour of dynamical systems could be very complicated and unstable, in more general science, the notion of "chaos" came up only around 1960, and somewhat as a surprise to scientists used to integrable models. Still today, at least in engineering, chaotic behaviour is often regarded as the exception (and not, as it is, in fact, the rule).

As a side remark, engineering is an interesting case in that respect. One could claim that a large part of engineering consists of the attempt to construct systems so that they can be described and controlled by efficiently computable models.

2.4 Interdisciplinary Work Is No Excuse for Diffusion Of Responsibility

A more general issue concerns the convergence of experiments and models. The different tasks in Fig. 2 require several different fields of expertise. In practice, one observes the danger of a certain diffusion of responsibilities. This problem is well-known in a single scientific project and may be avoided by project management. As soon as one starts to use models as tools for decision support and works in a broader setting, the problem of the diffusion of responsibilities becomes somewhat more pronounced.

On the one hand, modellers tend to complain about the "lack of data" and use this argument to justify the shortcomings of models. Project management must clarify that the limited availability of data may look like a bug but is actually a feature to deal with responsibly. That means: one must set up models in a manner that can get along with the available data.

On the other hand, there is a specific danger that those using the output of models do that by regarding the models as black boxes and putting trust in them without scepticism. The scepticism certainly includes the results produced by the model. A responsible and practical form of scepticism is constructing a broad range of plausibility checks. Those who use the model's output usually have a rather precise idea of what this output should look like, especially for some extreme choices of system parameters. Such checks can be pretty efficient, but one needs to organise them properly. Again, in the relatively narrow setting of a scientific research project, this practice is well established and can be implemented easily (probably because experimentalists are familiar with checking their setup by measuring some boundary cases). As soon as one uses models for general decision support, the employment of plausibility checks by extreme scenarios requires some management skills. Besides the fact that the communication between those implementing a model and those applying the results has to be organised (and is subject to limited resources, especially time), a psychological barrier has to be passed. Project management must clarify that models usually do not work right from the start but require certain debugging and polishing. We repeatedly encountered considerable criticism from those using the results of models at the initial stage of implementation. The line of argument was: “If your model cannot deal with a straightforward situation, how can we trust its results in a more complicated setting?”. Responsible project management makes it transparent at an early stage that an iterative calibration and refinement process is an instrumental part of modelling.

According to our experience, there is a potential misunderstanding about modelling: Sometimes, users are under the impression that modelling means implementing a small number of very fundamental relations and natural laws, which, if implemented correctly, will result in precise predictions in all possible settings of the input variables. There is a hidden danger in this optimistic perspective. If a model were omnipotent (at least within the frame of the system simulated) framework for prediction, the task of posing appropriate questions would be trivial. Say it in other words: if the modellers can simulate everything, the users can ask anything. However, models are most often much more specialised in the sense that the range of questions they can answer reasonably is limited. For the design of a model and a simulation, it is essential to know what questions the users want to answer. Responsible use of models requires carefully designing questions and being aware of the limitations of models. This holds especially if models are developed in a collaborative process of modellers and users. It is then the responsibility of the users to define in sufficient detail what types of results the model should be able to produce.

We illustrate this with examples from science and decision support. Within a specific scientific discipline, for example, solid-state physics, the relation between experimental observables, the model's variables, and the difference between qualitative and detailed quantitative statements are well understood by all partners. Discussions may still be necessary, but they build upon a tradition that is part of the discipline (see also Sec. 3.2). Again, the prototypic example is the Ising-model; see for example [13]. Interdisciplinary discussions are more demanding than intradisciplinary ones. A famous example is the application of Lotka-Volterra models on problems in ecology, e.g. the lynx – snow hare – cycle. A mathematical model can never predict in detail how populations develop, and it is also not possible to determine their exact size in the field. The question is then what one can learn from such model, i.e. what questions are answered by the model and what conclusions would be an overstatement. The discussion between modellers and ecologists should then focus on whether and how the model's output can be made helpful for ecological reasoning. Such discussions become even more necessary in a transdisciplinary context (s. Sec 4.1). We refer to epidemiology, where coarse-grained models get input data of somewhat limited quality (the data reported may be satisfactory for, say, influenza but do not match the needs for a novel disease because essential aspects are not analysed). Note well that there is a trade-off between the resolution of the model and the amount of input data necessary. A finer resolution would require even better data, see also the next section. For decision support during an ongoing pandemic, the modellers have to clarify that this trade-off exists and that predictions about, for example, the number of fatalities cannot be made (although giving bands is possible). In contrast, statements about, for example, the transition into an endemic state are possible with reasonable reliability.

We emphasise the importance of this discussion. In decision support, the scientists and modellers must realise that their partners from, say, politics are not necessarily familiar with the do's and do not's of a specific scientific discipline. Nevertheless, these partners have to communicate science-based decisions. In our view, modellers must actively support this communication and in a way that reflects the specific challenges of communication with a general audience.

2.5 Observables Are Not Variables

It may sound relatively trivial, but it turns out to be a fundamental management challenge to bring model variables and experimentally measurable observables into a relation. Or, to state it differently: in a project that uses models, variables and observables must be chosen to be related efficiently and require tools available to those working on the project. This matching process is complicated.

This complication results, first, from intrinsic reasons. Model users and modellers are not necessarily aware that variables and observables are linked by a process that includes several steps (for a still simplified overview, see Fig. 3). Here, $S(t)$ represents the system under consideration at the time t. We point out that by S, we don't understand a number or some other data structure but an actual system. The downward track represents experimentation. Some aspects of the physical system can be observed; these are the observables $O(t)$, which are still understood as physical phenomena. The relation/mapping between system and observables is given by a relation $O(t) = \Theta(S(t))$. The observables can be measured, which produces signals $M(t) = \Gamma(O(t))$. The signals are assumed as data structures, e.g., time series of numbers or digital images. We point out that, for example, the production of a digital image is again a process that involves many steps, but most of them are standardised and/or their limitations are purely technical and do not involve the issues we are raising here. Going upwards from S represents the modelling track. First, the real world has to be mapped to a mathematical model $R(t) = \Phi(S(t))$. $R(t)$ are mathematical objects of some kind. For obtaining a computable model, these mathematical objects need to be represented by some data structure $N(t)$ (which, at the end, is a finite bit string): $N(t) = \Psi(R(t))$. The data structures $N(t)$ are the model variables. We point out that this step requires some subtle considerations: The mathematical objects $R(t)$ need not to be conventional numbers but can be complex numbers, vector fields, manifolds, or other mathematical objects. The translation of, for example, a vector field into a data structure is a step that poses its specific challenges (the choice of a coordinate system, for example). Finally, model variables, by some processing, are transformed into data $C(t) = \Sigma(N(t))$, which one can compare (in a mathematical, means quantifiable sense) with the signals $M(t)$. The possibility for this comparison is instrumental for model validation and is the basis for model parameterisation.

Second, modellers and experimentalists have somewhat different objectives. The more precisely one can describe an object, the better it is for the model, or to be precise, the easier modelling is, at least from a conceptual perspective. Relying on probabilistic concepts or statistical quantities such as averages is done chiefly if enforced by limited computational power or memory space. However, data acquisition by experiment requires specialised equipment, time, and staff. It sometimes is even impossible (from a modelling perspective, in vivo data acquisition would be desirable for calibrating models of cellular processes, but the necessary experimental possibilities are still not always available).

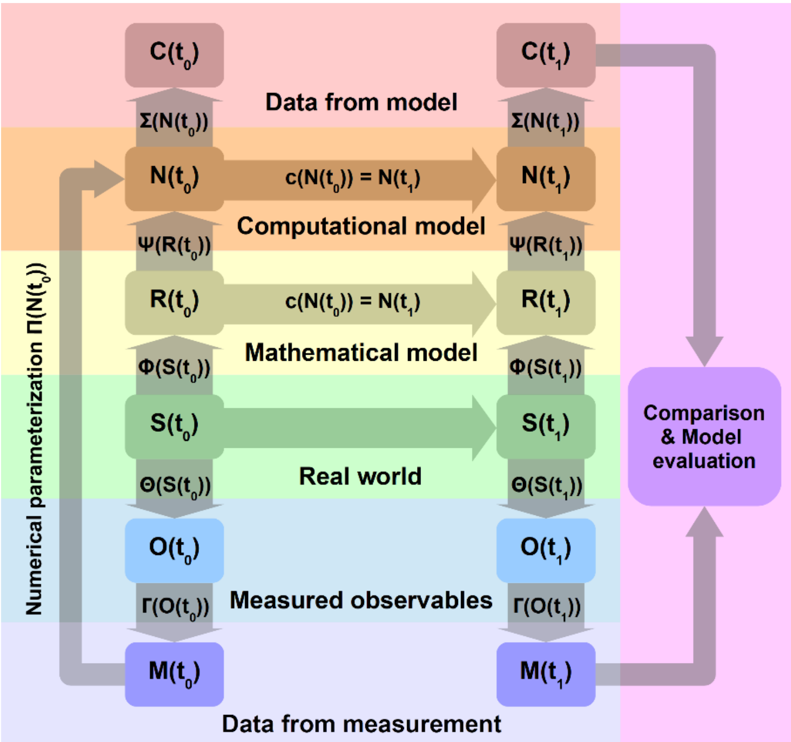

Fig. 3. Relating observables and variables. After comparison and model evaluation, the model and the experiments may be modified, s. Fig. 2.

The task of responsible project management consists of finding a position that matches given experimental boundary conditions. This means setting up a permanent negotiation process between modellers and experimentalists oriented on the possible and not on the desirable. In a scientific context, this may be, if not easy, generally accepted practice. As soon as one enters the realm of model-based decision support, the tension between available and desirable data can become considerable, with often only limited mutual understanding for, e.g., legal boundary conditions (to mention a non-technical limitation).

As a general observation, we note that model building tends to be output-oriented in the sense that modellers want to produce optimal results but often do not care sufficiently about the necessary input. In part, this is a reasonable practice because there is a historically grown division of labour between those developing methods and those applying the tools based on these methods. The developers tend to work with assumptions or "toy data" often generated based on assumptions. Responsible modelling requires balancing the desire for maximal output with the realities of the available input. We illustrate this with a recent example. Modelling the COVID -19 pandemics was an important challenge in the last two years. Modelling pandemics is a prototypical example for applying agent-based models. However, agent-based models constitute only a small part of all simulations in epidemiology [14]. Besides the fact that agent-based simulations are rather time-consuming, one must consider a fundamental problem. Agent-based models allow, in principle, to model very precisely the behaviour of representative batches of the population. But this precision comes at a price: one has to provide the according input data, say contact structures in a population. As it turns out, simpler, less data-hungry models may give a more reliable picture of the potential scenarios than models that would provide a very detailed picture if only fed with appropriate but, in practice, not available data.

3 Models In Science

There is a vast literature about the role of models in science; for a short and well-written overview, see, for example [15]. In our practice, two categorisations are of particular relevance for ethical considerations. First, one can use models for pre- and post-processing of data. The two uses pose different challenges concerning responsibility. Second, one can classify computational models in those trying to virtualise a given situation as precisely as possible (we call these models complete) and minimal models. The focus of minimal models is generality; one looks for general mechanisms and strips them from the details of specific systems. One explores those aspects of the dynamics independent of the details of this implementation.

3.1 Pre- And Post-processing Data

One can roughly categorise the role of computational models in science into two classes: Pre- and post-processing data of an experiment (whereby pre-processing includes those cases where one does not perform an experiment at all).

Post-processing happens in cases where one pursues one or several of the purposes discussed in the introduction. "Understanding the past" and "Giving meaning to data" are undoubtedly important cases, whereby meaning includes the transformation of sensory signals into measurements, s. Fig. 3, but also includes the support of the interpretation of data.

Pre-processing experimental data encompasses virtual testing, scenario building, and prediction. For example, microscopic models can relate fundamental properties of a system to some macroscopic observables. From a scientific perspective, such microscopic models are crucial for the reductionist program. For engineering, material models enabling virtual construction are vital tools in modern design processes.

From an ethical perspective, the distinction between post- and pre-processing is relevant insofar as in post-processing, the further use of outputs of an experiment stays in focus, whereas pre-processing may guide the implementation of the experiment itself. One could think that, because the experiment already happened, post-processing data is of minor ethical relevance. This is, of course, not true. For example, in most industrialised countries, the regulations for animal testing require ethical balancing, which means a justification of animal suffering compared to human benefit (see e.g. European directive 2010/63/EU[3]). The use of models can do both, enhance the benefit of data from an experiment, and change the type or reduce the amount of experimentation for the same benefit.

Somewhat colloquially said, the challenge of pre-processing lies in the fact that models start to guide experiments. Using models as guidance can be misleading in various ways: A system may seem to work in theory or *in silico*, but does not in reality (It is only slightly polemic to say that this is not a real problem but what experimentalists expect from modellers). Another danger is that a model may show something to be unfeasible, which some clever engineering can implement nevertheless. Even if we assume that the modelling has been done with all precautions established in the field or reasonably possible: All modelling is based on assumptions that are sometimes by no means obvious, and one tends to disregard them. In our experience, only a permanent and maximally transparent discussion of assumptions between all involved parties in a project can reduce the problem of (mis-) guidance. Again, we point out that this process is by no means easy and an actual intellectual task (One needs to pass the boundaries between different fields of expertise).

3.2 Minimal And Complete Models

Again, we discriminate models into two large classes, whereby we are aware that a continuous spectrum would be more appropriate. By a minimal model, we understand a model that is as simple as possible and contains only those variables of primary interest. In a minimal model, one deliberately simplifies external effects, environmental conditions, and complicated details of the interaction between the parts of the system under consideration. Such models turned out to be of enormous value to science. As examples, we note the Ising model in statistical physics or the Schelling model in the study of segregation processes (by the way, the mathematics behind these two models are closely related). Minimal models can shed light on qualitative aspects of the system's dynamics and, despite their seeming simplicity, can show rich emergent

[3] `https://eur-lex.europa.eu/eli/dir/2010/63/2019-06-26`

behaviour. However, the purpose and scope of the model needs to be clarified as soon as minimal models are used outside a narrow range of scientific investigation. Results produced by minimal models are often easy to interpret, mainly because the relation between cause and effect is transparent. The problem with minimal models is that it is difficult to transfer their results into real-world situations. What seems to be a plausible cause and effect relationship in a minimal model needs not be one in the real world.

Complete models try to represent as much of reality as is possible or necessary. Put costs aside, this seems to be a reasonable approach. However, such models contain a lot of parameters and turn out to be difficult to calibrate. The main difficulty of calibration is that if there are various parameters, their respective values may be underdetermined regarding the available data. At this point, one usually quotes John von Neumann, who is said to have said: “With four parameters I can fit an elephant, and with five, I can make him wiggle his trunk”. Today, partially caused by machine learning development, science developed several sophisticated procedures to cope with underdetermination. In our view, the problem with complete models is that the preparation of input becomes a highly complex task. This is first because the input tends to be large. Second, some types of input data are easier to get than others in reality, which leads to biases. Systems in which one couples different processes, some that one can parameterise by laboratory experiments and some that cannot be isolated (e.g. socio-technological systems, for which the technological processes are usually easier to parameterise than the social components) pose a particular problem. The more input parameters one has, the more difficult it becomes to evaluate the quality of the output because whether or not the limited precision of an input parameter has a relevant effect depends on the size of the statistical uncertainty and the influence of that parameter. Robustness analysis is often only of limited help because biases in input parameters, mainly if one deals with social systems, tend to be correlated.

We emphasise that phenomena observed in minimal models are generic in that they are not the result of some, potentially very special, circumstances of a specific situation. From this perspective, minimal models are general models.

4 Models In Decision Support

Models have a well-established and continually refined role in science, with a development based on practical experiences and theoretical studies. Especially with the growing digitalisation but also the awareness for data science, models get an essential role in decision support, be it in an emergency such as COVID-19 pandemics or be it in planning, e.g., in an economic or urban context.

One may argue that decision support is somewhat outside the boundaries of science, and one should not mix up scientific and nonscientific applications of models. However, we advocate that the distinction between scientific and nonscientific uses is somewhat artificial and neither beneficial for the goals of science nor modelling as a scientific field. First, the boundary between science and non-science is challenging to draw. In our experience, the discussion of what belongs to science rarely leads to a sensible conclusion and most often ends in a somewhat sterile dispute about definitions.

We see a considerable benefit in the broader use of models. First of all, we believe that models can contribute to decision-making. In addition, if there is bidirectional communication between users and modellers, confronting a model with reality can promote science by adding novel questions and initiating developments that have the advantage of dealing with testable scenarios. Often, external questions can promote interesting internal scientific developments. If at all, the distinction between pure and applied science can be characterised by the notion that pure science studies the internal workings of the system in deliberately simplified boundary conditions (the laboratory conditions), whereas applied science asks for what can be said scientifically about a system immersed into a complex real-world environment. Note well that in this reading, applied science is not just about applying the results of science. Applied science also aspires to do science outside the laboratory.

Understood in the manner described above, the authors are all involved in applied sciences, particularly science applied to decision support. This includes applications of modelling in purely scientific projects such as ACDC, in which a sophisticated interaction between modellers and experimentalists is a key project objective. This involvement also includes decision support in processes that include the whole society.

We investigate two large clusters of problems, the first one relating to the position of science in society and the second addressing some hidden aspects of optimisation. We conclude with some general observations, which we realised as being important.

4.1 Transdisciplinarity

We distinguish inter- and transdisciplinarity by assuming interdisciplinarity as communication and collaboration over the borders of different branches of sciences but within the general context of science. We understand transdisciplinarity as the interaction between the actors inside and outside science. For a detailed discussion of these concepts, see [16] or [17].

Science is a cultural practice with its own rules, language, codes of conduct, and signaling systems. As a scientist, one must accept that communication with the non-scientific sphere requires finding common grounds and using a common language. There are several challenges, which we will address in what follows.

Arrogance: Science And Democracy.

Science is undoubtedly one of the most successful collection of practices in human history. This fact is seducing. Scientists tend to justify the advantages of a scientific mindset and approach by the successes of science. Some conclude that because science works well, not only those doing science but everybody should act and think like a scientist. There is a danger of going even further and regarding those unfamiliar with what scientists call a scientific discourse as unfit for a general discussion.

We regard this as a naïve position. Although we are thoroughly convinced of the value of science and its methods, we are aware that either the concept of a “scientific discourse” must be stretched until it contains all forms of rational discourses or essential questions, e.g. about moral values or inter-subjective impressions as they occur in art, are not regarded as part of the discussion. For computational modelling, this means that even if one is convinced to have powerful tools that enable a well-grounded insight into natural and social processes, these insights are not above democratic processes. The scientists must make these insights a part of democratic negotiations that respect, for example, discussions about moral values. These discussions may still be rational but are not necessarily following what one usually regards as integral parts of the scientific method (evidence-based argumentation, falsifiability, etc.).

Analogies, Metaphors And Speaking Plain English.

Transdisciplinarity requires abstaining from jargon. We point out that this requires a lot of effort. In fact, “jargon” in science means the use of sophisticated formalised concepts, which in some cases require a level of abstraction that one can only master after intensive occupation with a specific topic and its formal apparatus. We emphasise the role of formalisation; often, the underlying idea of a concept is well accessible; the formalisation of the concept requires training and detailed knowledge about the formalisation itself (for a detailed discussion of this point in the context of the physical sciences, see [18].

There is a real problem: One cannot explain complex formal abstractions “in a sentence.” We see only three, partly connected, ways out of this:

3. Speaking in analogies: One compares something unfamiliar with something more familiar to the audience. We emphasise that there are various aspects to consider. First, when explaining models, analogies most often focus on functions, structures and dynamic behaviour. Second, they are rarely faithful. That can be helpful because they enable highlighting relevant aspects and neglect those of lesser or no relevance. On the downside, one always has to keep in mind the danger of overstraining analogies. Third, if A is an analogy of a process P, the audience may be familiar with A's formalisation but not with the one of P. This is a situation of particular interest because one can scrutinise the extent to which two formalisms are equivalent with some rigour, and one can state the limits of analogies. At the same time, one can profit from the power of formalised reasoning.
4. Explanation: One tries to explain formalisations. As desirable as it would be, this is often unfeasible.
5. Modularisation: One modularises the explanation of a complex process and makes the modules, their dynamics, and their interactions transparent. Thereby, the processes taking place on the level of modules should be clear to all participants of a discussion but not necessarily the internal dynamics of the individual modules. An example is a recipe: The act of cooking requires the realisation of a series of biochemical processes. The (evolved) practice of cooking modularised the process by using building blocks that are robust (small deviations of temperature, amount of

ingredients and the like lead to only small changes in the outcome; processes close to transitions such as caramelisation are for the advanced), standard cooking ware is employed, and there is usually no need to know about the chemistry or physics that takes place.

We emphasise two further issues. First, the relation between concepts and formalisation is not unidirectional. The formalisation itself may lead to an extension of the initial concepts. A famous example is antimatter (or, to be precise, the anti-electron, now termed positron), which P. A. M. Dirac introduced on purely mathematical grounds. Such concepts, which originate not from experience or experiment but as a consequence of mathematical reasoning, are particularly hard to convey to an audience without formal background.

Second, we point out the difference between analogies and metaphors. We quote H. U. Fuchs: "We all have access to abstract schemas that form through organism-environment interactions. Understanding something (or making something understandable) means bringing a description/explanation back to these fundamental abstractions/schemas (which are used in metaphors, and metaphorical webs, which, in turn, are used in narratives).", see also [19]. In order to explain processes/structures, one uses metaphors; by analogies, one compares and relates processes/structures.

A narrowed version of a metaphor is the idea of an abstract or conceptual data structure, say a vector. Vectors are inspired by the combined notion of length and direction. The concept of a vector is embedded in a network of other concepts, such as angles, rotations, parallelograms, dimensions, et cetera. Furthermore, our intuitive notion is complemented by a rigorous mathematical formalisation (which, in turn, inspired further concepts such as infinite-dimensional vectors for which we lack a complete intuition, see above). We discuss this example because it highlights a significant problem. The representation of a vector in a model (the respective model variable as it appears in Fig. 2) is usually an array of numbers. For communication, it is essential to understand the difference between the representation of a variable and the concept behind it. The representation carries much technical baggage, such as coordinate systems, which bury the idea of a vector.

In a communication, using the representation instead of the concept may be tempting for those familiar with the former; but it is usually not helpful for those lacking this technical familiarity. Where does the temptation come from? Speaking in plain English about representations is usually relatively easy (which is why we can "explain" representations even to computers using, from a linguistic perspective structurally simple, programming languages), whereas verbalising concepts requires hard work and accurate language skills.

Questions And Answers In Science.

Much postmodern critique of science tries to show that science is a social construct, and therefore, the scientific method has no privileged position for understanding the world. We cannot enter this discussion on a broader level; concerning computational models, we have to address some questions and points of critique:

1. As natural scientists, we take the existence of an "objective reality" as a fundamental assumption. However, one must not misunderstand a model for this reality. It is at least a point to keep in mind that a computational model is defined and bounded by many constraints (some are economical and therefore structurally social).
2. From the point of view of a natural scientist, the answers of science are at least approximations to objective truths; but the according questions are not. If one regards the development of science and particularly modelling in a transdisciplinary context as an interplay between questions, answers, refined questions, and consequently the further development of methods, complex models do have socially constructed aspects.
3. A computational model needs input data. As discussed in section 3.1, the input data can be subject to ethical evaluations. Since what one can compute is a function of the available input data, the ethical considerations concerning input data affect the possible modelling results.

Even if one does not share postmodern positions, the points above show that computational modelling, especially if applied in a context that includes partners from fields of expertise outside of science, is certainly affected by social and cultural processes.

4.2 Optimality: Give Options, Not Advice!

Decision support often aims to find an optimal strategy or implementation of procedures for a given task. One major challenge if one works outside a strictly scientific context is finding a proper definition of optimality. The result of any attempt to find an optimal solution depends on what one regards as desirable. As shown in what follows, the discussion of how one defines optimality has some aspects and lies in our view at the core of the ethical aspects of modelling. The problem is multi-layered: Even if one has a quantifiable desirable goal, there may still be several additional boundary conditions that one has to observe to establish procedures that yield optimal results and do this in a fair manner.

Optimality And Fairness.

In a purely technical context, optimality is quite often easy to define. Even then, one must be aware that proper balancing may not be trivial if there are different criteria for optimality (for example, efficiency and efficacy or quality and output in engineering). This is also recognised in a business context [20], where the discussion about KPIs (key performance indicators) has reached a high level of sophistication.

In a political context, optimality is most often a question about values. It is a hallmark of democracy that such questions have no general and definite answer (derived from some dogmatic set of principles) but are subject to permanent discussion in each case.

Optimality becomes even more involved if resources are limited, and one includes criteria considering the fairness of distribution. Here, computational models can be beneficial. As an example, we mention [21], a study in which the distribution of a limited number of defibrillators over different areas has been investigated. At first glance, optimality is easy to define and given by the number of saved lives. However, this would imply placing the defibrillators preferentially in urban areas, where many potential patients can profit from their presence. A fair distribution should not disadvantage those living in rural areas. In [21], a computational model was used to find a distribution pattern that maximises the number of saved lives and considers a fair distribution of resources. From the perspective of computational modelling, one has to quantify and combine two different criteria for optimality. The quantification makes implicit valuations explicit (compare with Sec. 2.2), and this itself poses non-trivial political challenges.

Ethics And Second-Round Effects.
The distribution pattern of defibrillators is an example of a static situation. Computational models are advantageous if they clarify ethical considerations in dynamic processes. Here, one has to distinguish between first-round effects and subsequent processes, which we summarise here by the term second-round effects (being aware that there are third-, and in general and nth-round effects). We illustrate this by a study [22] that used models to optimise the distribution of limited vaccines in influenza pandemics. A general goal is undoubtedly to maximise the number of saved lives. In the case of influenza, this implies that in general (there are exceptions, though) that the efforts should be focused on the most vulnerable. There are two different ways how the most vulnerable can be protected. First, they can get a vaccination. Second, they can be protected from infection by reducing the number of contacts with already infected ones. If one studies a "common influenza" and does not consider measures such as lockdowns or quarantines (which in 2017 looked outlandish), one can reduce the spread of the disease by vaccinating that part of the population first, which contributes most to the distribution of the infection. In general, the group of the most vulnerable (in the case of influenza, usually the elderly) and the group of the most critical spreaders are not identical. Whether direct vaccination (the first-round effects) or protection by reducing the spread of the disease and vaccinating the spreaders first (a second-round effect) results in a maximal number of saved lives depends on various parameters and can be studied by a computational model.

As we realised, the communication of such second-round effects is far from easy. We point out that whether one directly vaccinates the most vulnerable or protects them by stopping the spread of the disease always serves the same goal, namely the maximisation of the number of saved lives (which means the lives of the most vulnerable). However, one must carefully explain why protection of the most vulnerable sometimes may be most effectively achieved by the prioritised vaccination of a different part of the population.

That second-round effects occur can often be made clear with qualitative arguments. However, whether or not they can become prominent, even dominant, is usually a quantitative question. A computational model helps to show first that there are settings in which second-round effects are relevant and second, which factors influence the extent of this relevance.

Optimality And Observables.

In physics, a central principle states that the laws describing a process must be independent of the choice of coordinates one uses (That is an occasion where a subtle difference between models and simulations becomes apparent: one can formulate the physical laws describing the trajectory of an asteroid without referring to a specific set of coordinates. A simulation processes numbers and relies heavily on the choice of coordinates). Whereas one can firmly establish this principle in the natural sciences on mathematical grounds, the situation is much more complicated in socio-technical investigations.

We explain this with an example. In order to account at least partially for the variability of society, the individual members of the population are grouped into cohorts. Especially in medical settings, this usually happens according to age and medical preconditions. The interactions in the population are then also formulated based on these cohorts. One could use a different grouping, leading to different interaction schemes. Different groupings may represent the variability of society differently for the phenomenon under investigation and with different statistical quality. Whether age or, say, socioeconomic status is the best descriptor in a given situation is not always apparent, and the choice of the descriptors may influence the outcome of computation and the conclusions one draws from it. We emphasise that computational models that serve as tools for determining optimal solutions under consideration of ethical principles must be scrutinised for their dependence on the choice of input and model variables.

One may now ask for a determination of the best way to describe the process as a prerequisite for any use of computational models for ethical purposes. As reasonable as this sounds, it is pretty often not feasible. The evaluation of socio-economical data is difficult and expensive. As a modeller in decision support, one may be confronted with the fact that one has to work with the available data, which is not necessarily the data that would be best suited. It is, in our view, the central responsibility of the modeller always to point out that fact.

4.3 The Role Of Experts

In inter- or transdisciplinary projects, experts from different fields have to interact. Groups of experts, especially modellers, should carefully reflect their roles as soon as they become part of the decision-making process. We identified two main issues that we address in what follows.

Groupthink.

As pointed out in the introduction, one reason for using models is that they may help understand non-linear behaviour, emergent dynamics, and sometimes the appearance of seemingly counterintuitive phenomena. The challenge is distinguishing between those phenomena that are hard to understand but are real and those resulting from some possibly wrong assumption underlying one of the various aspects of the modelling process. We repeatedly observed an interesting process that belongs to the class of problems that one usually summarises by groupthink. The process works like this: a model gives a hard-to-understand result. The modellers, usually not experts in experimentation or observation, ask those familiar with the experimental aspects of the system on the consideration of whether or not the result of the model can be true. The experimentalists, not being familiar with the internal workings of the model, give a plausible explanation for what has been computed with the simulation belonging to the model. Such a line of argumentation that relies on incomplete knowledge may be the source of groupthink. Each member of the group regards his/her partial knowledge as justified by the partial knowledge of other experts. The problem thereby is not so much that all involved parties only have partial knowledge; this is unavoidable in interdisciplinary work. The problem is: What seems to be plausible in a discussion should be based on evidence or more detailed scrutiny.

The Position Of Experts In Complex Decision Processes.

In our view, the most critical problem of the use of models in decision support is the necessity of the experts to develop a proper understanding of their role in the decision-making process. Especially when models are used to evaluate or optimise ethical aspects in the decision process, experts tend to advise decision-makers. This advice is usually based on an already preselected set of simulation results. In our view, modellers must avoid this preselection in a proper decision-making process. The experts, especially the modellers, should understand their role as giving options for decision-makers. These options are then used to achieve a proper decision and represent a range of possible further actions. The modelling results should show the consequences and costs of different potential courses of action but should avoid guiding the decision in a specific direction by imposing a value system that is not transparent to the other parties in the decision process.

The other side of this is the potential tendency of decision-makers to diffuse responsibility by taking the results of modelling as such and not to apply a prioritisation or a valuation based on a transparent and ethically grounded evaluation scheme.

To say it in one sentence: For maintaining the integrity and transparency of decision making, science gives options, and decision-makers value and select them.

4.4 Computational Models As Tools For Discussion

In 2.2, we pointed out that models require turning implicit into explicit knowledge. This process is necessary for model building but is also of use in decision support. The discussion of concrete model assumptions and the possibility of studying their influence at least in a semi-quantitative way (for example, whether specific output variables are positively or negatively coupled with some basic assumptions?) helps to understand the emergent mechanisms in complex processes. The discussion between the “users” with expertise and domain knowledge and the modellers who perform simulations (always in the context of a given model) can result in a better understanding of the system as a whole and the emergent properties one observes (in reality and the model). In an ideal case, the interplay of experts (scientists or non – scientists) who discuss their experiences and verbal descriptions with modellers who turn these statements into formalised algorithms can improve the levels of insight into complex systems.

A further benefit is that an algorithmic description can be communicated in a way that is sometimes hard to achieve by prose. This communication is not only crucial for the interplay between modellers and other stakeholders; in our experience, the discussion of a model and its algorithms can also be beneficial for the collaboration of different stakeholders. We then use the model as a tool for discussion, and modellers act as intermediaries.

This function as a tool also applies to situations where one discusses qualitative aspects of system behaviour. For example, we take the question of tipping points, i.e. a qualitative change of system behaviour resulting from a small change in one or several parameters. A minimal model (s. Sec. 3.2) can help decide whether some generic dynamical properties are sufficient to produce such a tipping point. Knowing about such a tipping point is of value, even if we know that the numerical value at which the transition happens in a minimal model may differ considerably from the one in a specific and complete setting. This type of discussion is well known in physics (we again refer to the Ising – model, which is a minimal model of magnetism but shows some behaviours of phase transitions generically).

5 Models And Narratives

We focus the discussion on a central topic: The interpretation of results gained from models happens in a series of steps. This interpretation starts in the context of science, the place of production. Various methodologies, "best practices," and cultural habits exist in this relatively narrow social environment. Later, various instances transfer these results into a language that suits broader, even public communication needs. The transfer is not a translation; transfer is not only a matter of using "plain language." Instead, the communicator produces a "narrative" in which common analogies replace the system or process under consideration (see Sec. 4.1.2). We claim, however, that the scaffolds of narratives appear at an earlier stage in the process of model building.

We start with a central hypothesis, from which we derive / on which we base several questions. We cannot answer these questions in a definite manner, but we recognised them as central for discussing the relation between models and narratives.

5.1 Main Hypothesis On Models, Simulations, Storyworlds And Narratives

Although models are based on quantitative or qualitative scientific reasoning, how they are perceived and used in a context broader than that of science should be analysed with a range of tools from communication, journalism, literary analysis, and critics. In the narrow context of science, a model is a basis for mathematical reasoning. The function of a model is broadened, as soon as models and their results become part of the thinking and acting in politics, administration, and the wider public. Besides mathematical arguments, the model sets the stage for narrative elements. Again, taking up an idea of H. U. Fuchs, the model gives a story world, and the simulation is the backbone of a concrete story.

In what follows, we will use the term “narrative” instead of “story”. According to Collins dictionary, a narrative is “a story or an account of a series of events”, whereas a story is defined as “a description of imaginary people and events, which is written or told in order to entertain”.

In such a broad setting, "Reading the results of a model" is a non–trivial process that can no longer rely on some scientific subdiscipline's established standards and rules. Understanding the role of models/story-worlds and simulations/narratives and putting it into a social and ethical context requires a discussion that raises awareness of the reader's role and his/her background.

Good literature is more real than reality in the sense that a well-composed narrative contains "reality" in a more condensed and easier-to-follow form than just an account of what has happened where and when to whom. We probably all agree that writing a good narrative is a significant task. As soon as simulations are related to narratives, we strongly emphasise that communication profits from the inclusion of experts. However, we point out that one needs more than marketing (marketing is needed, but not only). One needs narrators and experts from literary studies who understand the complex relations between texts and readers.

5.2 Questions And Topics Relating To Politics And Operationalisation

Models And Novels

We compare a good model to a good storyworld: The model sets the stage for a narrative that, in some respect, is a streamlined image of reality but contains, concerning a specific set of topics, a sufficient representation of reality. Like a good novel, this narrative focuses on those parts of the dynamics relevant to the phenomena under consideration and neglects the others. As already explained in Sec. 4.2.3, socio-medical models often subdivide the population according to age. That leads to a picture in which "the elderly," "boomers," and "the younger" appear as actors. This is often reasonable but sometimes hides the fact that a similar subdivision according to socioeconomic status could be employed, which leads (literally) to a different narrative.

When one equips a model with a narrative, one needs to ask about the opportunity costs of invoking one specific narrative: The choice of the narrative one tells automatically implies that other narratives remain untold. Choosing a narrative (which happens already at an early stage in model-building when one chooses the model variables, see Fig. 2) must be done considering potential uses for ethical purposes later. Conversely, if one has a model and evaluates its use for ethical issues, one must ask whether the model is an appropriate stage for narratives that illustrate the ethical question under consideration.

Models, Narratives, And Communication Structures

In larger projects or organisations, models/storyworlds and simulations/narratives are embedded in communication structures. If one aspires to establish a smooth and correct interpretation of simulations in an organisation, a central question is: How can we avoid (maybe interest guided but probably more often unconscious) misinterpretations? The problem of misinterpretation is closely related to the "Give options, not advice!" – statement we discussed in Sec. 4.3.2. The narrative should inspire thinking and discussion but not predestine their outcome not justified by the facts.

It is not only about misinterpretation but about interpretation in general. Note well: As soon as one accepts that models and simulations go together with narratives, we have to accept generic properties of the latter as a part of the whole process. To express this colloquially: The fact that narratives can have many interpretations is not a bug but a feature of literature.

One further point is that how a narrative is understood depends on the culture in which that reading happens. If narratives transport/communicate the results of models, we should compose narratives with an awareness of the difficulty of writing stories in an intercultural context.

Secondary Literature

One usually values the primary texts higher than the secondary sources in literature and philosophy. This is different in the natural sciences, where almost nobody learns, for example, quantum mechanics via reading the original papers. This may be a pity in some specific cases where the original works are written by true masters of the field and contain deep insights. In general, however, the secondary literature clarifies basic concepts and uses a more accessible presentation and improved formalisation. Secondary literature in the natural sciences is quite often easier to understand and, therefore, more efficient in teaching sometimes rather technical ideas. One can explain this observation partly by considering that the authors of secondary literature have been in the same situation as the novice is when studying a new topic: One has to master an idea that one has not produced by oneself. We emphasise this, because, in our view, one must not regard the communicator as solely supportive. Those translating a model into a narrative contribute an essential part of knowledge in a transdisciplinary process.

5.3 The Role Of The Reader

One big difference between narratives in literature and narratives derived from scientific model-based simulations is the multiple authorship of the latter. Literary works with more than one author (not to speak of five, ten, or twenty) are almost nonexistent and, if at all, are certainly part of the experimental branch of literature. Concerning ethical considerations, this raises important questions. What is the individual author's responsibility for the narrative produced from the simulation results?

One can extend this question. If we compare the narratives related to models and simulations with other literary works, is there an ethics of literature that helps us understand how we can deal with models/simulations? On the one hand, there is the freedom of artistic work (which is an essential aspect of art). On the other hand, narratives/literature have a social effect. This means that the concept of a narrative may originate and primarily exist in the realm of art but can, as an intellectual object, be used outside art and, consequently, be subject to different standards than in the context of art. The question is then: Can we apply the methods of literary criticism to narratives (an object initially in art) but use it outside art?

Through the investigations of philosophers such as Roland Barthes, we know about the reader's role. One may or may not share the positions expressed in [23]; taking the reader's role seriously means that there is a responsibility of the author(s) and one of the reader. From a (or at least some) modern point of view, the text exists for its own, and the reader may well go beyond what the author had in mind. Whereas this "going beyond original intentions" - approach is most fruitful in, say, reading poetry, it is more problematic in model interpretation. If modelling results are embedded into narratives, the storytellers and the reader must be aware that what they read is a narrative, but the interpretation is not as free as in the case of a pure work of art. Whereas it is appropriate to take a piece of art as inspiration for own thoughts, ideas and emotions, a narrative for simulation results must be regarded as a vehicle of content. In reading such a narration, the reader is required to scrutinize her or his interpretations and try reading the text in the author's sense. The authors have the obligation to make this sense transparent.

The presentation of the results of a simulation in a scientific manner (means as tables and graphs) has its advantages insofar as the potential for misunderstandings is reduced. If modelling and simulation results are embedded into narratives, the reader or user of the simulation results shares the responsibility for correct reading and interpretation. That means, for example, that users are responsible for knowing the difference between a scenario and a prediction. On a somewhat higher level, users must be aware that the quality of the input data determines the quality of the output data. In general, users of modelling results must understand their role not only as a receiving one but as, in various aspects, a critical part of modelling ("critical" in the sense of "important", but also in the sense of offering critique to the modellers). This holds for individual readers but even more so for the media.

6 Concluding Remarks

In the introduction, we formulated three requests for responsible modelling. In the paper, we focussed the discussion on computational modelling.

The first request addresses the users of models. Doing science and acting in the scientific community requires acquaintance with and acceptance of a specific set of social practices. If one acts in an interdisciplinary context but still within science, one can build on these practices. As soon as one enters decision-support, the involved partner may add different boundary conditions or potentialities that alter the extent and range of ethical considerations. The discussions in Sec. 2-5 contribute all to this discussion.

The second request, R2, postulates that one needs a process-oriented understanding of modelling. This holds for the development of models, discussed in Sec. 2. The interplay between model builders and experimentalists poses challenges for project management. In our view, it is crucial to understand that the existence of these challenges is not "a bug, but a feature." Of course, experimentalists and modellers know about these difficulties. To overcome them is a part of the scientific process but needs some organisation; we need interfaces between different branches. It is relevant that one must not take the existence of these interfaces for granted, but their construction can be part of a project. The critical aspect is the relation between model variables and measurable observables. The potential tension between model variables and observables becomes even more critical if one works not only in an inter- but in a transdisciplinary context. Besides the fact that one can no longer rely on well-established conventions of scientific argumentation and practices, questions of language and communication become demanding tasks that may require using narratives as tools for conveying content. In Sec. 5, we discussed opportunities and potential problems one faces when working with narratives. It is crucial to realise an essential distinction between art and the natural sciences if one does so. Art inspires and conveys a mood, whereas the natural sciences explain and transport facts. The boundary between art and natural sciences is, at a closer inspection, quite blurry. Nevertheless, if one uses storytelling methods as part of complex decision support processes, one must keep in mind the different objectives of art and science.

Finally, request R3 asks to clarify the responsibility of the reader. The systemic nature of the use of models requires considering the different types of users, as discussed in Sec. 2 - 4. In Sec 5, we embedded these arguments in a discussion that emphasises the role of communication. Notably, the "responsibility of the reader" is a concept that, in the context of computational modelling, does not only apply to individuals but should also be extended to institutions, especially the media. Fortunately, responsible journalists (there are still many!) must criticise the results of computational models and develop an understanding of what a model can do.

As stated in the introduction, computational modelling probably faces historical opportunities. There is some loud mistrust of science, but it is a minority of the population that expresses it. The possibilities of data acquisition, computer technology, and a growing fundamental understanding of modelling offer the modellers the chance to have a tangible impact on society in fields ranging from personalised medicine over epidemiology and economic planning to climate change. However, the chance for impact brings the duty for responsible and ethical action. Acting responsible starts undoubtedly at the level of the individual scientist. In addition, we must implement social and administrative structures that allow the ethical use of computational models and actively promote them. In our view, it is crucial to recognise that such promotion must observe the lessons of transdisciplinarity and activate resources ranging from pure natural science over philosophy and cultural studies to art and politics.

Acknowledgements

This report brings together expertise from different research directions and projects. The coordination of efforts and the compilation of this report has been made possible by the project ACDC (Artificial Cells with Distributed Cores to Decipher Protein Function) that has been funded by the European Union's Horizon 2020 research and innovation programme under Grant Agreement No. 824060. Part of the presented context has been produced as contribution to the ACDC Summer School organised by the University of Trento and the Museo delle Scienze. The contributors to this report gathered at the WIVACE 2021 meeting in Winterthur, partially sponsored by the School of Engineering, Zurich University of Applied Sciences. Contributing expertise came from projects funded by the FOPH (Federal Office of Public Health) in Switzerland and Armasuisse.

References

1. Fleischmann, K. and W. Wallace, *Ethical implications of computational modeling.* The Bridge, 2017. **41**(1): p. 45-51.
2. Wallace, W.A., *Ethics in modeling*. 1994: Emerald Group Publishing Ltd. .
3. Shults, F.L., W.J. Wildman, and V. Dignum. *The ethics of computer modeling and simulation*. in *2018 Winter simulation conference (WSC)*. 2018. IEEE.
4. Noorman, M., *Computing and Moral Responsibility*, in *Stanford Encyclopedia of Philosophy (Spring 2020 Edition)*, E.N. Zalta, Editor. 2020.
5. Fellermann, H., et al. *The 2019 Conference on Artificial Life (full proceedings pdf)*. in *Artificial Life Conference Proceedings*. 2019. MIT Press.
6. Pfeifer, R. and C. Scheier, *Understanding intelligence*. 2001: MIT press.
7. Anderson, C., *The end of theory: The data deluge makes the scientific method obsolete.* Wired magazine, 2008. **16**(7): p. 16-07.
8. Domingos, P., *The Master Algorithm: How the Quest for the Ultimate Learning Machine Will Remake Our World* 2015: Penguin. 322.
9. Gunning, D., et al., *XAI—Explainable artificial intelligence.* Science Robotics, 2019. **4**(37): p. eaay7120.
10. Ghassemi, M., L. Oakden-Rayner, and A.L. Beam, *The false hope of current approaches to explainable artificial intelligence in health care.* The Lancet Digital Health, 2021. **3**(11): p. e745-e750.
11. Stauffer, D., *A Biased Review of Sociophysics.* Journal of Statistical Physics, 2013. **151**(1): p. 9-20.
12. Wigner, E.P., *The unreasonable effectiveness of mathematics in the natural sciences. Richard Courant lecture in mathematical sciences delivered at*

New York University, May 11, 1959. Communications on Pure and Applied Mathematics, 1960. **13**: p. 1-14.

13. Wolf, W., *The Ising model and real magnetic materials.* Brazilian Journal of Physics, 2000. **30**(4): p. 794-810.
14. Gnanvi, J., et al., *On the reliability of predictions on Covid-19 dynamics: A systematic and critical review of modelling techniques.* Infectious Disease Modelling, 2021.
15. Frigg, R. and S. Hartman, *Models in Science*, in *Stanford Encyclopedia of Philosophy (Spring 2020 Edition)*, E.N. Zalta, Editor. 2020.
16. Tress, G., B. Tress, and G. Fry, *Clarifying integrative research concepts in landscape ecology.* Landscape Ecology, 2005. **20**(4): p. 479-493.
17. Klein, J.T., *A taxonomy of interdisciplinarity.* The Oxford handbook of interdisciplinarity, 2010. **15**: p. 15-30.
18. Fuchs, H.U., *From Stories to Scientific Models and Back: Narrative framing in modern macroscopic physics.* International Journal of Science Education, 2015. **37**(5-6): p. 934-957.
19. Fuchs, H.U., et al., *How metaphor and narrative interact in stories of forces of nature.* Narrative and Metaphor in Education. Look both ways. London: Routledge, 2018.
20. Kenny, G., *What Are Your KPIs Really Measuring?* Harvard Business Manager, 2020.
21. Bak, M.A., *Computing Fairness: Ethics of Modeling and Simulation in Public Health.* SIMULATION, 2020: p. 0037549720932656.
22. Krütli, P., et al., *Prioritätenliste und Kontingentberechnung: Pandemievorbereitung in der Schweiz.* 2018.
23. Barthes, R., *La mort de l'auteur.* Manteia, 1968. **5**: p. 12-17.