Synthese (2024) 203:105
https://doi.org/10.1007/s11229-024-04518-x

EDITORIAL

Surrogative reasoning in the sciences

Rawad El Skaf[1] · Laura Felline[2] · Patricia Palacios[3] · Giovanni Valente[1]

Surrogative reasoning is an effective form of scientific thinking that is commonly adopted in various fields of science, ranging from high energy physics to medicine, artificial intelligence, and robotics. The importance of surrogate reasoning is that it enables one to make inferences that would otherwise not be feasible to reach with other methods. In fact, faced with epistemic limitations concerning the target systems under investigation, researchers often resort to other directly accessible, or even just more tractable, *surrogate systems*, from which they hope to gain information about the purported target. A prototypical example of surrogative thinking is the use of mouse models to make inferences about human diseases in medical sciences; likewise, another important example in physics is the experimentation on fluid analogues to learn about the essential properties of black holes.

The expression "surrogative reasoning" was first introduced in the philosophical literature by Swoyer (1991) to refer to the inferential process that enables one to draw conclusions about a given system from a theoretical model built to represent it. Yet, we can meaningfully extend the scope of this expression beyond scientific modelling, so as to encompass all those forms of reasoning whereby one makes inferences about the target on the basis of some surrogate system, like in the case of analogical thinking, or even by positing imaginary scenarios and unfolding virtual processes, as it happens with thought experiments and computer simulations, respectively. Adopting such epistemic tools is customary in scientific practice, and one may go as far as claiming that, under certain circumstances, they are possibly the only available method to generate hypothesis and formulate predictions about real systems. That is true especially when the character of the latter precludes direct or full empirical access, for instance because they are too far away (e.g. black holes), too small (e.g. elementary particles), too complex (e.g. the earth climate), too expensive to construct (e.g. real-scaled bridges), unethical or dangerous to experiment upon (e.g. pharmacology and human medicine).

✉ Rawad El Skaf
rawad.elskaf@polimi.it

1 Department of Mathematics, Politecnico di Milano, Milan, Italy

2 Instituto Marconi Lussu, San Gavino, Italy

3 University of Salzburg, Salzburg, Austria

Published online: 25 March 2024

Notwithstanding the widespread use of surrogative reasoning in the sciences, there is ongoing philosophical debate about its very foundations. First and foremost, it is unclear what, if anything, makes it legitimate to replace the target system with a surrogate, even just for inferential purposes. Arguably, the surrogate system should share some similar properties with the target, but it is to be determined, possibly on a case-by-case basis, what properties are supposed to be relevant as well as when the alleged degrees of similarity prove to be sufficient. Hence, a relation of similarity between the target and the surrogate, however it is established, may not yield enough ground to draw reliable conclusions. Moreover, another outstanding problem is that, insofar as the target system is empirically inaccessible, one may even lack the ability to check whether or not the conclusion of a surrogative inference is actually correct. As a consequence, one can hardly know to what extent scientific hypotheses generated by surrogative reasoning ought to be trusted at all, the more so given that there is no guarantee that a fact holds true for a system if it is just observed for its surrogate counterpart. The issue when an inference performed via surrogative reasoning is valid thus remains unsettled. Contessa (2007), in contrast with the inferential conception of scientific representation advocated by Suárez (2004), contends that surrogative inferences can be valid even though they are not sound (in the sense of their conclusion not being true for the target), at least insofar as the surrogate system provides its user with an epistemic representation of the target. However, what should constitute exact necessary and sufficient conditions for validity, and in particular whether they could be established independently from the epistemic state of individual subjects, is still a rather controversial matter. In addition to these problems, there are also specific philosophical issues affecting all forms of surrogative reasoning, which deserve to be treated separately.

The aim of the present Topical Collection is to critically survey the status of surrogative thinking as it is discussed in the current philosophical literature, by examining how it is effectively applied in concrete instances arising in various fields of science, including physics, biology, medicine, artificial intelligence, and archaeology. In our view, besides having philosophical value on its own, a case-by-case analysis can aid us to shed light onto the general problem of validity of surrogative inferences. In particular, our Topical Collection features original contributions focussing on scientific modelling, analogy, thought experiments and computer simulations, whose authors are experts on relevant sub-disciplines of philosophy of science. Let us comment in greater detail on such diverse forms of surrogative reasoning here below, and then conclude by describing the content of the included articles.

1 Scientific models

As Swoyer (1991) himself pointed out, scientific modelling is the prototypical form of surrogative reasoning. In this case, the role of the surrogate is played by a model, conceived as a purpose tailored entity, distinguished from the target system one wishes to investigate, that can be of different types: indeed, there is a large variety of scientific models, ranging from theoretical constructions and purely fictional objects to collections of phenomenological data and even other real systems (cfr. Frigg & Hartmann, 2020 for a detailed classification). It is commonly held that a model is meant to represent its target, or at least some its salient aspects. As such, the literature on scientific modelling is closely tied to that on scientific representation, as it is well explained by Frigg and Nguyen (2020). Among the many philosophical issues arising in that context, the interesting one here is what goes under the name of "Problem of Surrogative Reasoning", which concerns how to afford predictions and explanations about the target on the basis of the model chosen to represent it. Presumably, providing an adequate representation should be a sufficient condition for the validity of surrogative inferences: for, it seems that, if the conclusions drawn about the target turn out to be invalid, then that would show that the model is not able to represent it in any adequate manner. Even so, though, the problem of surrogative reasoning just gets shifted to that of providing an adequate representation, and the jury is still out concerning such a problem. For instance, according to the similarity conception (Weisberg, 2013), a model adequately represents a system just in case they are similar in relevant respects to some appropriate degrees. Yet, several objections have been levelled against this view. For one, it faces the already mentioned complication of identifying the relevant properties with respect to which model and target ought to be compared, which practically reduces to a context-dependent matter (Teller, 2001). Moreover, different available metrics tend to disagree on how to quantify the degrees of similarity thereby introducing an ambiguity in the formal treatment. On top of that, even if one were to settle on a suitable, well-defined similarity measure, there still remains a great deal of arbitrariness as to what would be the threshold value that would guarantee a representation to be adequate. Finally, it should be noted that from a structural point of view the relation of similarity is symmetrical, whereas the notion of representation goes in a one-way direction: that is, while model and target are mutually similar, it is the model that is supposed to represent the target but not vice versa (Goodman, 1972). It thus follows that representation via similarity, at least in its standard account, can hardly offer a definite solution to the problem of surrogative reasoning in the context of scientific modelling. That leaves the question open whether a more refined version of the similarity conception, or perhaps some other alternative approach like representation-as (Goodman, 1976; Elgin, 2010 as well as Frigg & Nguyen, 2020 for the DEKI account) can offer a more satisfactory solution.

2 Analogical reasoning

Analogy is a more specific form of inferential process of the surrogative kind, which is fruitfully used in scientific practice to formulate hypotheses, especially in domains where we have limited or no access to the system under investigation. The aim of analogical inferences is to derive conclusions about the unknown properties of a target, based on its similarities with a known source system, which plays the role of the surrogate. The argumentative scheme goes as follows: given a number of properties P(1), P(2),…,P(n) of the source being similar to the corresponding properties P*(1), P*(2),…, P*(n) of the target ("positive analogies"), if the former has an additional property P(n + 1), then the latter has a similar property P*(n+1) too. Of course, this purported conclusion does not follow with certainty, since it is not a deductive consequence of the premises, no matter how large the number of positive analogies may be. So, the best one can hope for is to derive plausible hypothetical conclusions. Yet, establishing the conditions of validity for analogical inferences appears rather problematic for a number of reasons, and in fact it is an outstanding issue of philosophy of science (Bartha, 2010). First of all, one ought to qualify the extent to which the properties of the target and the source are supposed to be similar. Secondly, the two systems cannot share all properties, else they would be identical: hence, one needs to explain why their dissimilarities (or "negative analogies") can, and should, be disregarded. Finally, for an inference to be licensed, it is expected that the n positive analogies be somehow connected to the additional n+1 property, but it is subject to debate what kind of connection could possibly convey plausibility to the formulated hypothesis. In order to delve into this point, let us recall a standard distinction introduced by Hesse (1966) between "material analogies", which are based on empirically observable properties, and "formal analogies", which establish mere structural isomorphisms between the properties of the target and the source. Hesse argued that a plausible inference should hinge upon material analogies that exhibit a causal connection across the relevant properties. The problem with her proposal, though, is that there are instances of successful analogies that do not depend on any underlying causal structure, but rather bear on statistical similarities. What is more, many analogical arguments, especially in physics, involve formal analogies: a typical example is the development of electricity after Newton, with the hypothesis that opposite electrical charges obey an inverse square law similar to that holding between distant masses in the theory of gravitation. To capture this fact, another proposal tracing back to Hempel (1965) advocates the notion of "nomic isomorphism", whereby the laws of two theories describing the target and source systems, respectively, have the same mathematical form, while the physical interpretation of the symbols differs in each theory. So, the driver of an analogical inference is clearly identified with the common nomic structure. Nevertheless, restricting to the formalism of the laws fails to take into account the fact that science does avail itself of empirically grounded similarities. All in all, in spite of their respective virtues, both alternative proposals to rely solely on material or on formal analogies have deep limitations. More recent approaches thus attempt to find a proper balance between these two types of analogies (see Bartha, 2019 for a review). Yet, it remains an unsettled issue whether, and how, a general scheme for analogical reasoning can be elaborated at all. In fact, authors such as Norton (2021) even go as far as claiming that an analysis of scientific analogies can only be made locally on a case-by-case basis.

3 Thought experiments

Thought experiments involve hypothetical or even counterfactual scenarios conceived to investigate specific features of reality, for instance by testing the consequences of a scientific hypothesis or even a full theory in a way that would not be materially possible, or at least convenient, to do. As El Skaf and Stuart (forthcoming) observed, just like scientific models also thought experiments raise ontological, semantic and epistemic issues. Concerning ontology we find several proposals in the literature.

That is, thought experiments are conceived differently as mere picturesque arguments (Norton, 1991), as narratives that enables readers to construct, and reason upon, mental models (Miščević, 1992; Nersessian, 1992), as limiting cases of actual experiments (Buzzoni, 2008; Sorensen, 1992; Stuart, 2016), or as real-world props that prescribe imaginings in a game of make believe (Meynell, 2014; Salis & Frigg, 2020). However that may be, the next issue from the point of view of semantics is whether, and in what sense, thought experiments can provide a representation of target systems. This question is much less addressed, with the notable exception is Nersessian (1992) who argues that it is just the ability of a mental model to represent some real-world phenomena, typically via a relation of structural similarity, that does the justificatory work in thought experiments. A long-standing discussion, instead, has surrounded the issue of epistemology. As it is commonly held that thought experiments give us new insight about the actual world, either directly or indirectly, by revealing something novel about our scientific theories, the question arises as to why are we at all justified to accept such new insight. In other words, the epistemic puzzle is to understand how we can gain knowledge about the world without acquiring empirical data, but only by entertaining a hypothetical or counterfactual scenario. Several different solutions have been proposed in the literature, of which we can only mention a few here. For instance, Brown (1991 [2011]) submits that a kind of thought experiments referred to as platonic provide us with a priori, non-empirical access to the laws of nature, by some sort of intuition or seeing in the mind's eye. On the other hand, Norton (1991) claims that thought experiments can always be reconstructed as deductive or inductive arguments, whereas defenders of the mental models account (e.g., Miščević, 1992; Nersessian, 1992) reject the view that the justificatory power of thought experiments can be reduced to the logical structure of their propositional content, so as to eliminate experimental details as irrelevant. El Skaf (2021), instead, argues that thought experiments have a two-fold function of revealing and resolving inconsistencies, with different epistemic force: that is, while the revelation of an inconsistency could be analysed as conclusive knowledge, its resolution is only conjectural. All in all, the philosophical debate on thought experiments, especially concerning their nature and epistemic value, appears far from being settled.

4 Computer simulations

The last tool of surrogative reasoning we wish to review here is that of computer simulations. As a matter of fact, the philosophical literature on computer simulations tackles similar issues as those raised for thought experiments: in fact, these two epistemic devices have been explicitly compared by a number of authors (Arcangeli, 2018; Chandrasekharan et al., 2013; Di Paolo et al., 2000; El Skaf & Imbert, 2013; Lenhard, 2018; Shinod, 2021). From an ontological perspective, it should be noted that there are various types of computer simulations, e.g. equation-based, agent-based, multi-scale and Monte Carlo, and while there is no univocal definition of this concept (cf. Winsberg, 2019), the core idea is to unfold virtual processes by feeding a model into a computer, whose ability to perform huge calculations allows one to efficiently treat the complexity of the mathematical equations, even just in terms of finding approximate solutions. As the objects featuring in the simulations are meant to be virtual counterparts of target systems, they play the role of surrogates aiming to provide indirect information about real-world behaviours. Thus, beside serving as convenient devices that facilitate our calculations, in the absence of empirical access to the target, whether because it is not available at the present time (e.g., the future climate) or because it is impractical to experiment upon (e.g. a bridge or some other construction), computer simulations have an essential epistemic function to enable predictions and explanations of real processes that would otherwise be impossible. That can in turn have an impact on important decisions to be taken in terms of policy-making, especially in scientific fields with socio-political applications. So, given that very often, due to the opacity of virtual processes, we are not even in a position to check directly whether they properly reproduce (actual or potential) real-world processes, the pressing semantic and epistemological questions to be addressed are the following: do computer simulations adequately represent their target? and, if that is not the case, why should we be justified to trust their outcomes in the first place? In the attempt to answer such questions, the philosophical literature gets divided between authors that aim at analysing computer simulations *per se*, regardless of their similarities or differences with other tools of surrogative reasoning, and authors that aim at reducing, in some sense or another, their analysis to scientific models, or to thought experiments, or even to laboratory experiments (both direct and analogue). Specifically, Frigg and Reiss (2009) claim that computer simulations cannot possibly raise new epistemic, semantic, ontological and methodological questions beyond issues already raised by more familiar tools (Frigg & Reiss, 2009); to the contrary, other philosophers of science (e.g., Humphreys, 2009, Winsberg, 2019) submitted that computer simulations do require a novel epistemology, which ought to be further articulated.

5 Summary of the contributions

As our discussion here above illustrates, all specific forms of surrogative reasoning are affected by unsettled philosophical issues. Indeed, the problem of determining the conditions of validity of surrogative inferences, through whatever form they are drawn, remains outstanding. The present Topical Collection purports to be an up-to-date contribution to the ongoing debate, by including novel articles that deal with case studies from different fields of science. Before proceeding to review each article in greater detail, we wish to thank the authors for their careful work, some of which was presented and discussed in dedicated workshops and lectures in the Institutions where the guest editors are based, namely Politecnico di Milano and University of Salzburg. Our gratitude also goes to the Editors of *Synthese* for accepting to host our Topical Collection and for kindly offering constant support during the entire editorial process.

Let us begin with a contribution on scientific modelling, namely “Confidence in Covid-19 Models” by James Nguyen (Stockholm University), where the author focuses on the epistemological models used to inform policy makers in the context of the COVID-19 pandemic. In this analysis, he stresses the idealized character of such models, which include agent-based simulations, and explains how this can cause important inaccuracies in the projections coming from the results of the models.

According to Nguyen, an outstanding problem that frequently arises when idealized models, such as Covid-19 models, are used to guide policy interventions has to do with a lack of transparency about the model-target distinction and the uncertainty associated to it. He then suggests two different strategies to cope with the said problem that are largely based on the fifth IPCC report in climate sciences: that is, including confidence judgements to each projection scenario and taking into account experts' judgement about how to "translate" model results into realistic projections. Although the author acknowledges that these strategies may introduce non-mechanical elements into the process, he argues that this does not necessarily threatens objectivity.

Another relevant paper is "Technology-Driven Surrogates and the Perils of Epistemic Misalignment: An Analysis from Contemporary Microbiome Science" by Javier Suárez (University of Oviedo) and Federico Boem (University of Twente), which examines specific case studies in microbiome science, such as the use of 16S rRNA sequencing. The authors maintain that the successful use of surrogate reasoning in this discipline depends on certain features of the system that is used as surrogate: in particular, an adequate choice of such features enables one to avoid a phenomenon sometimes observed within the scientific community, called epistemic misalignment, in which the information provided by the putative surrogate fails to confirm a hypothesis at stake in spite of apparent positive outcomes. Although in the present paper Suárez and Boem focus only on case studies in microbiome science, which is characterized by a combination of tools and methods coming from different fields, they contend that this conclusion can be extended to all cases in which technologically-driven surrogates are used for hypothesis testing. The importance of this result is that it appears to challenge the widely spread idea in philosophy of science that what grounds surrogate reasoning is principally the user's intention.

Next, we review two articles engaging with analogical thinking. The first one is John Norton's (University of Pittsburgh) "How analogy helped create the new science of thermodynamics", which discusses a case study from the history of physics. That is Sadi Carnot's ground-breaking work on the efficiency of heat engines, which allegedly took inspiration from the theory of ordinary machines previously elaborated by his father Lazare. In Norton's view, this example is meant to illustrate his own deflationary account of analogy, according to which there is no universal scheme for analogical reasoning, but acceptable inferences can only be licenced by the truth of some "facts of analogy" that are common between the target and the source systems. The heuristic analogy between Lazare's machines and Sadi's engines is warranted by the fact that they are both dissipative and lead to maximal efficiency in a sufficiently similar manner: for the former one requires that all shocks and percussion be minimized, and likewise for the latter one requires that all temperature differences between contiguous bodies be minimized. However, in Norton's account, when drawing the intended analogical inferences, one should remain alerted that target and source are still dissimilar in important ways. Indeed, in the historical example at stake, the conditions for maximal efficiency are satisfied just in the idealized case of reversible processes, which for Norton introduces a crucial disanalogy between machines and engines, in that such processes are compatible with the laws of mechanics, whereas they contradict the laws of thermodynamics.

The other article, titled "Evidence and analogy in Archaeoastronomy", that advocates a local approach to analogical inferences comes from Francesco Nappo, Giulio Magli and Giovanni Valente (Politecnico di Milano). These authors address an area of science virtually unexplored in the philosophical literature, namely archaeoastronomy, which is an interdisciplinary field that aims to reveal mysterious connections of archaeological sites with observable astronomical phenomena. As we lack direct information about ancient civilizations from the remote past, analogical reasoning is one of the most effective methods to formulate hypotheses about the original intentions behind their monuments. For instance, if the orientation of some building tracing back to an unknown civilization is such that illumination from natural light is maximal during a particular day, say a solstice, then by analogy with other temples having the same orientation but constructed by some more familiar civilization one is tempted to infer that even the target building was originally used for similar divination purposes. The author distinguish and classify different instances of analogical reasoning in concrete case studies of archaeoastronomy, and then proceed to evaluate the epistemic value of the alleged evidence supporting the relevant inferential processes.

Thought experiments occupy a central role in black hole physics, yet their epistemic function remains understudied in the philosophy of science. In their contribution titled "What can we learn (and not learn) from thought experiments in black hole thermodynamics", Rawad El Skaf (Politecnico di Milano) and Patricia Palacios (University of Salzburg) start filling this gap by focusing on two examples from black hole thermodynamics, namely Wheeler's Thought Experiment and Geroch's Engine. According to the authors, the main epistemic role of these thought experiments is to give us conclusive information about contradictory statements coming from different theories. In particular, Wheeler's Thought Experiment shows a tension between the no-hair theorem in general relativity, quantum mechanics and the second law of thermodynamics, which Bekenstein's (1972) attempts to assuage by introducing what is called Generalised Second Law. In turn, as Geroch's Engine reveals, such a law appears in contradiction with other statements of general relativity associated with the existence of an event horizon. El Skaf and Palacios then argue that thought experiments can also help us visualize ways of solving the revealed inconsistencies. Yet, a more conclusive resolution to these inconsistencies requires one to appeal to other epistemic tools, such as direct experimentation, analogue experiments or mathematical derivations.

The next article on the topic is "Putting the 'experiment' back to into the Thought Experiment'" by Lorenzo Sartori (London School of Economics). Here, the author tackles the radical disagreement between different accounts of thought experiments found in the literature, and he argues that it could be remedied by reintegrating the experimental component into the notion of thought experiments. The proposed move is done within a novel analysis of thought experiments as representations, which rests upon the distinction between internal and external validity: specifically, in Sartori's account, the former is best analysed by Walton's pretence theory, whereas the latter is best defined as accurate representation. To be sure, as he explicitly acknowledges, this proposal does not have the resources to fully assess if a given thought experiment is externally valid, since it does not tell us how we can determine whether it accurately represent something specific. Yet, Sartori defends that this is not really a problem for his own account: in fact, the question whether or not we are justified to extrapolate the outcomes of a thought experiment to its intended target partly goes beyond the scope of the thought experiment itself, just like the same question goes beyond the scope of models and laboratory experiments.

As for the topic of computer simulations, in their contribution "Computer simulations and surrogative reasoning for the design of new robots" Edoardo Datteri (Università di Milano-Bicocca) and Viola Schiaffonati (Politecnico di Milano) distinguish two kinds of simulation-supported surrogative reasoning, namely model-oriented and prediction-oriented, which concern the activity of an agent testing hypotheses about the behaviour of the target system (under fixed conditions) on the basis of a simulation system. Their difference lies in the fact that in the former the hypothesis is that the target's behaviour can be reproduced by a certain mechanism, whereas in the latter the hypothesis is that the target's behaviour will be of a certain kind. Armed with such a distinction, the authors go on to show that, in the context of robotics, designing processes tend to employ computer simulations *qua* concrete models to support model-oriented reasoning on the building of new robots. This analysis is then applied to van Eck's (2016) account of the counterfactual use of engineering theoretical models, so as to offer a reconstruction of computer simulations that can potentially aid to identify factors responsible for unreliable predictions.

The final contribution to our Topical Collection is the paper titled "Sharpening the tools of Imagination", wherein the author Mike Stuart (University of York) identifies several forms of surrogative reasoning such as thought experiments, models, diagrams, computer simulations, and metaphors, as "tools of the imagination". His main contention is that, since these tools prompt or focus our imagination, they enable us to imagine in a more useful manner. In order to enforce this point, instead of insisting on the traditional relation of representation between the surrogate introduced by a given tool and its target, Stuart adopts an adequacy-for-purpose approach with a consequentialist epistemology: a tool of the imagination is effective in prompting or focusing the imagination if that tool, and correspondingly the imagination, had, has, or will most probably have good (epistemic) consequences. The intended effectiveness is sensitive to the different contexts of application, depending on whether a given tool (i) already has known consequences (when we are analysing a historical case study), (ii) has unknown but foreseeable direct consequences (when we are analysing a case study at the cutting edge of science) or (iii) has unknown but foreseeable indirect consequences (when scientists don't have a particular problem in mind). It is worth stressing that Stuart's analysis is grounded on an empirical investigation with qualitative methods on how scientists themselves evaluate the effectiveness of imaginings in these different contexts.

References

Arcangeli, M. (2018). The hidden links between real, thought and numerical experiments. *Croatian Journal of Philosophy, 18*(1), 3–22.

Bartha, P. (2010). *By parallel reasoning: The construction and evaluation of analogical arguments*. Oxford University Press.

Bartha, P. (2019). *Analogy and analogical reasoning*. Stanford University Press.

Bekenstein, J. D. (1972). Black holes and the second law. *Lettere al Nuovo Cimento, 4*, 737–740.

Brown, J. R. (1991 [2011]). *Laboratory of the mind: Thought experiments in the natural sciences*. London: Routledge, Second Edition.
Buzzoni, M. (2008). *Thought experiment in the natural sciences: An operational and reflexive-transcendental conception*. Würzburg: Königshausen & Neumann.
Chandrasekharan, S., Nersessian, N., & Subramanian, V. (2013). Computational modeling: Is this the end of thought experiments in science? In M. Frappier (Ed.), *Thought experiments in philosophy, science and the arts.* Routledge.
Contessa, G. (2007). Scientific representation, interpretation, and surrogative reasoning. *Synthese, 74*(1), 48–68.
Di Paolo, E. A., Noble, J., & Bullock, S. (2000). Simulation models as opaque thought experiments. In M. A. Bedau, J. S. McCaskill, N. H. Packard, & S. Rasmussen (Eds.), *Proceedings of the 7th international conference on artificial life* (pp. 497–506). MIT.
El Skaf, R. (2021). Probing theoretical statements with thought experiments. *Synthese, 199*, 6119–6147.
El Skaf, R., & Imbert, C. (2013). Unfolding in the empirical sciences: Experiments, thought experiments and computer simulations. *Synthese, 190*, 3451–3474.
El Skaf, R., & Stuart, M. T. (forthcoming). *Scientific models and thought experiments. In Routledge handbook of philosophy of scientific modelling*. Routledge. https://www.routledge.com/The-Routledge-Handbook-of-Philosophy-of-Scientific-Modeling/Knuuttila-Carrillo-Koskinen/p/book/9781032071510
Elgin, C. Z. (2010). Telling instances. In R. Frigg & M. Hunter (Eds.), *Beyond mimesis and convention: Representation in art and science*. Springer.
Frigg, R., & Hartmann, S. (2020). *Models in science*. Stanford University Press.
Frigg, R., & Nguyen, J. (2020). *Modelling nature: An opinionated introduction to scientific representation*. Springer.
Frigg, R., & Reiss, J. (2009). The philosophy of simulation: Hot new issues or same old stew. *Synthese, 169*, 593–613.
Goodman, N. (1972). Seven strictures on similarity. In N. Goodman (Ed.), *Problems and projects.* Bobs-Merril: Indianapolis.
Goodman, N. (1976). *Languages of art* (2nd ed.). Hackett.
Hempel, C. G. (1965). Aspects of scientific explanation. *Aspects of scientific explanation and other essays in the philosophy of science* (pp. 331–496). Free Press.
Hesse, M. B. (1966). *Models and analogies in science*. University of Notre Dame. Press.
Humphreys, P. (2009). The philosophical novelty of computer simulation methods. *Synthese, 169*, 615–626.
Lenhard, J. (2018). *Thought experiments and simulation experiment*. Routledge.
Meynell, L. (2014). Imagination and insight: A new account of the content of thought experiments. *Synthese, 191*(17), 4149–4168.
Miščević, N. (1992). Mental models and thought experiments. *International Studies in the Philosophy of Science, 6*(3), 215–226.
Nersessian, N. J. (1992). In the theoretician's laboratory: Thought experimenting as mental modeling. *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association, 2*, 291–301.
Norton, J. D. (1991). Thought experiments in Einstein's work. In T. Horowitz & G. Massey (Eds.), *Thought experiments in science and philosophy* (pp. 291–301). Rowman & Littlefield.
Norton, J. D. (2021). *The material theory of induction*. BSPS Open.
Salis, F., & Frigg, R. (2020). Capturing the scientific imagination. In A. Levy & P. Godfrey-Smith (Eds.), *The scientific imagination* (pp. 17–50). Oxford University Press.
Shinod, N. K. (2021). Why computer simulation cannot be an end of thought experimentation. *Journal for General Philosophy of Science., 52*, 431–453.
Sorensen, R. (1992). *Thought experiments*. Oxford University Press.
Stuart, M. T. (2016). Taming theory with thought experiments: Understanding and scientific progress. *Studies in History and Philosophy of Science Part A, 58*, 24–33.
Suárez, M. (2004). An inferential conception of scientific representation. *Philosophy of Science, 71*(5), 767–779.
Swoyer, C. (1991). Structural representation and *surrogative reasoning*. *Synthese, 87*, 449–508.
Teller, P. (2001). Twilight of the perfect model model. *Erkenntnis, 55*, 393–415.
Van Eck, D. (2016). *The philosophy of science and engineering design*. Springer.
Weisberg, M. (2013). *Simulation and similarity: Using models to understand the world*. Oxford University Press.
Winsberg, E. (2019). *Computer simulations in science*. Stanford University Press.