

Scientific understanding through big data: From ignorance to insights to understanding

María del Rosario Martínez-Ordaz
(École Centrale de Lille, Federal University of Rio de Janeiro)
martinezordazm@gmail.com

Forthcoming in *Possibility Studies & Society*.

Abstract

Here I argue that scientists can achieve some understanding of both the products of big data implementation as well as of the target phenomenon to which they are expected to refer –even when these products were obtained through essentially epistemically opaque processes. The general aim of the paper is to provide a road map for how this is done; going from the use of big data to epistemic opacity (Sec. 2), from epistemic opacity to ignorance (Sec. 3), from ignorance to insights (Sec. 4), and finally, from insights to understanding (Sec. 5, 6)

Keywords: ignorance, ignorance of theoretical structure, insights, scientific understanding, big data.

1 Introduction

Among scientists, there is a shared impression that, in the last couple of decades, science has moved from being computationally aided to being data-driven (Cf. Zhou *et al.* 2019: 1018). This transition is seen as changing radically the ways in which knowledge is achieved, novel phenomena are reached, and discoveries are made, among other things. Methodologically speaking, this transition is often reduced to the incorporation of big data and the necessary tools to work such data –like Artificial Intelligence (AI), Machine Learning (ML), Deep Learning (DL), neural networks, etc.. Epistemically speaking, this transition has undoubtedly been linked to novelty, increasing in scope and depth of scientific knowledge, and the strengthening of the web of knowledge.

And while philosophers of science have systematically paid attention to how these changes affected disciplines like astronomy and ecology –which have a very long history of being data-driven, it is a fact that at this point, almost all scientific disciplines are making important use of these novel techniques. For

instance, in molecular and materials science, big data and machine learning have started playing a central role in the discovery and testing of *catalysts*:¹

Catalysts are used in many industrial processes. Traditionally, the optimal design of catalysts has been empirical or has mostly depended on experimentation. Quantum chemical calculations provide the possibility for first-principles catalyst design. However, the large computational cost limits their application to relatively simple reactions and to a small number of catalyst candidates. With the rapidly increasing amount of available experimental and computational data, as well as the development of catalysis informatics, catalyst structure and activity relationships can now be well described using ML models, which are very useful for catalyst development (...) It has been shown that compared with traditional computational and experimental trial-and-error approaches, ML methods possess great potential for accelerating the discovery of high-performance heterogeneous catalysts. (Cf. Zhou *et al.* 2019: 1023).

In this context, ML has become key for the development of the so-called *data-intensive materials design* (or *inverse design*) protocol, in which the desired function of a material is specified beforehand, and then candidates are extracted from a database (which can be either computational or experimental). This allows scientists to predict “the properties of substances, including those of unknown molecules/materials (...)” Therefore, if a single physical property is dominant in governing the performance of a material, molecular/materials informatics serves as an ideal tool to identify a novel functional material” (Toyao *et al.* 2020: 2261). This helps scientists not only to run analyses faster and more efficiently, but also provides them with novel information about scientific objects that might still remain unknown to them.

The use of ML methods for the design, discovery, and testing of catalysts, has changed experimental practices as well as helped scientists broaden the scope of what they could consider observable and knowable. In this sense, it is clear that ML and similar methods have become central not only for the furthering of scientific research in fundamental (and mostly theoretical) sciences but also in those with key roles for dealing with pressing practical issues –environmental, health-related, among others. This change has generated the impression that science is growing in such a way that we, human agents, are every day more capable of understanding the world than we were the day before. But what would that mean to understand the world through science?

Scientific understanding —henceforth, ‘understanding’— consists “of knowledge about relations of dependence. When one understands something, one can

¹A *catalyst* is a substance that expedites chemical reactions and that is not consumed or modified during such reactions. *Catalysis* is the process of modifying the rate of a chemical reaction by adding a catalyst. Catalysis is ubiquitous and can be found in the human body up to large-scale industrial processes. Nowadays, its study and further understanding are crucial for dealing with pressing environmental issues like the production of clean fuels and other chemical products, and it ultimately allows the mitigation of undesired gas or liquid emissions to the environment (see Melián-Cabrera 2021).

make all kinds of correct inferences about it” (Ylikoski 2013: 100). Thus, understanding is a matter of relating doxastic bodies to make specific domains clearer, and in this sense, it consists in building more exhaustive and better-integrated pictures of reality. Given its comprehensiveness, understanding is commonly considered to be one of the ultimate goals of science.

However, with the emergence and implementation of new technologies in different scientific disciplines, the possibility of achieving understanding has reduced significantly. The increase in both the amount of data and the speed at which it is collected generates scenarios in which scientists are profoundly ignorant of both (i) how to interpret the reliability and the content of the data, as well as (ii) how different doxastic bodies about it relate to one another. This leaves scientists in a precarious position when pursuing understanding and gives rise to the dilemma of either accepting that our traditional view on understanding should change in light of new scientific methodologies or committing to the fact that some novel technological resources pull us away from understanding them, their outputs and their target objects —leaving understanding unattainable in some areas of scientific research.

In this paper, I aim at drawing connections between the use of big data in the sciences and issues of epistemic opacity, ignorance, and understanding. Particularly, I focus on the circumstances under which agents can overcome their ignorance and achieve understanding when using big data in their disciplines. I explain how scientists deal satisfactorily with the challenges that ignorance poses for understanding in contexts of big data implementation in the empirical disciplines. The general aim of the paper is to provide a road map for how this is done; going from the use of big data to epistemic opacity (Sec. 2), from epistemic opacity to ignorance (Sec. 3), from ignorance to insights (Sec. 4), and finally, from insights to understanding (Sec. 5 and Sec. 6).

2 From big data to epistemic opacity

A significant part of the current success in science is due to the implementation of extremely complex technological resources, in particular, the use of big data and computational methods. Intuitively, such a level of technological complexity should make us feel assured that the *products* obtained through these technologies are in many senses reliable.² Nonetheless, when looking at these products with much more attention, one would notice that there is very little certainty about the types of epistemic commitments that one can and should endorse toward them.

This section is devoted to addressing some of the main epistemic challenges from the implementation of big data in empirical sciences.

²In what follows, I use the label “products of big data” to encompass a variety of outputs resulting from heavy computational methods implemented to deal with big data; such outputs include predictions, measurements, and classifications, among others. These results should at best be interpreted as models of phenomena, assuming (at a starting point) that their connection with the target domain is significant mediated.

2.1 Preliminaries

Let data be anything that can be soundly recorded in a relational database respecting semantic and pragmatic requirements. “The semantics require that the recordings be understood as true or false statements. The pragmatics suggest that we favor recording what seems to be concrete facts (i.e., singular and relatively weak statements) and that interpreted recordings to be true statements” (Fricke 2014: 652).

Data should meet, at least one of, the following criteria:

- Data whose *raw* form is so large that we must qualitatively change the way in which we reduce, store, and access it.
- Data whose *reduced* form is so large that we must qualitatively change the way in which we interact with and explore it.
- Data whose *structure* is so complex that our current tools cannot efficiently extract the scientific information we seek. (STcI 2016: 5)

In addition, when concerning big data, there are five main characteristics that datasets possess: *volume* (the amount of data that is being managed, measurable in terabytes, petabytes, and even exabytes), *velocity* (the data generation rate and the processing time requirement), *variety* (the data-type, which can be structured, semi-structured, unstructured, and mixed), *veracity* (how accurate or truthful a dataset or a data source may be) and *value* (the possibility of turning data into something useful).

Furthermore, the *computational complexity* of a particular task (problem) results from the number of resources required for its realization (solution). The resources commonly considered when determining the computational complexity of a task are time and space –and when addressing the complexity of algorithms one should also take into account *bit complexity* (the number of operations on bits that are needed for running such algorithm), and communication (the amount of *communication* between the executing parties), among others. The combination of the different characteristics of big data makes the processing of the data extremely computationally complex.

The complexity of big data analytics, and the management of big data, has strengthened the exploitation of AI, ML, DL, among other computational resources. In recent decades, the evolution of both big data and artificial intelligence has had a large impact on the methodological grounds of any scientific discipline, nowadays both are implemented not only to gather evidence but also to explore alternative scenarios and their consequences, to identify interesting and novel (possible) results of theories, models, and experiments, among others.

AI is a rapidly evolving field that involves various domains, such as reasoning, knowledge representation, and machine learning (ML). Machine learning has been widely implemented for numerous drug

discovery applications pertaining to large data sets. It uses various algorithms and techniques to recognize templates and patterns within the given data set (...)ML methods have been classified under two broad subcategories, supervised learning and unsupervised learning methods.(Tripathi et al. 2021: 1440).

That said, it is clear that big data does not come solely with large amounts of data, but with the implementation and improvement of resources such as AI, ML, DL, neural networks, and their mutual combinations, among others. As a matter of fact, and contrary to what its name suggests, the salient feature of big data is not the quantity of information that is gathered. The amount of data that is managed is an important aspect of it, however, what characterizes big data are “the methods, infrastructures, technologies, and skills developed to handle (format, disseminate, retrieve, model and interpret) data. These developments generate the impression that data-intensive research is a whole new mode of doing science, with its own epistemology and norms” (Leonelli 2014: 2). The combination of these changes generates different types of *epistemic opacity*.

2.2 Epistemic opacity

O is *epistemically opaque* to an agent, in a particular context, if the agent ignores all the features of O that are relevant to a specific task within the context. Depending on the object of the epistemic opacity, one can recognize at least two types of opacity present when agents carry out or scrutinize complex computational tasks: opacity about the status of the products of such tasks and opacity about the procedures that underlie those tasks.

- **Opacity regarding the status of the products:** consists of a lack of clarity on whether the models that are created by computer-based methods are substitutes for empirical observations and experimental results, or if they are closer to theoretical abstractions and idealizations (Cf. Barberousse and Vorms 2014, Morrison 2015: Chap. 7). This opacity has an effect on the doxastic commitments that scientists are justified to have towards the products of computer-based methods, and how trustworthy they consider them to be.
- **Opacity regarding the procedures:** consists of a cognitive agent ignoring all of the epistemically relevant elements of a particular process (Cf. Humphreys 2009: 618). When a process is epistemically opaque to an agent, this often has the effect of undermining the strength of the agent’s justification for each of the steps of the process –as well as weakening the justification for the outputs of the process. This type of opacity can be overcome if the agent knows when a step in a procedure is relevant (weak transparency) as well as when it is not (strong transparency) (Cf. Boghossian 1994).

While many of our daily processes are opaque to us to different degrees, what is

special about big data practices and the corresponding computational processes, is that some of these processes are essentially opaque.

- **Essential Epistemic Opacity:** “A process is essentially epistemically opaque to X if and only if it is impossible, given the nature of X, for X to have access to and be able to survey all of the relevant elements of the justification” (Duran and Formanek 2018: 651). This type of opacity can equally be about either the steps within computational processes or the resulting status of the outputs of such processes.

While the types of epistemic opacity that have been described above could be present in many other contexts that do not involve big data, what is relevant for the case of big data implementation is that, once a process or a product is essentially opaque, agents working with it are (or at least, should be) in trouble when addressing their trust in it. This is, on the one hand, essentially opaque computational processes –even when known to be so, are implemented because they are key for achieving scientific success, either novel predictions, measurements, etc. On the other hand, these processes “are so fast and so complex that no human or group of humans can in practice reproduce or understand the processes” (Humphreys 2009: 618); weakening the agents’ capability for rationally justifying their trust in them.

Take the case of contemporary discoveries made in catalysis using big data and data-driven computational techniques. On the one hand, big data and these techniques have been crucial for streamlining the discovery of novel high-performance materials. On the other hand, the data that underlies the identification of such materials is often treated by implementing extremely complex computational processes –that go beyond human computational capabilities. The combination of the above makes the nature of such discoveries very unique.

While not all big data implementation entails epistemically opaque processes and results, there are some that are necessarily considered to be the result of *black-box models*. Black box models are extremely computationally complex models whose internal logic is not readily interpretable, this is, in general, the processes carried out within them are unknown to the agents. The large majority of black box models currently used in catalyst materials discovery are essentially epistemically opaque for human agents; these models include Gaussian process models and neural networks. “The motivation for using these approaches is that, in many cases, the design space of possible catalysts is too large to be studied using quantum chemical methods alone. ML models serve as computationally efficient surrogates to minimize expensive quantum chemical calculations, enabling an accelerated screening of the catalyst design space” (Esterhuizen *et. al.*, 2022: 175).³

The above leaves chemists often able to satisfactorily address the relevance of the discoveries made by using black box models, but rarely meeting the same

³For an analysis of black box methods and interpretable machine learning in heterogeneous catalysis see [Esterhuizen *et. al.*, 2022] and for an extensive discussion on epistemic opacity in computational chemistry see [Wieber and Hocquet 2020]. It is important to stress that not all black box models use ML nor all of them are subject of essential epistemic opacity.

success when having to disentangle the processes followed in such discoveries. This is, at the same time that big data has revolutionized the methodologies of catalyst materials discovery, it has done so by incorporating new veils into the epistemology of the discipline.

It is important to notice that while computer scientists have extensively researched ways to convert black-box models into glass-box models, these attempts even if successful do not suffice to provide scientists with the exact knowledge that they were initially looking for. For instance, *Explainable Artificial Intelligence* (XAI) aims at producing possible reconstructions and explanations of processes and methods that are essentially opaque to humans, these explanations are most of the time only alternatives —this is, are made in such a way that they are accessible for humans to understand and trust but not necessarily to 'open' the actual black box of the processes that were carried out.⁴

Summing up, in big data contexts, epistemic opacity (either essential, about products, or processes) comes as a direct result of the complexity of the computational tools that are used in collecting, reducing, and structuring immense amounts of data:

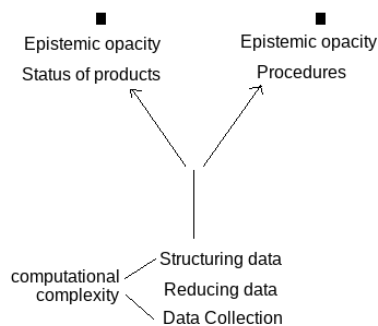


Image 1. From computational complexity to epistemic opacity.

I now turn to explain how all these challenges relate to a more familiar one: ignorance.

3 From epistemic opacity to ignorance

This section is devoted to explaining how epistemic opacity relates to one particular type of ignorance: ignorance of theoretical structure with reliable consequences.

⁴For a comprehensive analysis of the epistemic advantages of XAI see [Páez 2019].

3.1 Opacity and ignorance of theoretical structure

Intuitively, epistemic opacity resembles that of ignorance. Traditionally, ignorance has been characterized as a "lack of knowledge".⁵ Epistemic opacity about the status of a procedure's output might be considered a case of factual ignorance, given a certain proposition, lacking knowledge of whether it refers to a fact, or even if it is true. Opacity regarding the procedures might seem like an instance of procedural ignorance, just not knowing a procedure. Yet, this translation of epistemic opacity into ignorance is a bit simplistic, it makes epistemic challenges of computational procedures into cases of ordinary epistemic problems. The question is whether there is anything special about the ignorance involved in big data practices. Here I contend that the ignorance that underlies big data implementation is of an inferentialist-spirited type: *ignorance of theoretical structure with reliable consequences*. The reasoning is the following:

- First, take a structure of a theory to be, broadly constructed, a set of inference patterns that, when put together over the elements of the theory, constrain it and allow it to make sense of both the content of the theory as well as the domains on which the theory is correctly applied.
- Second, ignorance of theoretical structure has been characterized as lacking knowledge of

the (relevant) inference patterns that scientific theories allow for. When ignoring (the relevant parts of) the theoretical structure of a theory, scientists are not capable of grasping abstract causal connections between the propositions of their theory, they can neither identify the logical consequences of the propositions that they are working with nor can explain under which conditions the truth value of such propositions will be false. (Martínez-Ordaz 2020: 12)⁶

⁵I am fully aware that there is an ongoing philosophical debate about the different ways to approach and understand ignorance. However, given the nature of the discussions endorsed here, the adoption of this characterization of ignorance will suffice for the purposes of the paper. Nonetheless it is important to mention that there are two major philosophical frameworks for the study of ignorance. The first is of an analytic epistemological nature and the second is of a much more naturalist scope. The former encompasses rival views on ignorance as a lack or absence of knowledge and ignorance as a lack or absence of true belief (Le Morvan and Peels, 2016: 12). The latter approaches ignorance as a more complex phenomenon that emerges from epistemic, cognitive, psychological and even social features. For extensive discussions on this latter approach, see [Arfini, S. 2019] and [Arfini and Magnani 2022].

⁶When ignoring a segment of the theoretical structure of a theory or model, one cannot identify the inferences that, within it, lead to adequate predictions and explanations, among other interesting products. In addition, the ways in which the (epistemic and logical) reliability of such a theory or model is preservable are unknown. The inferentialist spirit of this type of ignorance lies in the fact that ignorance of theoretical structure prevents scientists from either inferring the value of certain propositions, identifying whether distinct sequences of properties refer to the same object, or explaining inferential procedures that underlie the management of the information.

- Third, big data implementation requires the gathering of data of different types (images, redshifts, time series data, and simulation data, among others) coming from sources of also very different kinds, it isn't surprising that the sets of data are not always fully compatible. To solve this issue scientists rely on extremely complex computer-based resources to filter and structure the information; yet, when these resources go further human cognitive abilities, this has the effect of scientists losing track of the inferential mechanisms that determined the later structure of the data. When scientists "cannot provide inferential explanations about why an output obtains, they are not ignoring only a specific recipe, they are ignorant of how the bits of data relate to one another –at least, inferentially; and this is indicative of ignorance of theoretical structure" (Martínez-Ordaz 2022: 127). It is important to notice that the presence of ignorance of theoretical structure often is the source of other, more traditional, instances of ignorance, such as factual, objectual, and procedural ignorance.
- Fourth, the standard reason for which scientists tolerate these high degrees of ignorance is the quality of the most successful outputs of computer-based processes. This is, it is well-known that the incorporation of big data into the empirical sciences comes with new levels of epistemic opacity; however, scientists are willing to pay this price because the use of these resources helps science to grasp distant or complex objects that without them would have never been at our reach, it also helps us to conceive novel scenarios and to 'witness' new phenomena, among other things.⁷

The most successful outputs of big data implementation in science, those that justified our reliance on big data, are often (1) novel in their fields, (2) empirically adequate, (3) fruitful (crucial for the development of related research programs), and (4) hold possible evidential relations with models or theories within the discipline (cf. Martínez-Ordaz 2022).

∴ Therefore, the ignorance that underlies big data implementation in the sciences has two salient components: its inferential nature and the reliability of the products obtained through big data.

From the outset, I want to be clear about the dialectic. I am not claiming that all successful products and processes of big data implementation are epistemically opaque for us. Yet, in those cases in which epistemic opacity is present, scientists are ignorant of the theoretical structure that constrains the building of such products. I am aware that there are many more epistemic problems associated with complex computational procedures and epistemic opacity, problems that often deal with these issues from a computational perspective. However, here I only focus exclusively on the challenges that scientists as individual human epistemic agents might find when working with big data in their disciplines.

⁷There is a sense in which these technological tools are used to overcome our own cognitive and physical limitations; therefore, their use is only motivated and justified if they can help us to do so if they can allow us to increase the rate of success in science.

Going back to current practices in catalysis. In 2021, the *Schoenebeck Research Group* reported having predicted 21 phosphine ligands using unsupervised machine learning with only five experimental data points (along with in-silico data);⁸ remarkably, such a set of ligands included never made ones (Cf. Hueffel et al. 2021). One of the most interesting features of this prediction was that the phosphine ligands form air-stable Palladium(I) dimers, whose geometry and air stability were over the ones of Palladium(0) and Palladium(II) species, and this made them very promising catalysts (Cf. Welter 2021). Yet, the novelty of this prediction is stressed by the fact that their chemistry is still not well understood yet, and while the implementation of ML helped to predict new theoretical entities, it also helped to highlight the important gaps of explanatory knowledge that might exist around them.

The combination of the above leaves chemists in a peculiar position in which they can address the relevance and reliability of the discovery for the discipline, even to the point to develop new lines of research around such a discovery; but at the same time, struggle explaining the internal logic of the models through which the discovery was obtained as well as the theoretical framework in which this discovery fits.

Summing up, while the epistemic feature of big data implementation that is the easiest to spot is epistemic opacity (either about products or procedures), what underlies it is an ignorance of theoretical structure.

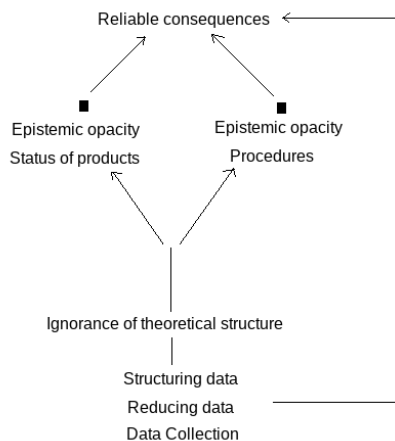


Image 2. Epistemic opacity and ignorance of theoretical structure with reliable consequences.

⁸The most salient feature to take into account with respect to unsupervised ML techniques is that they “can be applied to recognize patterns in datasets without requiring training of the algorithm with labeled data (and therefore without the known outputs, such as experiments). The learning process provides insights that are fundamentally different from traditional analyses, as they are derived purely by the “machine” without “human” guidance” (Hueffel et al. 2021: 1136).

To repeat: I take this section to have shown that there is a strong relation between epistemic opacity and ignorance, they are not the same but when epistemic opacity is present, it causes different presentations of ignorance of theoretical structure with reliable consequences. First, the reason why agents fail at figuring out the status of certain products or at determining the steps that are followed in procedures is that they lack knowledge of the inferential constraints of the building of these products and procedures. Second, when the outputs of big data are extremely successful —either because of their novelty, their accuracy, their scope, etc., they are commonly considered epistemically reliable; even if coming from opaque processes. Of course, successful consequences of big data implementation can also come from not opaque processes; however, what justifies the toleration of epistemic opacity, is the salient reliability of some of the products that result from epistemic opaque scenarios.

4 From ignorance to insights

The word ‘insight’ is commonly used to indicate either an epistemic product or a belief-formation process. Regardless of its usage, however, insight typically denotes a sense of envisioning a solution to a problem through an opaque process that remains unclear to the agent who experienced it. This creates a seeming inconsistency in the nature of insights: on one hand, they often engender a feeling of certainty, even in the absence of clear justification; on the other hand, the opacity of the process by which insights are formed should logically diminish one’s trust in them.

Similarly, agents may form strong beliefs based on the outputs of big data, even when they are aware of the opacity of the processes that produced these outputs. Such high levels of trust are typically based on the perceived reliability of the outputs at a given moment, such as their accuracy, innovativeness, precision, or overall utility.

This section deals with the question of which doxastic commitments that epistemic agents endorse towards products of opaque or unclear processes. In order to do so, here I explore the connection between epistemic opacity, ignorance, and insights.

4.1 Insights

Insight consists of a sudden realization or discovery of a solution path that allows one to solve a problem. The inferential mechanism that underlies the building of insights is ampliative, this is, the result is novel to the agent who experiences the insight. While there is common agreement on the fact that these are the key elements of insights, it remains unclear the complexity of the entity referred to as “insight”.

The term “insight” is used in literature to refer to either the solution of a particular problem or to the process through which that solution is achieved. The latter definition, which encompasses mostly the process of insight, is based on

Pierce’s characterization and has been endorsed by philosophers and cognitive scientists. Pierce’s work was instrumental in identifying the type of reasoning involved in producing insights, which is often ampliative, creative, and outside-the-box, and it is often present in abductive contexts. Furthermore, insights are typically accompanied by a feeling of surprise. When experiencing an insight, an individual may report having discovered new information that they believe to be efficient, trustworthy, or functional for solving a specific problem. However, it is also acknowledged that this information has been obtained through an unclear path. Contemporary studies that view insight as a process typically emphasize the cognitive elements that limit mental processes related to discovery and innovative problem-solving. While these studies aim at providing detailed explanations of how human agents produce insights, they often neglect two important factors: the potential role of external aids in facilitating insights, and the attitudes that epistemic agents should adopt towards them.

In contrast, the understanding of insight as an epistemic product pertains to the type of commitment (epistemic) agents have towards the solution that is produced through a given inferential path. When looking at insights as epistemic products, one centers the attention on the commitments that agents have towards the solution of a problem that resulted from following a very creative and unclear path. At first glance, insights involve forming the belief that “X is the solution for Y problem”.⁹ This understanding of insights suggests that this type of beliefs consist of three key components:

- they are formed through an unclear or unrigorous process as a response for a given problem,
- they appear to be strong and robust enough to guide our acceptance or rejection of other beliefs.
- Additionally, insights are an indication of our grasping of a specific problem, object, domain, or phenomenon.

Given the normative perspective taken in this paper, I adopt the second approach to insights, which retains both the opaque nature of belief formation processes and the strong epistemic commitment to the achieved solution of a problem. These features are essential to insights, regardless of one’s preferred general view of them. However, considering insights as beliefs allows us to avoid the elements of human psychology involved in the reasoning that gave rise to them and to offer an interpretation that brings together (human) epistemology and advanced technological implementation.¹⁰

The fundamental characteristics of insights suggest that, at their best, they present a conflict, and at their worst, a contradiction. How can a belief that is endorsed so strongly also result from an unclear process? Moreover, how can insights be considered rational beliefs? These questions arise from the apparent

⁹Or simply “X is the case”

¹⁰A similar discussion on the epistemic merit of outputs of insights-like processes can be found in [Aliseda-Llera 2023. In Spanish].

tension between the strength of insights and the lack of clarity in the process that generates them.

First of all, human agents endorse insights so strongly because they allow us to evaluate some epistemic virtues of specific sets of beliefs considering the ways in which other beliefs relate to them. But this endorsement does not mean that agents never give up insights, what it means is that they do it only when they are faced with either the falsity of the insight or with strong evidence of the incompatibility between the specific insight and a set of core beliefs (that are better supported than the insight in question). Furthermore, for agents to be rational when endorsing a specific insight they have to explicitly regard it as being “trustworthy”. One is rational when trusting something (or someone), if such a trust is justified, at least, either truth-directly or end-directly.¹¹

This considered, beliefs formed through insights share similarities with both scientific hypotheses and knowledge. Scientific hypotheses are proposed explanations about a particular domain that result from a problem-solving process. They are tentative statements that are subject to testing and revision in light of new data, yet they also serve as a starting point for new lines of research. In contrast with hypotheses, when endorsing an insight, the process of testing their truth is not as pressing; as a matter of fact, in the case of insights, what matters the most is how they are taken as a starting point to both deal with the problem that gave rise to them as well as pursuing further research. Insights share some similarities with factual knowledge (knowledge of facts), mainly, both involve believing that something is the case. Yet, in the case of insights, there is also an acknowledgment that the belief forming process is unclear or unknown. This recognition sets insights apart from knowledge, which is typically grounded on the possibility of providing justification for the reliability of the process that gave rise to the belief. As such, insights represent a unique type of belief that is both tentative and based on a process of problem-solving, yet also acknowledges the limitations of its formation.

Finally, it is important to emphasize that because the rationality of the endorsement of insight comes from gathering evidence in favor of either its truth or its role in the achievement of specific goals, its rational character is only temporary. This is, through this search of evidence either we succeed or we get to a point in which we have to admit that the strength of the insight should have degenerated in absence of new evidence that supported either its truth or its value for meeting goals.

4.2 Ignorance and insights in big data

Very often, the most salient scientific achievements that are produced through big data implementation are taken as touchstones for the development and pur-

¹¹One can say that S ' trust in an insight i is rational only if S has gathered enough evidence in favor of the truth of claim i is *trust-worthy* –despite the fact of ignoring the processes through which i has obtained. In addition, S will be rational when trusting i , if to assume the trustworthiness of i plays an essential role for Y , and Y is S ' goal in the relevant context (Cf. McLeod 2020: Sec. 2.1).

suit of novel research lines. And, regardless of the presentation of such findings, whether they are predictions, measurements, etc.; they are often indicative of knowledge of objects or phenomena. Nonetheless, the large majority of them are obtained through processes that are opaque and unclear for human agents. This resembles the situation of insights.

Take for instance the case of the prediction of 21 phosphine ligands obtained by the Schoenebeck Research Group when employing unsupervised machine learning algorithms. First, it is clear to the scientists that this result was produced via epistemically opaque methods. In particular, some of the most novel outputs were very likely to have been never investigated without the implementation of ML techniques (or similar tools) –as the algorithms followed a research route very distant from those intuitive for the experts (Cf. Hueffel et. al, 2021: 1138). However, the accuracy of the predictions is considered to be remarkable and the prediction in itself is taken as revealing something about an object in the world that deserves further analysis (Cf. Hueffel et. al, 2021; Welter 2021).

With regard to the scientists’ commitments towards (some of) the successful products of big data implementation, there is a peculiar combination of a strong acceptance of their trustworthiness with the awareness of the opaqueness of the processes through which they were obtained. This gives the impression that the belief-forming mechanisms that are used to bring beliefs around the trustworthiness of the outputs of big data are similar to those underlying the beliefs about the trustworthiness of a solution to a given problem –in insight-contexts. It is important to notice that scientists’ confidence over certain products of big data is not formed via hunches –as we tend to imagine that insights in our daily life are. Nevertheless, the beliefs behind this confidence satisfy the conditions for insights by recognizing that the products that determine their content were formed through unclear and opaque processes, as well as by being strongly accepted to the point in which they help to determine the acceptance/rejection of other beliefs.

In these scenarios, it seems significantly challenging to assess the trustworthiness of the outputs of big data and computer-based methods. First of all, because there is no clarity about the status of many of these outputs, it becomes hard to appeal to their truth in order to establish their trustworthiness. In addition, due to the opaqueness of the procedures behind them, it is impossible for us to ‘infer’ the truth of the products by tracking down the steps through which they were built. Furthermore, determining their trustworthiness by appealing to the role that they might play is a possibility, particularly, in those in which the output was being explicitly searched for. Yet, the large majority of saliently novel products of big data implementation are unexpected discoveries that end up grounding new research programs –which, most of the time, aim at explaining them. So, to establish their trustworthiness by pointing to their role in scientific research seems a quite difficult task.

When seeking ways to justify the trustworthiness of a particular novel output, it often becomes obvious that the scientists’ acceptance of the output comes from the fact that successful products of big data, most of the time, provide scientists with “access to empirical phenomena especially if those phenomena that

wouldn't be accessible to humans without the aid of big data and computational processes, and enhances the achievement of objectual knowledge regarding such phenomena.” (Martínez-Ordaz 2022: 128). This is, in order to justify the end-directed trustworthiness of a product of essentially opaque processes, the most common way to do so is to indicate which is the role that this output might play in the achievement of further epistemic products such as knowledge and understanding—even if in the long run, the output is discovered to be not true.

Summing up, when the products of big data implementation are extremely novel, scientists are inclined to (strongly) endorse them even if they ignore the inferential patterns that constrained the products' building.

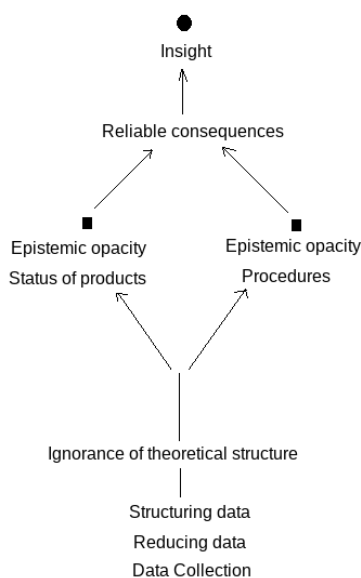


Image 3. From ignorance to insights.

The reliability of such products clashes significantly with the fact that agents are ignorant of not only the steps followed by procedures but, more importantly, of the inferential constraints of such products and procedures. So, the fact that simultaneously certain products can ground beliefs that are crucial for scientific development, and at the same time, scientists cannot explain the basis of their building, leaves us with the impression of these beliefs being insights. It is important to say that, for the purposes of the next section, only those beliefs that are grounded by or contain part of the reliable products are to be considered insights—leaving undetermined the role that beliefs about other consequences of datasets may play.

In the following sections, I deal with the role that these insights play in the moving forward of science.

5 From insights to understanding

This section addresses the basic notions behind scientific understanding. In addition, it also scrutinizes some of the challenges to the achievement of understanding in big data contexts.

5.1 The basics of understanding

Scientific understanding consists in putting together bits of knowledge in such a way that the result is a cohesive picture of a specific domain —at least, a more cohesive one than the resulting ones from each of these bits of knowledge alone. Considering its integrative nature, understanding is often seen as the ultimate goal of science. The (ideal) scientist is expected to be capable of knowing, explaining, and understanding her theories, the phenomena that such theories depict, as well as the procedures that are followed in her discipline.

Yet, understanding a theory or a particular phenomenon is very different, epistemically speaking, from understanding a procedure. This indicates that there can be two main types of understanding in science, a theoretical understanding and a practical one (cf. Bengson 2018). The former is rooted in the acquisition and exercise of explanatory knowledge; this is, knowing why certain things occur the way they do within (or according to) a particular theory or model. The latter relates to the ability to perform tasks in a successful way, and it is often linked to procedural knowledge.

Another salient feature of scientific understanding is the combination of a strong psychological component with an objective one.

- The former is the feeling of *grasping*: Take grasping to refer (at least) to “a cognitive state bearing some resemblance to scientific knowledge of some part of the explanatory nexus” (Khalifa 2017: 11). The feeling of grasping refers to the sense of satisfaction that comes with realizing one has acquired the ability to put together bits of information that shed light on parts of an explanation.¹²

However, this sensation depends solely on the individual agent’s experiences, which often can be misguided. For this reason, the objectivity of understanding comes from requiring

- that what is grasped is a fragment of reality (Cf. Elgin 2007: 35). While this condition could inspire many philosophical discussions, the basic idea is that the agent should possess significant evidence of the content of her understanding being grounded in the world. This is, for the case of empirical sciences, legitimately understanding an empirical domain would require that the agent can interact with it in a satisfactory way —explaining

¹²I am fully aware that the precise characterization of grasping and its relation with understanding is still object of philosophical debate. This, however, should not drive the attention away from the fact that the feeling of grasping is a combination of a psychological element with an epistemic one, this is, is an epistemic feeling. For general takes on epistemic feelings see [Arango-Muñoz 2014] and [Greely 2021].

it, predicting it, accounting for its parts and the ways in which they relate, etc.

Additional to these two features, understanding also requires order and coherence, which when combined, allow *intelligibility* to emerge. First, because understanding is an integrative task, order is key. As it was explained in the previous paragraph, the difference between legitimately understanding something and just having the feeling of grasping it lies in the mind-independent grounds of such a feeling. This suggests that, while the same elements could be arranged in many different ways, only some of them are privileged regarding their correspondence with the domain. So the identification of these orders and structures is a necessary component of the objectivity of understanding. Second, coherence results from the combination of consistency, compatibility, and reinforcement (Cf. Elsamahi, 2005). A cluster of bits of knowledge is *consistent* if and only if it is impossible to form a contradiction from them. Two bits of knowledge within a cluster are *mutually compatible* if they are mutually consistent and they ‘talk’ (at least partially) about the same domain; this, of course, strengthens the motivation for their later union. Two bits of knowledge in a cluster *reinforce* each other if either one provides a “rationale” for the other or if, at least, one supports the basic assumptions of the other or explains it (Cf. Elsamahi, 2005). Third, *intelligibility* is “the value that scientists attribute to the cluster of virtues (of a theory in one or more of its representations) that facilitate the use of the theory for the construction of models” (De Regt 2009: 31).

Considering all of the above, to understand something is to be able to order the components of what has been understood in a coherent way (Cf. Bengson 2018: 19). Scientific understanding is gradual, and can always be improved (with regards to either its depth or its scope). Understanding is an extremely valuable product because it requires an exhaustive effort to be attained. In addition, it used to be thought that understanding should be both *factive*, meaning that its content ought to include only true propositions, and *explanatory*, this is, the acquisition of understanding should follow the prior acquisition of explanatory knowledge.¹³ These two criteria are independently motivated, but they are mutually reinforcing as they head in the same direction: the epistemic robustness of understanding.

5.2 The conflicts between understanding and big data

When big data is implemented in scientific contexts, our first intuition would be that it will help us to gain an understanding of novel phenomena or at least contribute to the improvement of the understanding that has been previously

¹³One might wonder whether there is a meaningful difference between explanation and explanatory understanding. While this is a topic of significant controversy among philosophers, a difference between the two is that scientific explanation concerns either causes or mechanisms that underlie a particular phenomenon, while understanding encompasses a broader comprehension of the underlying principles, patterns, and relations of dependence of a given phenomenon.

gained. However, according to many epistemologists, this is not necessarily the case. Unfortunately, there are three conflicts between understanding and big data that have been put forward in the literature: the first concerns the relation between understanding and explanation, the second, the relation between truth and understanding, and the third focuses on the challenges for the identification of relations of dependence.¹⁴

- **Explanation.** Since Hempel, understanding has been continuously linked to explanation g (Cf. Kvanvig 2003; Grimm 2006, 2014; Kelp 2014; Sliwa 2015; Lawler 2016, 2018). The most salient cases of scientific understanding are those that involve the previous acquisition of explanatory knowledge about the phenomenon that will be later understood. Now, because the outputs of big data implementation are only in the form of correlations, and because correlations do not suffice for an explanation, philosophers of science have concluded that these outputs won't suffice for the achievement of understanding.
- **Factivity.** Another element of scientific understanding that conflicts with big data are the truth value of the content of understanding –the so-called factivity condition. Not only because of its alleged relationship with explanation but also due to its objective component, scientific understanding has been traditionally considered to only include true propositions. If we expect to grasp a segment of the actual world, this can be only possible if the elements that constitute what we understand are true in the actual world. According to this view, we are allowed to use fictions, idealizations, and other non-true items when seeking for understanding; yet, this does not imply that those items are in any significant way part of what we have understood, only that they were useful tools (Cf. Lawler 2020). If epistemic opacity surrounds the status of the outcomes of big data and, because of this, it is impossible for the agents to determine whether a specific outcome should be seen as a punctual description of reality, as an abstraction, etc.; one will also be unable to include it in the content of understanding around a target phenomenon. This challenges both the understanding of the output as well as its inclusion in the understanding of the target phenomenon.
- **Relations of dependence.** Because agents are ignorant of the theoretical structure behind the datasets that originated the outputs, these agents cannot identify the relations of dependence from which the outputs are obtained. And in this sense, it seems impossible to understand both the outputs and their associated procedures. Furthermore, traditionally, and in spite of big data implementation being a source of groundbreaking results,

¹⁴I am aware of the fact that there are important ongoing debates concerning these aspects of understanding in general. But considering the purposes of this section, and because in the corresponding literature, the connections between understanding and explanation have been presented as some of the most serious obstacles to the achievement of understanding in big data contexts, here I do not discuss the alternative standpoints.

the understanding of its products and processes, as well as of the phenomena described through these products, seems extremely complicated –if not impossible. It is important to clarify that I am not saying that the traditional view on understanding and big data is that the products of big data implementation are useless for the pursuit of scientific understanding. What has been said in the literature is that these products cannot alone promote our understanding, that the processes that generated them cannot be understood, and that they cannot play a central role in the achievement of understanding.

Furthermore, traditionally, and in spite of big data implementation being a source of groundbreaking results, the understanding of its products and processes, as well as of the phenomena described through these products, seems extremely complicated –if not impossible. It is important to clarify that I am not saying that the traditional view on understanding and big data is that the products of big data implementation are useless for the pursuit of scientific understanding. What has been said in the literature is that these products cannot alone promote our understanding, that the processes that generated them cannot be understood, and that they cannot play a central role in the achievement of understanding.

6 From insights to understanding (II)

This section addresses how the ignorance of theoretical structure can be (partially) overcome allowing some understanding to be achieved. Here, I particularly deal with the role that insights play in this matter.

6.1 The road to understanding

As it might be obvious to the reader at this point, the possibility of achieving understanding requires, at least, the partial overcoming of the scientists' ignorance. In particular, as the underlying ignorance in big data practices is of a relational (inference) nature, and because understanding is also a relational phenomenon, the overcoming of the former seems necessary for the achievement of the latter. The resources that scientists have to do so are their theoretical frameworks (independent theories, models, etc.), their most important scientific observations, as well as the salient products of big data implementation that have been accepted as groundbreaking –and around which scientists have formed insights.

The main claim here is that the reliable outputs of big data implementation are keystones in the achievement of scientific understanding, without them, understanding would be unattainable to scientists.

The road that takes scientists from ignorance of theoretical structure to scientific understanding can be broadly described through the following five steps:

1. **Acknowledgment of ignorance.** The starting point is for scientists to recognize that the key elements of the ignorance that underlies big data implementation are inferential relations. This acknowledgment leads to seeing how the inferential nature of their ignorance prevents them from explaining the procedures that they guide, cannot determine the constraints of the set of information that gave rise to the successful outputs, and cannot determine the truth value of such outputs.

Going back to the Schoenebeck Research Group’s prediction of 21 phosphine ligands. First, it is necessary to say that the reason why this research was conducted using opaque resources was that

To accurately predict the favored speciation of catalysts on the basis of mechanistic and quantum mechanical considerations, it is necessary to have precise knowledge of the various potential species in solution that may (or may not) form, their coordination states (with or without solvent), spin or charge states, and potential dynamic interconversions. Such information is rarely accessible in full, and it is therefore not surprising that there is to date so little understanding of the factors that dictate catalyst speciation.(Cf. Hueffel et. al, 2021: 1134).

For instance, gray-box methods like High-Throughput Experimentation (HTE) and SubGroup-Discovery (SGD) require around 100-10 000 experimental data points to be able to satisfactorily navigate the number of possible materials (which is practically infinite) and arrive at a neat identification of the needed catalyst material (Cf. Foppa et. al. 2021). Because alternative methods, including insight-driven strategies, do not suffice for this speciation challenge, scientists were in need of employing opaque methods to do so. As a result, they were fully aware of their ignorance about the road taken by the algorithms, and the ways in which, from only five experimental data points, it was possible to arrive at such a novel prediction.

2. **Identification of reliable consequences.** A crucial element for the rational toleration of ignorance and epistemic opacity is the identification of the payoffs of big data implementation in their disciplines. What justifies such toleration is the identification of the outputs of big data implementation that are considered to be extremely reliable and crucial for the development of science and the explanation of the role that they play that matter.¹⁵
3. **Formation of beliefs grounded around the most reliable consequences (insights).** This involves both beliefs, about the reliability

¹⁵The acknowledgment of ignorance of theoretical structure with reliable consequences allowing scientists to become aware of the target problem to solve: the identification of inference patterns that allows us to make sense of, at least, the reliability of the most salient products of big data implementation.

of these consequences as well as beliefs grounded in the consequences themselves. These beliefs are considered to be insights. On the one hand, because they are firmly endorsed by scientists because of their role in scientific development –this warrants the doxastic strength of insights. On the other hand, because their truth cannot be stated by tracing the quality of the output (through the procedure that generated it) nor by ‘checking’ its relation with the actual world –which provides the unclear/opaque/unrigorous basis of the belief.

Going back to our case study. The novelty of the prediction resulting from the implementation of unsupervised machine learning for the identification of palladium catalyst, was taken by the scientists as extremely reliable. As a matter of fact, it was taken as “a clear demonstration of the power of machine learning techniques to accelerate catalyst development with suggestions that are beyond a scientist’s intuition. Our future efforts are directed at exploring the potential of the new dimers in catalysis” (Cf. Hueffel et. al, 2021: 6). This should be taken as indicative of both the identification of the prediction as novel and reliable as well as the acceptance of such a prediction as reliable –at least for the purposes of guiding future research round its most novel results.

- **Identification of inferential patterns.** The (partial) overcoming of ignorance of theoretical structure requires the identification of particular inference paths that connect the reliable outputs of big data with the best theoretical frameworks at the scientists’ disposal and the most entrenched observations (Cf. Martinez-Ordaz 2022). The search for these paths requires taking the big data outputs, the frameworks, and the observation as fixed points that are assumed to be true.

In the case study described above, there are two mutually complementary ways in which the identification of inferential patterns should be carried out. On the one hand, the theoretical path, which consists in taking the palladium(I) dimers (particularly, the one that had never been made), whose chemistry has not been well studied, as their synthesis seemed unpredictable, and search for adequate embeddings within theoretical and experimental models. On the other hand, the methodological path consists in assessing the reliability of the methods used for the production of the prediction of the palladium (I) dimers. For this matter, scientists might employ tools like Interpretable ML (IML) methods, that aims at “translating the hidden patterns identified by ML models into interpretable information formats can lead to testable theories and hypotheses, further advancing scientific understanding.” (Esterhuizen et. al. 2022: 175). The combination of these two paths will shed light on the epistemological status of the output taking into account its relevance within the discipline as well as the method that was used for its production.

- **Building networks of understanding.** Scientific understanding consists of building *networks* that successfully connect our scientific beliefs

about the world and allow us to obtain a detailed map of specific regions of it. The most reliable outputs of big data implementation, the theories, and the most robust scientific observations work as the nodes of the network; meanwhile, the inference paths that the scientists have selected to connect them shape the network and determine its strength and its robustness. Furthermore, these inferential paths also define the (logical) constraints of the network of understanding. The importance of doing such an integration of the product of big data into a broader network is that of strengthening its value within a particular theoretical view on the domain. This is especially important taking into account the epistemic limitations of tools like IML and XAI; whose reconstructions and explanations are, most of the time, only alternatives to what actually happened in the black box. “As helpful as interpretation tools might be, ML cannot eliminate the role of catalysis scientists in advancing scientific theories and hypotheses. (...) We believe that, if possible, the best practice is to use features that align with earlier physical explanations, as the interpretation is likely to be more insightful if it reinforces or connects to pre-existing domain knowledge.” (Esterhuizen et. al. 2022: 182)

That said, the general picture is the following:

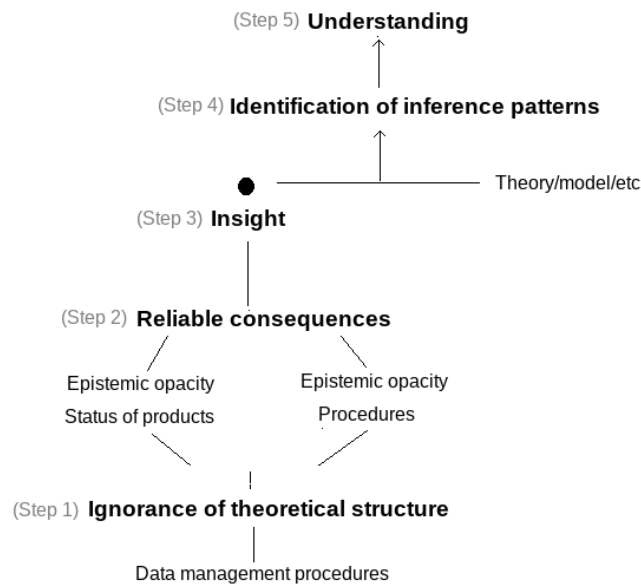


Image 4. From ignorance to understanding.

Now, there are crucial questions to be addressed: which type of understanding is achievable in big data practices? In the following paragraphs I focus on this issue.

6.2 Interpreting understanding

Let's deepen the particularities of the understanding that results from following the road that I just described above.

First, the integration of the insights from big data implementation into specific theoretical frameworks is crucial for both the understanding of the target phenomena as well as of the corresponding product of big data. And because of the nature of the ignorance that scientists deal with in these contexts, such integration has to take place, at least, at an inferential level. Second, when connecting the successful products of big data implementation to a specific theory or model, scientists identify or produce the logical bridges (inference patterns) that would make the product trustworthy, and therefore, legitimize the insights about its reliability. This constitutes the building of a specific *logical space*.

This logical space is constrained by the conditions according to which the selected products of big data are reliable and the target phenomenon is coherently described by the theory or model. It is important to notice that the resulting logical space is constrained by the theoretical framework, the insights around the trustworthiness of the big data product as well as the scientists' previous knowledge of the target phenomenon; and because of this, it will be narrower than the ones build taking into account only one of these components. Yet, these logical spaces can only tell a possible story, a story that *might* be the case.

How informative really are both the relations of dependence that *might* constrain the trustworthiness of the product of big data as well as those that *might* constrain the information about the target phenomenon? And more importantly, do they suffice for understanding? These questions are grounded in the fact that while the identification of inference paths might lead scientists to a cohesive picture of a particular domain using the outputs of big data implementation, this does not mean that this particular picture is in any relevant sense connected to the target phenomena. Therefore, as the integration of elements does not suffice for understanding, it is not clear that the finish line described above matches any type of understanding. However, the building of a logical space around the insights into the trustworthiness of the products of big data provides scientists with a particular type of scientific understanding: *modal understanding*. It is often said that someone has a modal understanding of X when that person knows how to navigate the possibility space associated with X (Cf. Le Bihan 2017: 112).

In big data contexts, a modal understanding of the reliability of the products of big data is achieved when scientists are able to determine under which circumstances, theoretically and empirically speaking, these products are reliable. As well as when they are able to navigate the associated possibility space in order to connect the reliability of such products with one of the other similar future outcomes of technological implementation. For the case of the target phenomenon "to achieve a modal understanding of the behavior of novel objects in an established theoretical domain would be to determine the set of possible worlds that correspond to the generic structural features assumed by the the-

oretical view that such a cluster of data substantiates.” (Martínez-Ordaz 2022: 131).

It is important to emphasize that while *insights* do not suffice for understanding, they work as fixed nodes within a structure, playing a crucial role in the constraining of the possibility space associated with the outputs of big data as well as with the target phenomenon. The resulting logical space is built around the insights that scientists have about the outputs of big data. Furthermore, because this type of understanding is surrounded by the procedural epistemic opacity that underlies the building of big data products, its scope is narrow. Modal understanding in this case refers to the identification of the inferential conditions under which it would make sense to use a theoretical framework together with outputs of big data implementation to justify, explain and use our empirical knowledge about a target phenomenon. This shouldn’t be confused with having acquired knowledge of relations of dependence between doxastic bodies that we know occur in the *actual* world.

If what has been said here is on the right track, it helps to motivate discussions about the role of scientific community for the achievement of understanding. Because with the advance of big data a growing inter/transdisciplinary collaboration is needed, this suggest that the social component will also play a significant role in the acceptance and rejection of products of big data, as well as in the building of networks of knowledge and scientific understanding around those products. And although this issue is extremely important for assessing the particularities of scientific understanding in data-driven scientific research, it is beyond the scope of this work.¹⁶

7 Final remarks

Here I addressed the possibility of achieving scientific understanding in contexts of big data implementation in science. I focused on those cases in which reliable products of big data are obtained through essential epistemically opaque procedures.

Pace traditional views on scientific understanding, I argued that understanding was available for scientists even under these conditions and that certain outputs of big data implementation play a crucial role in this matter. In order to do this, I first explained that the prevalent types of epistemic opacities found in big data implementation are the result of the way in which computational complexity prevents scientists from knowing the inferential constraints of some procedures and products (Sec.2). I identified this lack of knowledge as ignorance of the theoretical structure of the datasets, the procedures that are used

¹⁶Thanks to a reviewer for raising this issue. Discussions on the importance of transparency and communication in (collective) epistemology and ethics of AI can be found in [Russo, Schliesser and Wagemans 2023: Sec. 4]. Also discussions on the importance of communication between expert-communities and non-experts for the acceptance/rejection of the outputs of black box algorithms can be found in [Durán and Jongsma 2021]. Finally, see [Cesar-Jimenez, forthcoming -in Spanish] on the value of AI and human collaborative interactions for epistemic purposes.

to manage these datasets and their products. I next explained the way in which despite scientists being ignorant in this sense, they are capable of identifying reliable outputs of big data implementation—which in the most fortunate cases are considered to be groundbreaking (Sec.3). Furthermore, I explained that the beliefs that are built around these reliable outputs are insights (Sec. 4) and that they are crucial for the later building of networks of (modal) understanding (Secs. 5, 6).

Acknowledgments.

I am indebted to Moisés Macías-Bustos for rich conversations on these issues and his extensive feedback on different versions of this paper. Thanks to Carlos César-Jiménez, Selene Arfini, Federica Russo, Atocha Aliseda and Alexiadis Vaíos for valuable comments on these issues. Thanks to the anonymous referees for their suggestions. And thanks to the editors of this volume, Selene, Wendy and Samantha, for their patience, help and understanding.

An early version of this paper was presented in the workshop “Enacting Chance: Ignorance, Insight and Intuition” in 2021.

Funding.

This research has been financed, in part, by Agence Nationale de la Recherche (ANR) under the project ANR-21-CHIN-0005-01 and by the Programa Nacional de Pós-Doutorado PNPd/CAPES (Brazil).

References

1. Aliseda-Llera, A. (2023): “Towards a characterization of clinical reasoning: an epistemological perspective” (Hacia una caracterización del razonamiento clínico: una perspectiva epistemológica) in A. Aliseda-Llera, C.M. Calderón-Aguilar, and M. Villanueva (eds) *Filosofía de la medicina: discusiones y aportaciones desde México*, UNAM.
2. Arfini, S. (2019). *Ignorant Cognition. A Philosophical Investigation of the Cognitive Features of Not-Knowing*. Springer.
3. Arfini, S. and L. Magnani (Eds.) (2022): *Embodied, Extended, Ignorant Minds: New Studies on the Nature of Not-Knowing*. Springer International Publishing.
4. Bowden, E. M., Jung-Beeman, M., Fleck, J., Kounios, J. (2005). New approaches to demystifying insight. *Trends in Cognitive Sciences*, 9(7), 322–328.
5. Barberousse A. and M. Vorms (2014): “About the warrants of computer-based empirical knowledge”, *Synthese* 191(15): 3595–3620.

6. César-Jiménez, C. (forthcoming): “Conversational Agents and Hybrid Intelligence” (Agentes Conversacionales e Inteligencia Híbrida), in K. González (ed).
7. Chen, Y., Argentinis E. and G. Weber (2016): “IBM Watson: How Cognitive Computing Can Be Applied to Big data Challenges in Life Sciences Research”, *Clinical Therapeutics* 38 (4): 688-701. de Regt, H. W. (2009): “Understanding and Scientific Explanation” in de Regt, H.W., S. Leonelli and K. Eigner (eds.) *Scientific Understanding: Philosophical perspectives*, Pittsburgh University Press: 21-42.
8. Durán, J. M., and N. Formanek (2018): “Grounds for Trust: Essential Epistemic Opacity and Computational Reliabilism”, *Minds and Machines: journal for artificial intelligence, philosophy and cognitive sciences*.
9. Durán, J.M. and K. R. Jongsma (2021): “Who is afraid of black box algorithms? on the epistemological and ethical basis of trust in medical AI”, *J. Med Ethics* 47: 329-335.
10. Esterhuizen, J.A., Goldsmith, B.R. Linic, S. (2022): “Interpretable machine learning for knowledge generation in heterogeneous catalysis”. *Nature Catalysis* 5, 175–184.
11. Elsamahi, M. (2005): “Coherence between theories”, *Canadian Journal of Philosophy* 35(2): 331-352.
12. Feigelson, E.D. and G.J. Babu (1997): “Statistical methodology for large astronomical surveys”, in McLean, B., D. A. Golombek, J. J. E. Hayes, and H. E. Payne (eds.) *New Horizons from Multi-Wavelength Sky Surveys: 363-370*.
13. Floridi, L. (2011): *The Philosophy of Information*, Oxford UK: Oxford University Press.
14. Floridi, L. (2012): “Big data and their epistemological challenge”, *Philosophy Technology* 25(4): 435–437.
15. Floridi, L., Fresco N. and G. Primiero (2015): “On malfunctioning software”, *Synthese* 192(4): 1199–1220.
16. Foppa L, Sutton C, Ghiringhelli LM, De S, Löser P, Schunk S, et al. (2022): “Learning design rules for selective oxidation catalysts from high-throughput experimentation and artificial intelligence”. *ACS Catalysis*, 12, 4, 2223–2232.
17. Fricke, M. (2015): “Big data and its epistemology”, *Journal of the Association for Information Science and Technology* 66(4): 651–661.
18. Hueffel, J.A., T. Sperger, I. Funes-Ardoiz, J.S. Ward, K. Rissanen, and F. Schoenebeck (2021): “Accelerated dinuclear palladium catalyst identification through unsupervised machine learning”. *Science* 374, 1134-1140

19. Humphreys, P. (2009). The philosophical novelty of computer simulation methods. *Synthese*, 169(3), 615-626.
20. Khalifa, K. (2017): *Understanding, Explanation and Scientific Knowledge*, OUP.
21. Le Morvan, P. and R. Peels (2016): "The Nature of Ignorance: Two Views" in R. Peels. and M. Blaauw (eds.) *The Epistemic Dimensions of Ignorance*, Cambridge University Press: 12–32.
22. Leonelli, S. (2012): "When humans are the exception: Cross-species databases at the interface of biological and clinical research", *Social Studies of Science*, 42(2): 214–236.
23. Leonelli, S. (2014): "What difference does quantity make? On the epistemology of Big data in biology", *Big data Society* 1(1): 1-11.
24. Leonelli, S. (2020): "Scientific Research and Big data", in E. N. Zalta (ed.) *The Stanford Encyclopedia of Philosophy*, <https://plato.stanford.edu/archives/sum2020/entries/science-big-data/>
25. Martínez-Ordaz, M. del R. (2020): "The ignorance behind inconsistency toleration", *S.I. Knowing the Unknown*, *Synthese*.
26. Martínez-Ordaz, M. del R. (2022): "Is there anything special about the ignorance involved in Big Data practices?", in Lundgren, B. L.. and N. Nuñez-Hernández (Eds.) *Philosophy of Computing*, *Philosophical Studies Series*, Vol. 143.
27. Melián-Cabrera, I. (2021): "Catalytic Materials: Concepts to Understand the Pathway to Implementation", *Industrial Engineering Chemical Research*. 60, 51: 18545–18559. Morrison M. (2015): *Reconstructing Reality: Models, Mathematics, and Simulation*, New York, NY: Oxford University Press.
28. Napoletani, D., Panza M. and D. C. Struppa (2014): "Is Big data Enough? A Reflection on the Changing Role of Mathematics in Applications", *Notices of the American Mathematical Society*, 61(5): 485–490.
29. Páez, A. (2009): "The Pragmatic Turn in Explainable Artificial Intelligence (XAI)", *Minds and Machines* volume 29: 441–459.
30. Pierce, C. S. (1992): *Reasoning and the Logic of Things: The Cambridge Conferences Lectures of 1898* (Ed. by K. L. Ketner), Harvard University Press.
31. Ross, W., and Vallée-Tourangeau, F. (2022). Insight with stumpers: Normative solution data for 25 stumpers and a fresh perspective on the accuracy effect. *Thinking Skills and Creativity*, 46, 101114.

32. Russo, F., E. Schliesser, and J. Wagemans (2023): "Connecting ethics and epistemology of AI", AI Society, Springer.
33. Saha B. and D. Srivastava (2014). "Data quality: The other face of Big Data," 2014 IEEE 30th International Conference on Data Engineering: 1294-1297.
34. Sterner, B. and N. M. Franz (2017): "Taxonomy for Humans or Computers? Cognitive Pragmatics for Big data", Biological Theory, 12(2): 99–111.
35. Swan, M.(2015): "Philosophy of Big data: Expanding the Human-Data Relation with Big data Science Services," 2015 IEEE First International Conference on Big data Computing Service and Applications: 468-477.
36. Symons, J., and Alvarado, R. (2016): "Can we trust Big data? Applying philosophy of science to software", Big data & Society, 3(2), 2053951716664747.
37. Tripathi, M.K., Nath, A., Singh, T.P. et al. (2021) "Evolving scenario of big data and Artificial Intelligence (AI) in drug discovery". Mol Divers 25, 1439–1460 .
38. Welter, K. (2021): "Unsupervised machine-learning tool could accelerate catalyst discover", Chemistry World, 02 December 2021: <https://www.chemistryworld.com/news/unsupervised-machine-learning-tool-could-accelerate-catalyst-discovery/4014842.article>
39. Wheeler, G. (2016): "Machine Epistemology and Big data", in McIntyre, L. and A. Rosenberg (Eds.): The Routledge Companion to Philosophy of Social Science. Routledge: 321-329.
40. Wieber, F. and A. Hocquet (2020): "Models, Parameterization, and Software: Epistemic Opacity in Computational Chemistry". Perspectives on Science; 28 (5): 610–629.
41. Ylikoski, P. (2009): "The Illusion of Depth of Understanding in Science", in H. De Regt, S. Leonelli and K. Eigner (eds.) Scientific Understanding: Philosophical Perspectives, Pittsburgh University Press: 100-119.
42. Zhang, Y. and Y. Zhao (2015): "Astronomy in the Big data Era". Data Science Journal, 14(11): 1–9.
43. Zhou, T., Z. Song and K. Sundmacher (2019): "Big Data Creates New Opportunities for Materials Research: A Review on Methods and Applications of Machine Learning for Materials Design" in Engineering, Volume 5, Issue 6: 1017-1026.