Why We Love Pictures (for the Wrong Reasons)

A lesson from the picture of a black hole

Lorenzo Sartori

September 12, 2025

Paper accepted to the 2024 edition of the Philosophy of Science Association (New Orleans) and for publication in the related special issue in *Philosophy of science*. This version of the article has been accepted for publication after peer review but is not the Version of Record and does not reflect post-acceptance improvements, or any corrections.

Abstract

In this paper, I first show that similarity accounts of scientific pictures fail with more realistic cases of scientific pictures. My primary case study is the picture of a black hole, from which I develop an interpretation-based account of picture representation analogous to how models represent: a picture represents a designated target system iff, once interpreted, it exemplifies properties that are then imputed to the target via a de-idealising function. Then, I show that justification of the inferences from pictures crucially depends on their causal mechanisms of production, in contrast with the standard justificatory strategies we employ for model inferences.

Wordcount: 4700 words (bibliography excluded).

Keywords: pictures, scientific representation, justification, black hole

Contents

1 Introduction

Pictures are ubiquitous in science. Astronomers analyse telescope images to study stellar phenomena, medics use X-ray and MRI scans for diagnosis, and epidemiologists employ heatmaps to predict disease spread. This raises a question: how do pictures convey knowledge about the world they represent?

In section 2, I examine Meynell's (2013) similarity-based account of pictorial representation and argue it is inadequate, focusing on the black hole M87* image. Section 3 provides context on how this picture was produced. Section 4 develops an interpretation-based account of this image, showing it aligns with the DEKI account of representation (Frigg and Nguyen 2020). Finally, section 5 explores how epistemic justification differs for pictures like M87* compared to models, emphasising the role of production history in the former.

2 The mirage of similarity

In the literature on depiction, both in aesthetics and the theory of images, there is a long-standing tradition that focuses on similarity to explain how pictures represent the real world (Wollheim 1987, Hyman 2006, 2012, Peacocke 1987, and Gombrich 1960, 1982).[1] However, applying the similarity view to scientific pictures has been largely unexplored by its proponents in the context of scientific models (e.g., Giere 2004, 2010, and Weisberg 2013, Ch. 8).[2] An important exception is Meynell (2013), who explicitly clarifies the role of similarity in scientific pictures and visual representations. I take her view as a reference for critically analysing the similarity account in this context.

Meynell's account combines a critique of Perini's (2005) Goodmanian conventionalism with a constructive proposal inspired by Willats's (1997) work in psychology. While Meynell accepts that Perini's approach may work for linguistic or schematic visual representations, she argues it is insufficient for "dense" pictures, like photographs, scans, and astronomic images. She proposes instead a similarity-based approach, informed by psychology, perception theory, and geometry.

Following Willats, Meynell (2013, 338) argues for a two-step relationship between a picture and its target. First, picture primitives (e.g., lines, points, and coloured areas) are associated with scene primitives, which are the basic shape elements in the scene—3D (lumps, sticks, slabs), 2D (surfaces), 1D (edges), or 0D (corners). Second, these scene elements are connected to the target system in the real world.

Meynell then characterises both the relation between picture primitives and scene primitives, and the relation between scene primitives and the target. The aim is to show that both steps are based on (objective) similarity and how our perception works, in contrast with the position defended by Perini and Goodman, who focus on interpretation and convention. Both steps, she claims, are based on objective similarity and perceptual mechanisms. The first step relates picture primitives to scene primitives via geometrical projection, while the second step ties the scene to the target through visual or perceptual similarities: the scene represents the target because it visually resembles what we would see if we directly observed

---

[1] As an entry point in the literature, cf. Lopes (1996).

[2] Structuralist accounts of representation, e.g., Da Costa and French (1990, 2000) and van Fraassen (2008), focus on similarities concerning mathematical structures, but are therefore less practical for pictures. Isaac (2019) adapts structuralism to but he argues that the relation of isomorphism in this context has to be understood allegorically and not literally. For these reasons, I set structuralist accounts aside here.

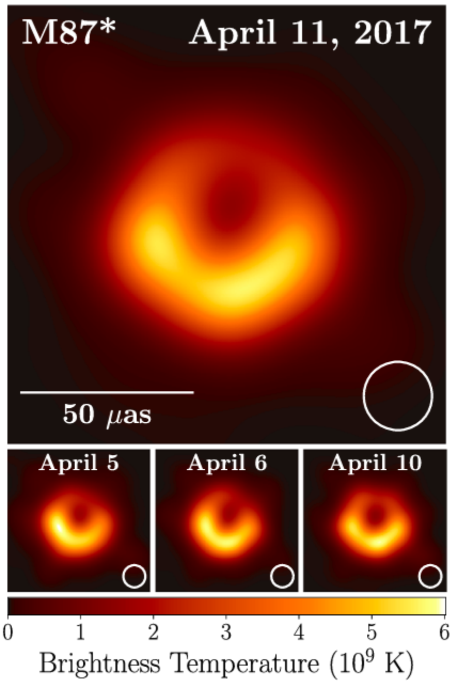

Figure 1: Picture of the black hole M87*. In EHTC (2019a, p. 5).

the target.

Assume Meynell's account succeeds for photographs, realistic paintings, and simple geometrical figures.[3] The point I want to make here is that even if the similarity view were successful for more mundane cases, the account is inadequate when applied to more complex examples of visual representations used in science. Let us then consider the picture of the M87* black hole produced by the Event Horizon Telescope (EHT).

First, according to Meynell, we should be able to reconstruct the scene represented by the picture by geometrically projecting the picture primitives (coloured regions) onto the scene's basic elements. Yet even knowing the geometrical projection (e.g., perspectival), it is unclear what the scene itself would be, as projections are neutral with respect to what is projected. Second, if similarity between the scene and target is defined in terms of visual appearance, we would infer that the black hole is a reddish-yellow doughnut in the middle of nothing—an obviously false depiction.

In contrast, what we need is a function that associates colours with what they represent—levels of radiation intensity. This requires interpreting the picture, translating visual properties into non-visual ones. Thus, on closer inspection, similarity is not the core concept for understanding how such pictures epistemically function; the critical role is played by interpretation. Once colours are translated into radiation intensity levels, there is no meaningful sense in which the picture and target are visually similar. This point generalises to other scientific pictures, as seen in examples from Elkins (1999) and Tufte (1997).

---

[3] Critics have argued that similarity is reflexive and symmetric, while representation is not (Goodman 1976, Suárez 2003). While valid, these objections have been already discussed (Giere 2004, 2010; Weisberg 2013; Frigg and Nguyen (2020, pp. 31-50) and references therein)). My critique focuses on dense visual representations, which Meynell specifically addresses.

3 One step back: imaging a black hole

There is already an emerging literature on the epistemology of the picture of M87* and in general of black holes (Skulberg 2021; Muhr 2023; Doboszewski and Lehmkuhl 2023; Doboszewski and Elder 2024). Curiously enough, however, philosophers of science have not studied this picture *as a picture*. That is, they have not focused on the features of this picture as a representation, namely as an object allowing surrogative reasoning about its target system. I want to suggest that, while the epistemological analyses conducted so far are crucial, they remain incomplete. Indeed, a study of the picture as evidence presupposes an analysis of how the picture of a black hole is supposed to be "read" as a representation. In this sense, my analysis will also be a useful contribution to the general epistemological enquiries about black hole pictures.

In order to make such an analysis of the picture of M87* as a representation, one needs to first provide some details on black holes and how we image them.

Imagine an object, far away in space (ca. 54.8 million light years from the milky way) with the mass equivalent to around 6.5 billion times the Sun but compressed so that its size is comparable to our Solar System. This physical object is what scientists think that probably lies at the centre of the Messier 87 galaxy (Gebhardt *et al.* 2011; Walsh *et al.* 2013): a supermassive black hole that the astronomers call M87*.[4]

Black holes are astronomical objects predicted by the general theory of relativity (Einstein 1915; Penrose 1965) and central to issues concerning the unification of GR with quantum physics (Hawking 1976; Giddings 20017). While the very definition of a black hole is a matter of dispute (Curiel 2019), many of the technical aspects concerning this problem are irrelevant for my analysis, which intends to work well irrespective of their exact theoretical definition.

The gravitational pull created by the black hole is so strong that, if something gets actually too close, it is irremediably swallowed into the black hole: not even light escapes (Schwarzschild 1916), and that is what gives the black hole its name. There is then a boundary beyond which even photons cannot escape, and we call the line delineating this point of no return the *event horizon* of a black hole. Nothing escapes, so nothing can be observed[5] when it is beyond the event horizon: we can only make theoretical hypotheses on what happens beyond that line.

Fortunately for us, the black hole attracts all sort of matter and energy from its surrounding. The first observational confirmations of the existence of black holes were due to the very fast and small orbits of stars around a centre of gravity where no observable object was reported (Harms *et al.* 1994). Furthermore, and more importantly for our purposes here, there is something that we can observe in the external proximity of the event horizon. There, orbits at incredible speed what is called an *accretion disk*, namely a tremendous amount of matter, mostly ionised gases, burning at a temperature ca. 1 to 10 billion degrees Kelvin. Because of its high temperature, the accretion disk irradiates many forms of radiation among which light. Most of this radiation, of course, travels at wavelengths that cannot be perceived

---

[4] Supermassive black holes distinguish themselves from the far smaller black holes originating by the implosion of a star. Supermassive black holes are thought to exist in the centres of nearly all galaxies (Lynden-Bell 1969; Kormendy and Richstone 1995; Miyoshi *et al.* 1995), including our own (Eckart and Genzel 1997; Ghez *et al.* 1998; Abuter *et al.* 2018).

[5] Here, I use the term "observation" in a technical sense, encompassing any form of measurement – it should thus not be restricted to human vision alone.

by the human eye but can still be measured by our measurement devices.

The fundamental theory to measure astronomic phenomena is interferometry.[6] An interferometer gives a measure of the intensity of radiation of an electromagnetic source by decomposing the original light in two beams and then calculating the phase differences between them. In this way, we can measure the relevant distribution of the radiation from a radiating source – in principle, even from an astronomic one like the accretion disk of the black hole M87*. These measurements are crucial, because the distribution of radiation can give us insight on the dimensions and shapes of what lies within the event horizon, which is usually called the *shadow* of a black hole.

For four days in April 2017, seven telescopes in different locations on the globe were pointed towards the centre of the Messier 87 galaxy and measured the radio signals coming from that region of spacetime. The idea was to synchronise all the telescopes so that they could be used as one single telescope. The resulting "lens" of this composite telescope, even though fragmented, had the width of the entire planet Earth. One and half petabytes of interferometric data were collected for each night of observation, that is, the greatest amount of data in the history of science for single experimental measurement. These data were then fed to a supercomputer which integrated the data of each single telescope. The data was further calibrated (the Event Horizon Telescope Collaboration (EHTC) *et al.* 2019b).

Then, four teams of researchers were created to independently produce a visual output from the data. The four teams worked autonomously, and they were not allowed to talk with each other. Two teams ended up using the so-called CLEAN algorithm, while the other two teams used two different versions of the so-called Regularized Maximum Likelihood (RML) family of algorithms: the algorithm SMILI and the algorithm EHT-imaging, the latter created specifically in the context of the Event Horizon Telescope measurement (EHTC *et al.* 2019c).

Then, the four teams compared their output images, which all exhibited two important structural features: a ring shape with more intense brightness in the south region of the black hole, and the diameter of the ring estimated around 40μas (*ibid.*, 9). Four images were then produced from each algorithm pipeline, one for each night of observation. As a further step in making these pictures more reliable, all these four images were further blurred to obtain a "common, conservative resolution" of each of them (*ibid.*, 20). Finally, to further emphasise the common features of the images produced by the three different pipelines, the scientists produced an average picture for each of the four days (*ibid.*, 21).[7]

The image that we eventually obtain from this procedure is, basically, a heatmap. A heatmap is a visual representation of data where values are represented by colours, and the spatial coordinates on the map are to be translated in other properties of the represented phenomenon. The easiest case is when the spatial properties on the map are translated in spatial properties of the target system via a geometrical projection. However, heatmaps can use colours to represent any sort of property or quantity. What we need is a way to systematically interpret the visual properties of the heatmap into the relevant properties we are actually measuring in the target system. In section 4, I draw the basic elements of such an account.

---

[6] The standard reference here is Thompson *et al.* (2017). More details about the specific methods that the EHTC employ are provided in the Science and Technology sections of the Event Horizon Telescope website.

[7] This is of course a very simple reconstruction. The reader can find all the details in the six articles published by the EHTC team reported in the bibliography.

4 An interpretation-based account of scientific pictures

The picture of M87*, I argue, interpreted as such, refers to M87*. Following Goodman (1976), I take this referential relation from a symbol to an object to be denotation, namely, the referential relation between a name and its bearer. In order to denote M87*, though, we need to look at the picture as not, say, mere pixels coloured in a specific way: we need to interpret the material instantiation, or *carrier*, of the picture, as a radiation-heatmap of the electromagnetic field surrounding M87*. Interpretation here can be understood as a function I that maps colours of the picture into levels of radiation intensity of an electromagnetic field, where black is interpreted as lower levels and yellow as higher. So, we have three elements: the carrier, the heatmap, and the target. The picture, once interpreted, represents M87* as a radiation-heatmap – or, alternatively, it is radiation-heatmap-representation of M87*.

The heatmap resulting from our interpretation of the carrier is not supposed to be read as a truthful description of the target, as it may involve idealisations and distortions with respect to the target system. Again, I use Goodman's (1976) terminology and I call the radiation-heatmap a Z-representation. Distinguishing these three elements (carrier, Z-representation, and target system) is crucial to appreciate the type of reasoning involved when we make inferences about the actual black hole on the basis of the picture. This reasoning is surrogative in nature (Swoyer 1991): reading the picture as a representation requires us to study a system in order to make inferences about another system.

It is nevertheless useful to study the Z-representation in question because it highlights certain properties of the target that we would have not been able to detect if we had just looked at the raw interferometric data collected by our telescope. Using the terminology of Goodman's (1976) and Elgin (1983, 1996) we can call this function *exemplification*. Technically, an object exemplifies a property A if it instantiates A and refers to A by instantiating it. A typical example is a swatch in a tailor shop. The swatch possesses many properties (say, being rectangular, being produced in Thailand, weighing 1.5 grams...) but, in the context of the shop where people decide how their clothes should be, the swatch refers only to a certain set of properties: colour, texture, material, and so on. By referring to them, it makes them salient: these properties become, or should become, more epistemically accessible for a user or observer.

The same happens with the radiation-heatmap: by abstracting away informational noise and resulting by an interpolation of the original sparse data, this picture summarises and highlights certain prominent features of M87*, particularly, the remarkable shadow of the black hole, its dimensions, and an asymmetry in radiation intensity between the Southern and Northern areas of the accretion disk. For example, from the scaled dimensions of the pictures, astronomers can infer that the dark object at the centre cannot be a naked singularity or a wormhole, but it is more likely to actually be a supermassive black hole as it is theoretically predicted by the general theory of relativity (Bouman 2020).

We can go further and, coupling the picture properties with our models of black holes and "information on the inclination angle, [...] derive the sense of rotation of the black hole to be in the clockwise direction, i.e., the spin of the black hole points away from us" (EHTC 2019a, p. 9). Also, the EHTC authors explain the brightness asymmetry in the South region of the black hole "as relativistic beaming of material rotating in the clockwise direction as [...] moving toward the observer" (*ibid.*). The South-North asymmetry in radiation intensity, then, is explained as a case of the so-called Doppler effect.

The point of all this is of course to eventually impute some of these exemplified properties to the target system, the actual black hole M87*. When I talk about imputation, I simply mean property attribution, with no assumptions about whether this attribution is correct or not.

Sometimes, the properties exemplified by a representation can be imputed to its target system unchanged, but other times this imputation will require some form of de-idealisation. To distinguish this further interpretive activity from the interpretation function I, let us call *key* the function that translates idealised properties of the representation into non-idealised properties imputed to the target.

In the case of the picture of M87*, there are many keys at work. One is simply a scale factor, multiplying the dimensions of the object in the picture into the actual dimensions that we expect M87* has, based on the picture together with our knowledge of the distance between us and the centre of the Messier galaxy. Another important key is a geometrical projection that translates the two-dimensional spatial properties of the heatmap into a set of three-dimensional ones.[8] A final interesting type of key seems in place in the process of blurring the final visual outputs of the algorithms in order to decrease precision but increase reliability. Here, a key should be used to de-blur the picture if our aim is to attribute more fine-grained, precise properties to the actual black hole.[9]

From this reconstruction, one can see that no appeal to similarity, perception or psychology has been made: given the presence of the I and the key, neither the picture as a carrier nor the picture as heatmap need to be similar to the target to represent it. At the same time, the account still retains Meynell's basic intuition about the role of geometrical projections in visual representation. This role, though, is inserted in a more general framework entirely based on interpretation (in the various fashions of the I-function, denotation, the selective process of exemplification, and the key). Interpretation here is arbitrary but not random: while heavily theory-laden, our interpretation of the picture strictly depends on the way in which the project was in fact produced.

Once the picture is interpreted as a heatmap, four basic elements have been highlighted in my analysis: denotation, exemplification, keying-up, and imputation. These are the basic ingredients of the so-called DEKI account of scientific representation (Frigg and Nguyen 2020), that takes its name exactly from those ingredients and it has been applied to many cases of scientific models. More formally, according to DEKI, a model system M (that is, a carrier C endowed with an interpretation I) is an epistemic representation of a designated target system T iff four conditions apply:

(i) M *denotes* T,

(ii) M *exemplifies* properties P_1..., P_n,

(iii) P_1..., P_n are associated with a second set of properties Q_1..., Q_n via a *key*,

(iv) Q_1..., Q_n are *imputed* to T.[10]

---

[8] This geometrical translation will have to account for the complex geometry of the spatiotemporal region under investigation, where distances in the pictures are distorted with respect to actual distances.

[9] This type of deblurring key also seems the converse of another type of keys used, for example, in modelling the climate. There, experts downgrade the likelihood of a certain result to factor in the high level of uncertainty about the model ensemble's reliability (IPCC 2010). The issues concerning how this practice can be epistemically justified are numerous and deep. See Harris (2021, pp. 245-261).

[10] The DEKI account is very complex and an exhaustive analysis of it lies outside the scope of this paper. Interested readers can find all the details in Frigg and Nguyen (2020, 159-214).

At this level of analysis, then, it seems correct to say that a mechanically produced picture like the picture of M87* functions as an epistemic representation in the same way as scientific models represent their target systems.

5 From semantics to epistemology: measurement vs. model

So far, I have focused on how to interpret the picture as a representation and characterised this representational function in the terms of the DEKI account. However, the account has two important shortcomings. First, it is skeletal by design: it needs to be completed with the specifics of each case study. The analysis of the black hole picture just offered provides the relevant details on how to apply DEKI to this specific case study. Second, the account remains silent on the justification of our inferences from the picture to the target system. This is because the account correctly acknowledges that the justificatory roots of our inferences lie outside the single representation system. In this respect, pictures are exactly like models, because even for inferences drawn from a model about a target system, the only way to justify our inferences is to support them with reasons extrinsic to the single model system (theories, observations, and other models).

The similarity between pictures and models, though, ends here. For, I want to suggest, in the case of pictures like the picture of M87* the root of the justification of both our interpretation of the picture and of our inferences about its target is the history of production that causally connects a given picture to its designated target system. And this is not usually the case with models.

Let us assume that, like in our example, it is impossible to observe the target system directly. If one wants to convince me that the picture of M87* is giving me (approximately) correct results about the actual black hole, they will have to explain to me how the picture was produced, and how the visual output we are looking at is causally dependent on the target system.

For example, Doboszewski and Elder (2024) analyse the picture in terms of robustness analysis, by showing that, first, the multiple algorithms employed for the imaging converged on similar results even if taking different procedures and assumptions. Second, the algorithms also exhibited reasonable sensitivity to different data. This second type of robustness was assessed by testing the three algorithms against synthetic images, showing different geometrical shapes, which the algorithms had to reconstruct as with the picture of the black hole. By doing this, the researchers obtained two results. First, they identified the fiducial parameters, that is, those parameters that allowed a more faithful reconstruction of the original image. Second, they proved some robustness of the algorithms by showing that they were sensitive to the input image: the outputs were really different for each synthetic image and the black hole picture, showing that there was a relatively strong counterfactual dependence of the visual output on the original source.

This two-faces robustness analysis was necessary to secure a reliable counterfactual relation between the data and the visual output obtained by applying the algorithm, and consequently, the accuracy of the latter with respect to the former.

The causal relation that I suggest lies at the root of the inferential stability[11] from pictures

---

[11] Cf. Roskies (2008) for similar considerations applied to MRI scans.

to targets more generally. The more numerous and complex are the steps in the causal chain of producing the picture from the target, the more difficult it will be to justify the inferences we draw from the former to the latter. Here, I am not arguing that the picture of M87* is, in fact, epistemically reliable. My point is just that *if* one wants to assess such reliability, one has to look at its production, and how this causally relates the target with the interpreted visual output.

Again, it is important to notice that the notions of reliability, accuracy, and success that I am employing here do well without any appeal to similarity between the picture and the target system: what counts is the counterfactual stability between the target and the visual representation endowed with an interpretation, which in turns depends on the causal mechanisms connecting them.

These considerations about the role of the justificatory role played by the causal relation between target and representation can be extended to many other scientific pictures. The justification of the inferences drawn from photographs, astronomic pictures, maps, X-ray and MRI scans all require assessment of the causal, counterfactual stability between the target and the visual output in question. Furthermore, this can be generalised to other examples of representations that we do not normally call pictures. The colour displayed on a litmus paper can be reliably interpreted as the indication of the level of acidity of a chemical solution only if one can give a story about the counterfactual stability between the chemical features of the liquid under measurement and the resulting colours appearing on the paper. The same applies to clock dials representing time, thermometers' reading of the surrounding temperatures, and even model organisms when, instead of being used to represent properties of other organisms (Sartori 2023a; see also below), are used to represent toxicity levels in the surrounding environment.[12] That is why I find useful to label all these instances of representations as "measurement representations", and not focus too much on the pictorial or imagistic features.

Nothing of the sort of what I have said about justification in the case of the picture of M87* applies to other forms of representations, like models. A model system is usually constituted by a set of assumptions on that system (an abstract object or a material one), often in interaction with each other. Let us take the simple case of an assumption that is expressed by a certain functional relation between two quantities. There are many ways in which we can justify this assumption. It may directly derive from more general theory in the relevant discipline. Or our assumption may be a simplification of a more general functional relation that however is intractable in its current form (e.g., an equation with no analytic solutions). Here, the justification follows from our reason to hold the original formula, plus some further reason to consider the simplification acceptable. Alternatively, our assumption may boil down to a hypothesis abstracted away from data, perhaps via an abductive inference.

In all these cases, the justification of the assumption will be more or less provided on the basis of previously acquired knowledge. However, the assumption could also be something completely new, detached from theory and experiments. The justification of that assumption will then solely depend on the success of the model as a whole. Success can take many forms: empirical adequacy, unification, explanation by providing an underlying mechanism. The more the model proves itself successful, the more we can justify its further application as an epistemic surrogate system. However, as it should be evident, there is no appeal to causal

---

[12] For examples of the sort, see Abdullahi *et al.* (2022) and Holsopple *et al.* (2023).

relations between the target and the model in our justification of the inferences from the latter to the former.

What I have said about models apply to other sorts of representations. Schematic diagrams, for example, seem to be justified in a similar way. Elsewhere, I have also argued that scientific thought experiments (Sartori 2023b) and model organisms (Sartori 2023a) also function as representations of phenomena in the sense expressed by the DEKI account. For both thought experiments and model organisms, I want to suggest that their style of justification is model-like.[13]

Nevertheless, it is important to remind the reader that the characterisation of the causal relation holding between target and representation in the instances of measurement-like representations is still based on theoretical assumptions and previously acquired empirical knowledge. I do not want thus to undermine the theory-ladenness of our interpretations of images. For even the assessment of the hypotheses about the causal relations in play in our production of images will strictly depend on the theoretical framework we are assuming in the first place. In a similar vein, I do not want my focus on causation here to foster the idea that pictures are somewhat more “objective” representations than models: whether an image is causally linked to its target and how accurately so can, and is, a matter of dispute, even among experts.

6 Conclusion

In this paper, I argued that similarity accounts of pictures fail with more complex scientific cases like the picture of the black hole M87*. I then offered an interpretation-based account of this picture, where the interpretation of the picture was properly anchored to the way in which the picture was produced. I show that, *qua* representation, this picture does not differ from other types of representations like scientific models. The difference, instead, concerns the justification of the inferences we draw about the designated target. In the case of pictures like M87*, the inferences can be justified only by an appeal to the causal process of production of the picture, while this is not a common justificatory strategy in the case of scientific models. In conclusion, if we love pictures, it is not because of an alleged similarity with their targets, but because of the causal mechanisms that, in many cases, grant a solid counterfactual relation between the target system and the visual output.

References

Abdullahi, M., et al. (2022). *Daphnia* as a sentinel species for environmental health protection: A perspective on biomonitoring and bioremediation of chemical pollution. *Environmental Science and Technology*, *56*, 14237–14248.

Abuter, R., et al. (2018). Detection of the gravitational redshift in the orbit of the star S2 near the Galactic centre massive black hole. *Astronomy and Astrophysics*, *615*(L15), 1–10.

Bartha, P. F. (2010). *By parallel reasoning.* Oxford University Press.

---

[13] One may want to relate this type of model-like reasoning to reasoning by analogy (Hesse 1963; Bartha 2010). However, this should not be taken as a concession to the similarity view: how analogies are drawn depend on the interpretation of both the representation system and the target system.

Bouman, K. L. (2020). The inside story of the first picture of a black hole. *IEEE Spectrum.* https://spectrum.ieee.org/the-inside-story-of-the-first-picture-of-a-black-hole

Curiel, E. (2019). The many definitions of a black hole. *Nature Astronomy*, *3*(1), 27–34.

Da Costa, N. C., & French, S. (1990). The model-theoretic approach to the philosophy of science. *Philosophy of Science*, *57*(2), 248–265.

Da Costa, N. C., & French, S. (2000). Models, theories, and structures: Thirty years on. *Philosophy of Science*, *67*(S3), S116–S127.

Doboszewski, J., & Elder, J. (2024). Robustness and the event horizon telescope: The case of the first image of M87. *arXiv preprint arXiv:2401.16323.*

Doboszewski, J., & Lehmkuhl, D. (2023). On the epistemology of observational black hole astrophysics. In N. M. Boyd et al. (Eds.), *Philosophy of astrophysics: Stars, simulations, and the struggle to determine what is out there* (pp. 225–247). Springer.

Eckart, A., & Genzel, R. (1997). Stellar proper motions in the central 0.1 pc of the Galaxy. *Monthly Notices of the Royal Astronomical Society*, *284*(3), 576–598.

Einstein, A. (1915). *Sitzungsberichte der königlich preußischen Akademie der Wissenschaften.* Deutsche Akademie der Wissenschaften zu Berlin.

Elgin, C. Z. (1983). *With reference to reference.* Hackett Publishing.

Elgin, C. Z. (1996). *Considered judgement.* Princeton University Press.

Elkins, J. (1999). *The domain of images.* Cornell University Press.

Event Horizon Telescope Collaboration, and others. (2019a). First M87 event horizon telescope results. I. The shadow of the supermassive black hole. *The Astrophysical Journal Letters*, *875*(1), L1.

Event Horizon Telescope Collaboration, and others. (2019b). First M87 event horizon telescope results. III. Data processing and calibration. *The Astrophysical Journal Letters*, *875*(1), L3.

Event Horizon Telescope Collaboration, and others. (2019c). First M87 event horizon telescope results. IV. Imaging the central supermassive black hole. *The Astrophysical Journal Letters*, *875*(1), L4.

Frigg, R., & Nguyen, J. (2020). *Modelling nature: An opinionated introduction to scientific representation.* Springer.

Gebhardt, K., et al. (2011). The black hole mass in M87 from Gemini/NIFS adaptive optics observations. *The Astrophysical Journal*, *729*(2), 1–13.

Ghez, A. M., et al. (1998). High proper-motion stars in the vicinity of Sagittarius A*: Evidence for a supermassive black hole at the center of our galaxy. *The Astrophysical Journal*, *509*(2), 678–686.

Giddings, S. (20017). Astronomical tests for quantum black hole structure. *Nature Astronomy*, *1*(0067), 1215–1220.

Giere, R. N. (2004). How models are used to represent reality. *Philosophy of Science*, *71*(5), 742–752.

Giere, R. N. (2010). An agent-based conception of models and scientific representation. *Synthese*, *172*(269), 269–281.

Gombrich, E. (1960). *Art and illusion.* Phaidon.

Gombrich, E. (1982). *The image and the eye.* Phaidon.

Goodman, N. (1976). *Languages of art.* Hackett.

Harms, R. J., et al. (1994). HST FOS spectroscopy of M87: Evidence for a disk of ionized gas around a massive black hole. *Astrophysical Journal, Part 2-Letter*, *435*(1), L35–L38.

Harris, M. (2021). *Conceptualizing uncertainty: The ipcc, model robustness and the weight of evidence* (Ph.D. Thesis). London School of Economics and Political Science.

Hawking, S. W. (1976). Breakdown of predictability in gravitational collapse. *Physical Review D*, *14*, 2460–2473.

Hesse, M. (1963). *Models and analogies in science*. Sheed; Ward.

Holsopple, J. M., et al. (2023). Assessment of chemical toxicity in adult *Drosophila melanogaster*. *JoVE (Journal of Visualized Experiments)*, *193*, e65029.

Hyman, J. (2006). *The objective eye: Color, form, and reality in the theory of art*. University of Chicago Press.

Hyman, J. (2012). Depiction. *Royal Institute of Philosophy Supplement*, *71*, 129–150.

Intemann, K. (2010). Guidance note for lead authors of the IPCC fifth assessment report on consistent treatment of uncertainties. *Intergovernmental Panel on Climate Change*.

Isaac, A. (2019). The allegory of isomorphism. *Pismo Awangardy Filozoficzno-Naukowej*, *2*, 1–23.

Kormendy, J., & Richstone, D. (1995). Inward bound - the search for supermassive black holes in galactic nuclei. *Annual Review of Astronomy and Astrophysics*, *33*, 581–624.

Lopes, D. (1996). *Understanding pictures*. Oxford University Press.

Lynden-Bell, D. (1969). Galactic nuclei as collapsed old quasars. *Nature*, *223*, 690–694.

Meynell, L. (2013). Parsing pictures: On analyzing the content of images in science. *The Knowledge Engineering Review*, *28*(3), 327–345.

Miyoshi, M., et al. (1995). Evidence for a black hole from high rotation velocities in a sub-parsec region of NGC4258. *Nature*, *373*(6510), 127–29.

Muhr, P. (2023). The “cartographic impulse” and its epistemic gains in the process of iteratively mapping M87’s black hole. *Media & Environment*, *5*(1), 341–357.

Peacocke, C. (1987). Depiction. *The Philosophical Review*, *96*(3), 383–410.

Penrose, R. (1965). Gravitational collapse and space-time singularities. *Physical Review Letters*, *14*, 57–59.

Perini, L. (2005). The truth in pictures. *Philosophy of Science*, *72*(1), 262–285.

Roskies, A. L. (2008). Neuroimaging and inferential distance. *Neuroethics*, *1*, 19–33.

Sartori, L. (2023a). Model organisms as scientific representations. *The British Journal for the Philosophy of Science*. https://doi.org/10.1086/728259

Sartori, L. (2023b). Putting the ‘experiment’ back into the ‘thought experiment’. *Synthese*, *201*(34), 1–36.

Schwarzschild, K. (1916). *Sitzungsberichte der königlich preußischen Akademie der Wissenschaften*. Deutsche Akademie der Wissenschaften zu Berlin.

Skulberg, E. (2021). *The event horizon as a vanishing point: A history of the first image of a black hole shadow from observation* (Ph.D. Dissertation). University of Cambridge.

Suárez, M. (2003). Scientific representation: Against similarity and isomorphism. *International Studies in the Philosophy of Science*, *17*(3), 225–44.

Swoyer, C. (1991). Structural representation and surrogative reasoning. *Synthese*, *87*, 449–508.

Thompson, R. A., Moran, J. M., & Swenson, G. W. (2017). *Interferometry and synthesis in radio astronomy*. Springer Nature.

Tufte, E. R. (1997). *Visual and statistical thinking: Displays of evidence for decision making*. CT: Graphic Press.

van Fraassen, B. (2008). *Scientific representation: Paradoxes of perspective*. Oxford University Press.

Walsh, J. L., et al. (2013). The M87 black hole mass from gas-dynamical models of space telescope imaging spectrograph observations. *The Astrophysical Journal*, *770*(2), 1–11.

Weisberg, M. (2013). *Simulation and similarity: Using models to understand the world.* Oxford University Press.

Willats, J. (1997). *Art and representation: New principles in the analysis of pictures.* Princeton University Press.

Wollheim, R. (1987). *Painting as an art.* Thames; Hudson.