

# Beyond transparency: computational reliabilism as an externalist epistemology of algorithms

Juan M. Durán<sup>[0000-0001-6482-0399]</sup>

**Abstract** This chapter is interested in the epistemology of algorithms. As I intend to approach the topic, this is an issue about epistemic justification. Current approaches to justification emphasize the transparency of algorithms, which entails elucidating their internal mechanisms –such as functions and variables– and demonstrating how (or that) these produce outputs. Thus, the mode of justification through transparency is contingent on what can be shown about the algorithm and, in this sense, is *internal* to the algorithm. In contrast, I advocate for an *externalist* epistemology of algorithms that I term *computational reliabilism* (CR). While I have previously introduced and examined CR in the field of computer simulations ([42, 53, 4]), this chapter extends this reliabilist epistemology to encompass a broader spectrum of algorithms utilized in various scientific disciplines, with a particular emphasis on machine learning applications. At its core, CR posits that an algorithm’s output is justified if it is produced by a reliable algorithm. A reliable algorithm is one that has been specified, coded, used, and maintained utilizing *reliability indicators*. These reliability indicators stem from formal methods, algorithmic metrics, expert competencies, cultures of research, and other scientific endeavors. The primary aim of this chapter is to delineate the foundations of CR, explicate its operational mechanisms, and outline its potential as an externalist epistemology of algorithms.

---

Juan M. Durán  
Department of Values, Technology and Innovation  
Faculty of Technology, Policy and Management  
Delft University of Technology  
Jaffalaan 5  
2628 BX Delft  
The Netherlands, e-mail: j.m.duran@tudelft.nl

## 1 Introduction

The use of algorithms for scientific purposes is delivering remarkable results. A couple of examples will suffice to illustrate this. In molecular biology, AlphaFold can predict protein structures with atomic accuracy in cases where no similar structures are known [1]. In medicine, BenevolentAI has combined structured and unstructured biomedical data sources to identify rheumatoid arthritis drugs like *baricitinib* as citerapeutics for COVID-19 symptoms [6]. In an increasingly number of cases, algorithms have successfully extended cite class of tractable chemistry, biology, physics, and medicine, broadening cite range of modeling and experimental capabilities available to researchers.

Yet, unlike other methods, the algorithm’s scientific merits cannot be easily determined by association with a body of scientific knowledge, by adequacy to empirical data, or by diverse theoretical constructs –such as explanation and observation. This is for a variety of reasons. Algorithms are epistemically and methodologically opaque [51], making it difficult to associate a given algorithm and its output with the general scientific canon. Likewise, empirical phenomena are often temporarily, spatially, or cognitively inaccessible for validation of the algorithm, potentially casting doubts over any representational value of these systems.

When confronted with these issues, philosophers and computer scientists gravitate towards *transparency*, an umbrella term capturing diverse methods linking the internal mechanisms and properties of algorithms to their outputs [26, 33, 16]. To see how transparency works, consider BenevolentAI. At its core, this algorithm is a search engine that combines structured and unstructured biomedical data sources, drug industry data, and automated retrieval of information from diverse scientific research papers. The data is curated and standardized via data analysis and data fabric. It is then fed into knowledge graphs that structure the data into relationships between diseases, genes, and different drugs [23]. Richardson led the team that used BenevolentAI to identify rheumatoid arthritis drugs – notably *baricitinib* – as suitable therapeutics for COVID-19 symptoms [18]. To justify Richardson’s belief in the scientific value of this output, partisans of transparency would focus on showing how *baricitinib* is rendered from procedures integrating biomedical data, instantiation of key variables, function calls identifying structural relationships within the algorithm, relevant conditional statements, and other algorithmic operations. Another way to make BenevolentAI transparent is via a *knowledge graph* [18, 30]. This visualizes how *baricitinib* inhibits AAK1 (associated with interrupting the COVID-19 virus’ passage into cells) and JAK 1/2 (critical for signal transduction pathways), and how *baricitinib* binds with GAK (known to decrease certain viruses’ infectiousness). This knowledge graph also provides reasons to consider drugs like *fedratinib*, *sunitinib*, and *erlotinib* as less effective and, depending on the case, unsafe. For instance, it is shown how these drugs only inhibit AAK1 and neither decrease the chances of cell infection (by binding with GAK) nor inhibit cytokine signaling (by inhibiting JAK 1/2) [18, 30].

Are Richardson and his team justified in believing that *baricitinib* is a medically valid outcome for the issue at hand? What reasons do they have to discard other drugs

as either less effective or unsafe? What supports their claim that BenevolentAI is a reliable system for the intended purposes? These are questions about the epistemic reliance on ML and the justification of their outputs. To a great extent, transparency provides answers to these questions. This paper, however, is an effort to provide an alternative answer, one that does not depend on methods for the transparency of the algorithm. More specifically, this paper lays the groundwork for *computational reliabilism* (CR), a reliabilist epistemology centered on algorithms, aimed at justifying their outputs.

As the name suggests, CR borrows bits and pieces from epistemological reliabilism, notably Goldman's process reliabilism [19, 50]. However, the version of CR I develop here draws from, but also expands on, my previous work on computer reliabilism for computer simulations [53, 4]. With these ideas in mind, this chapter is divided as follows. Section 2 presents and discusses two epistemologies of algorithms: one that is internal to the algorithm (e.g., transparency) and one that is external to the algorithm (i.e., CR). As expected, these epistemologies have different modes of justification, which are exemplified in section 2.1. The example provided is only intended to motivate CR. In section 3, I lay the groundwork for *computational reliabilism* (CR). Here, I present three types of *reliability indicators* (type-RIs) that credit reliability to algorithms. These are (1) type<sub>1</sub>-RI technical performance of algorithms (subsection 3.1.1), (2) type<sub>2</sub>-RI computer-based scientific practice (subsection 3.1.2), and (3) type<sub>3</sub>-RI social construction of reliability (subsection 3.1.3). In section 4, I briefly take stock of my findings and suggest further lines of investigation that substantiate the merits of CR. In gist, this article invites us to reflect on a crucial but often overlooked question: under what conditions are researchers justified in believing algorithms' outputs? My answer is that reliability comes through myriad methods, practices, and processes at diverse stages of specification, coding, use, and maintenance of the algorithms.

## 2 Internalist and externalist epistemologies for algorithms

A central motivation for seeking justification is that algorithms are often epistemically opaque. This concept has two distinct but related interpretations. The first interpretation addresses how algorithms involve multiple complex elements (functions, variables, decisions, data, etc.) in their specification, coding, execution, and maintenance. This means that little can usually be said about how these algorithms cluster data, which criteria are used for creating categories, and overall why algorithms behave the way they do. This interpretation is captured in the epithet 'black-box' algorithms as a way to express how far removed algorithms sometimes are from human insight. The second interpretation sets the focus on our limited capacities to say something meaningful about the output of an algorithm [51, 649]. That is, no human being (or group of human beings) can know which functions, variables, decisions, data, etc. are relevant to a given output.

Whereas the first interpretation focuses on the algorithm as an opaque method, the second highlights our cognitive, epistemic, and other limitations in making knowledge claims about the algorithm’s output. Both interpretations, I believe, can be cast as human agents lacking proper justification for the algorithm’s output. Here justification is taken to be epistemic, and understood in the general sense of a belief being formed in the proper manner. Thus, either because the algorithm is a black-box or because human agents are cognitively limited, there is no basis for claims about the proper formation of a belief about whether the algorithm’s output is true, has scientific value, or can be epistemically trusted.<sup>1</sup>

Under this heading, transparency surfaces as a promising epistemology of algorithms. It first requires uncovering the inner mechanisms and properties of the algorithm, and then linking these to its output. Human agents are justified by successfully revealing the functions, values, etc., that produced the algorithm’s output ([55, 51, 27, 24, 26]). Recall from the introduction that BenevolentAI utilizes *knowledge graphs* to visualize how the algorithm favors baricitinib over other drugs. Having access to how the knowledge graph works and that it rendered baricitinib justifies the belief in the algorithm’s output. Another example is LIME: a general algorithm that accounts for the predictions of any classifier by locally learning an interpretable model. Formally, LIME produces a model  $g \in G$ , where  $G$  is a class of potentially interpretable models (e.g., linear models, decision trees, falling rule lists). In practice, if an algorithm predicts that a patient has the flu, LIME can highlight the symptoms in the patient’s history responsible for the prediction. ‘Sneeze’ and ‘headache’, for example, are key variables used by LIME. Indeed, they are flagged as net contributors to the flu prediction. In contrast, ‘no fatigue’ is a variable used as evidence against the prediction [16]. Let us note in passing that many forms of explanatory AI (XAI) provide a rich source for transparency, as they often involve tracking back the *path-dependency* of the algorithm that relates a given function (or set of functions), variables, etc., to its output [15].

Thus understood, transparency purports justification as an *internal to the algorithm* matter. That is, the justification of our beliefs that the output is true depends exclusively on some form of surveying the inner workings of the algorithm. To put this idea more or less formally,

---

<sup>1</sup> For simplicity and continuity with the literature on justification, I shall talk of a belief as being “true” – or “false” –, scientifically valid – or sound –, trusted – or doubtful. However, I will refrain from defending a full-blown realist or anti-realist position. I believe this is a debate that philosophers interested in algorithms need to address at some point – see the Chapter 3 by Casey in this volume. Thus, to believe that the algorithm’s output  $\bar{o}$  is to take it that  $\bar{o}$  is true, has scientific value, can be epistemically trusted, etc. Following Elgin [46], *truth* will not be understood as (absolute) correspondence with reality, but rather as being sufficiently adequate for the purposes at hand, allowing for practical engagement, scientific progress, and understanding (also [11]). Note that talking in these terms doesn’t mean that an agent  $S$  must explicitly believe the proposition that  $\bar{o}$  is true, since the latter is a different and higher-order belief. That is to say, mere belief that the algorithm’s output  $\bar{o}$  is true doesn’t require possession of the concept of “truth”. Equally important is to note that our beliefs are not necessarily occurrent at any given time, that beliefs come in degrees of strength and confidence, and are historically situated, incremental, and perspectival [48]. Thanks go to Jack Casey for the close reading of my assumptions and for saving me from making further mistakes.

**Definition 1** A human agent  $S$  is justified in believing the algorithm's output  $\bar{o}$  just in case: a) it is shown, directly or indirectly, the algorithmic path-dependency to  $\bar{o}$ ; and b)  $S$  has reasons to believe that the path-dependency to  $\bar{o}$  are the case.

Opposing this view is computational reliabilism (CR), here presented as an *external to the algorithm* epistemology consisting of identifying (formal) methods, algorithmic metrics, expert competencies, cultures of research, and the like that make up our best epistemic and normative efforts to specify, code, and maintain reliable algorithms. I shall call these *reliability indicators* (RIs) and, as I shall explain in section 3, they can be divided into *types* and *tokens*. By construction, then, CR does not depend on showing the internal mechanisms and properties of the algorithm. Instead, it depends on reliability indicators that are external to the algorithm. A primer working definition can be,

**Definition 2** A human agent  $S$  is justified in believing the algorithm's output  $\bar{o}$  if and only if  $\bar{o}$  was rendered by a reliable algorithm. A reliable algorithm is one that produces true outputs  $\bar{o}$  most of the time. To this end, the algorithm must have been specified, coded, and maintained through diverse reliability indicators.

While I will dedicate a large portion of this chapter to the characterization of type and token reliability indicators, a preliminary conclusion can be drawn now: we can say that transparency and CR have different justificatory modes. According to the former, we have justification by having access to the inner workings of the algorithm. According to the latter, we have justification by identifying methods (formal and otherwise), metrics, expert competencies, cultures of research, and the like external to the algorithm that make up our best epistemic and normative efforts to increase the algorithm's reliability.

As a final attempt to illustrate these two justificatory modes, let me briefly present and discuss an example of an algorithm that classifies individual suspects as {criminal; non-criminal} based on their facial traits. It will be shown that transparency justifies the belief that a given suspect is a criminal –or a non-criminal– whereas CR flags the algorithm as unreliable and therefore lacking justification for such beliefs. It goes without saying that this example is only meant to contrast these two justificatory modes. No conclusions about their individual value as epistemologies are intended to be derived from it.

## 2.1 Merchants of mistrust

In 2016, computer scientists Xiaolin Wu and Xi Zhang developed a Convolutional Neural Network (CNN) that analyzed over 1,850 ID photos and classified them as {criminal; non-criminal}.<sup>2</sup> About 1,120 of these photos were of people with no criminal convictions, and the remaining were of people who were either wanted

<sup>2</sup> It is worth noticing that Wu and Zhang's use of photos from actual people, in contrast with other approaches that use synthetically generated photos [20, 31].

for crimes or convicted of crimes. The CNN's operation was simple. It picks out facial traits (e.g., distance between the eyes, length and curvature of the mouth) and classifies each photo as {criminal} or {non-criminal}. No other concept or category was operational. Despite this – or perhaps because of this – the predictive accuracy measured using the Area Under the Receiver Operator Characteristic Curve (AUC-ROC) was very impressive: Wu and Zhang measured 0.9540 accuracy in the classifications. This means that the CNN was able to successfully classify faces of individuals as being {criminal} or {non-criminal} approximately 95% of the time [44, 2].

To further validate their algorithm and rule out that such a high predictive accuracy resulted from overfitting, Wu and Zhang retrained the CNN on a dataset where the labels 'criminal' and 'non-criminal' were assigned randomly as negative and positive instances with equal probability. For the retraining case, the CNN failed to distinguish between the two categories, plummeting the average classification's accuracy to 48%, with a false negative rate of about 51%, and the false positive rate close to 50%. Wu and Zhang also accounted for problems related to unbalanced datasets, choice of photos (light, angle, over and under exposure, clothing, etc.), and other issues pertaining to accuracy. To most algorithmic standards, these results speak in favor of a reliable CNN capable of consistently classifying the photos in question.

Wu and Zhang naturally defend the scientific merits of their algorithm. To their mind, as to many, high predictive accuracy means that the algorithm's outputs have scientific value, are true, etc. It is thus no coincidence that they confidently announce the "law of normality for faces of non-criminals" [44, 8]. But high predictive accuracy is no standard for claims about scientific value or truth. One could argue that while the Ptolemaic model exhibited high predictive accuracy in its measurements, the model fundamentally misconceived and misconstrued planetary motion. Additionally, we know that predictive accuracy can be manufactured by carefully selecting the input data and calibrating variables and functions in the algorithm to some desired degree. For example, finding optimal values for hyperparameters (number of hidden layers, batch size, choice of activation function, etc.) is fundamental for having faster convergence, high accuracy, and overall better results. Now, algorithms allow multiple optimal hyperparameter configurations depending on datasets, purposes of the algorithm, and tasks [25]. Furthermore, optimal configurations for one algorithm do not typically translate to others, making them incompatible in many different ways [38]. As a result, selecting optimal values for hyperparameters, along with the best configuration for a given algorithm, is largely a matter of human decision. Without further provisions in place, such as ensuring compliance with scientific standards, professional integrity, and standardized measurements for the optimality of hyperparameters, predictive accuracy can (relatively easily) be manufactured.

These authors would insist that high predictive accuracy grants scientific value to their CNN's outputs. To further defend this, they retrained the parameters of every layer in the CNN while also modifying the architecture [7, 3]. As a result, the high accuracy in the output remained at the same levels. In fact, the CNN correctly picks out specific facial attributes from photos, and then classifies them

into the appropriate category ca. 95% of the time.<sup>3</sup> But again, taking high predictive accuracy as an indication of the reliability of automated inference on criminality algorithms is problematic. It confounds justification of a technically correct output with the justification required for believing that output.<sup>4</sup> Wu and Zhang have no justification for believing that someone is a criminal based on facial traits alone. This is the case regardless of how accurate their algorithm is at picking out and classifying photos.

In this context, transparency does not seem to be of much help for justification. When Wu and Zhang try to justify their outputs on high predictive accuracy, they look at what their AUC-ROC values are telling them. This means that specific inner functions and properties of the CNN responsible for the output will support the justification. But justifying the CNN's output using the same functions used to produce them is epistemically circular and inadmissible for the proper formation of beliefs. Rather, these functions only speaks of the algorithm's robustness, and only in a very limited way. It follows that Wu and Zhang can pin down the functions and properties of their CNN that account for the high predictive accuracy, but at no point can they use those functions and properties alone for claims about justification. In other words, transparency here does not help to distinguish what we are compelled to believe from what cements that belief.<sup>5</sup>

### 3 Computational reliabilism (CR)

Claims about justification find a home in CR, a branch of process reliabilism where subject *S* is justified in believing output  $\phi$  if the algorithm is reliable (see definition 2 on page 5) [53, 4, 50]. An algorithm is reliable when it produces  $\phi$  that are true rather than false most of the time. It is important to note that reliability here is not merely a matter of track record but rather about the algorithm's propensity to generate true  $\phi$  in most cases. Now, the debate in epistemology over the most suitable version of reliabilism is extensive and cannot be addressed here. Suffice it to say that I favor *propensity reliabilism* over Goldman's *frequentist reliabilism*, aligning with Alston, for whom "[a] reliable instrument is one that *would* usually deliver favorable results over an appropriate range of cases *if and when* they occur" [59, 6]. I will not discuss this point further as the relative frequency or propensity of CR are unproblematic

<sup>3</sup> Despite these efforts, the system's high accuracy remains questionable. While there is no evidence of output manipulation, one can't help but wonder whether the system would maintain the same level of accuracy when faced with a larger and more diverse datasets.

<sup>4</sup> What is operating here is the distinction between output accuracy, which is concerned with the correctness of the final results produced, and procedural accuracy, which is concerned with the execution of steps and adherence to methods.

<sup>5</sup> As suggested earlier, transparency is a broad concept that admits different interpretations. A partisan of *post-hoc* explanatory AI, for instance, could argue that the algorithm was not taking into account scientifically salient aspects of the pictures. By means of this, one could in principle identify high-level features that refer to domain knowledge (e.g., a list of criminality-based characteristics) and thus have grounds – reasons, evidence – for claims about justification.

for the purposes of this chapter. The question is rather, how to confer reliability to an algorithm. To achieve this end, CR utilizes *reliability indicators* (RIs) as markers of methodological, cognitive, social, and epistemological competence. RIs are any algorithmic-related methods, metrics, practices, domain-specific knowledge, and the like with a reliability-conferring property. Although I will neither discuss the nature of this reliability-conferring property nor how it operates, a simple example should illustrate that it is not of a ‘spooky’ kind –rather, it is very familiar to many philosophers of science [57].<sup>6</sup> Consider the microscope. Claims about its reliability stem from, say, the use of the laws of optics for its construction and calibration, the effective observation of entities also dependent on the researcher’s prior knowledge, and a scientific community with the background education and capacity to accept or reject an observation. The proper workings of the instrument, the knowledge of the right methods, and the social validation of an observation all confer reliability to the microscope. The RIs I shall discuss shortly play similar reliability-conferring roles, as they amount to accessible scientific practices, methods, cultures of research, scientific debates, and other (more or less) scientifically-grounded activities. Nothing spooky about that. The real challenge, rather, is to be as precise as possible in identifying RIs for specific cases. Here is where this chapter falls short. However, this is for a good reason. Recall that my only pretense with this chapter is to lay down the groundwork for an externalist epistemology of algorithms, and therefore my treatment of CR will be very general. The reader interested in concrete applications of CR to different domains is cordially invited to read [9] for cases on medicine and healthcare, and [13] for forensic science.

### 3.1 Reliability Indicators

For conceptual clarity, I distinguish between *type*-RIs and *token*-RIs. While the former refers to a unique category of indicators, the latter refers to an individual occurrence for that category. With this distinction in mind, the following *type*- and *token*-RIs are at the heart of CR:

- *Type<sub>1</sub>-RI - Technical performance of algorithms* focuses on the specification, coding, execution, maintenance, and other technical features that contribute to the performance of the algorithm (e.g., high accuracy and low rate of errors, but also tolerance to domain change, repurposability, reusability, modularity, etc). In this sense, typical cases of *token<sub>1</sub>-RI* include practices and protocols for collecting, curating, storing, distributing, and analyzing data; the use of out-of-distribution data and data augmentation, parametrizations; benchmarking; choice of architecture; treatment of algorithmic kludges [56]; recasting [2]; error treatment, and other techniques pertaining to achieving the desired performance of algorithms. Within this *type*-RI could also be included a justification for the employment of

---

<sup>6</sup> This is, of course, not to say that CR faces problems (See Chapter 5 by Alvarado in this volume).



said practices, metrics, and methodologies, along with the specific circumstances in which the algorithm is specified and coded.

- *Type<sub>2</sub>-RI - Computer-based scientific practice* focuses on securing algorithmic-based scientific research. It results from the operationalization and implementation of scientific concepts, causal structures, models and theories, laws and law-like principles, taxonomies, but also scientific metaphors and intuitions, values (epistemic and otherwise), idealizations, abstractions, and representations. This type-RI intends to capture the degree to which scientific units of analysis are implemented and operationalized into the algorithm. The selection and justification of domain knowledge are equally crucial in enhancing an algorithm's reliability. Let us note that the viability and success in doing so largely depend on algorithmic-related decisions, such as programming language choice, the use of formal techniques like verification methods [36], and the utilization of sub-modeling and multi-modeling [2].
- *Type<sub>3</sub>-RI - Social construction of reliability* focuses on broader goals related to accepting – or rejecting – algorithms and their outputs by diverse communities (e.g., scientific, academic, the general public), the realization of intended values and goals, and the overall assessment of the algorithm's scientific merits. This occurs through token<sub>3</sub>-RI such as debates, experimenting and testing, replicability of results, and other forms of intellectual exchange.<sup>7</sup>

Under this heading, CR is understood as a family of reliability-eliciting algorithmic-related indicators capable of crediting an algorithm as a reliable belief-forming method. It is important to note that by accepting a reliabilist epistemology, one also accepts the propensity likelihood that governs the reliability of a process. This translates into acknowledging that algorithms can occasionally be inefficient, contain errors, be unsuitable for specific purposes, misrepresent, and compute incorrect results. If failures perpetuate over time, the relative propensity governing CR shifts, rendering the algorithm ultimately unreliable.

Furthermore, proponents of CR take note of human cognitive limitations in accessing some token-RIs, which conditions the claims about the reliability of an algorithm. They also need to accommodate the fact that token-RIs are neither absolute nor universally applicable. Not all token-RIs are credited, relevant, and applicable under the same criteria, nor does the same token-RI equally apply to all algorithms. CR is thus understood as perspectival, provisional, and subject to corrections, with no particular token-RI considered to have an all-or-nothing reliability-conferring property.

Thus understood, token-RIs come in degrees. The degree to which one token-RI is more relevant than another, or contributes to the overall reliability of the algorithm will depend on the context in which the algorithm is specified, coded, used, and maintained. It will depend on the epistemic and non-epistemic values and goals at

---

<sup>7</sup> Let me echo what Heather Douglas [29] persuasively argued: scientific and computational practices, as presented in type<sub>1</sub>- and type<sub>2</sub>-RI, along with the social processes tailored to them, as presented in type<sub>3</sub>-RI, are neither reducible to one another nor completely uncoupled. This sentiment applies here as well.

stake. It will also depend on the culture of specifying, coding, maintaining, and using the algorithm of a given community [47]. In this sense, no individual (set of) token-RI can guarantee the reliability of all algorithms. Furthermore, even under the assumption that some token-RI is suitable for a given algorithm, this does not ensure that our reliability claims are eternally warranted. Old token-RIs can lose their appeal as new ones come to light. For these reasons, this chapter holds no pretensions to claim the completeness of the various type- and token-RIs presented here. Further arguments could be given on the need for additional type- or token-RIs not discussed here, or that some indicators are somewhat misplaced, or that some others need replacement. None of this is to say, however, that there are no stable type- and token-RIs that apply across many reliable algorithms. In fact, most of the token-RIs discussed next maintain, to my mind, a permanence in time despite changes and fine-tuning that occur with new technological and scientific developments.

Finally, I recognize that CR may not be readily accepted by everyone. As a reliabilist epistemology, one might feel that it still needs to address a few concerns. For starters, there are issues pertaining to the relevance and availability of type- and token-RIs, potential conflicts emerging among token-RIs, and their precedence, order, and weight. Unfortunately, these issues will not find a complete answer here. To my mind, it is the richness and urgency of this problem that requires putting into practice demands for an account at least as complex as the one presented here. In this sense, CR does not provide, nor intend to provide, absolute assurances. Instead, CR aims to highlight that our best epistemic efforts can be geared towards the reliability of algorithms. Little more can be expected given the fallibility and limitations of human cognition. Taking note of these caveats, I now discuss a few types- and token-RIs in more details.

### 3.1.1 Type<sub>1</sub>-RI: Technical performance of algorithms

In earlier versions of CR [42, 53, 4], RIs mainly focus on the specification, implementation, tractability, and overall performance of algorithms. For instance, the first three token<sub>1</sub>-RI discussed in [4] put forward defining criteria for assessing the utility value of algorithms and their outputs *qua* computational methods. Consider *validation* procedures as an example.<sup>8</sup> In automated diagnosis, algorithms are used for patient prognosis. One way to increase our confidence in the output is to compare the disease progression as indicated by the algorithm with clinical data from prior patients that share the same endotype or phenotype [10]. This practice validates the synthetic data rendered by the algorithm with empirical data collected via diverse scientific methods (e.g., observation, experimentation, intervention, measurement, and others). The utility value of the algorithm is then considered appropriate if validation standards are satisfactory.

In this respect, subjecting algorithms' outputs to validation methods increases – or reduces – our confidence in the reliability of an algorithm, as it is a good

---

<sup>8</sup> It is important to recognize that various forms of verification and validation exist, each conferring different degrees of reliability. Decisions must be made [58].

indication of the algorithm's accuracy and margin of error. Validation methods also give a fair sense of the capacity to generalize the algorithm from the training data to new, undiscovered data. From a scientific perspective, validating algorithms also contributes to the rigor and reproducibility of research, ensuring that findings are based on sound methods.

Now, it should be expected that validation methods encompass a variety of techniques and methods.<sup>9</sup> As such, they are not all appropriate for the same goals. This means that a given validation techniques cannot be simply applied to different algorithms without prior critical discussion. There must be agreement on how suitable a given validation technique is for the algorithm and data in question, as well as the purposed goals and tasks [41, 32]. This is an often overlooked aspect of the social dimension of engineering the performance of algorithms. In [4], we argued for a *history of (un)successful implementations* that affords this interpretation. The idea is simple and intuitive: good practices with visible success –such as high accuracy, low margin of error, ease of implementation, and formal verification– tend to endure over time, while less successful practices tend to be eradicated.

To illustrate this token<sub>1</sub>-RI a bit further, consider *design prototyping*, a sub-field of software engineering that assists developers in assessing alternative design strategies and deciding which is best for a particular goal. Since there are no standard methods for choosing the best strategy, researchers need to compare the requirements of the algorithm with various design approaches to evaluate which one possesses the best characteristics for fulfilling the intended objectives. The example I used in previous publications is a computer simulation involving networking. For this, there are different topologies: ring, star, tree, and mesh. In order to pick the most suitable one, diverse performance characteristics need to be evaluated to see which topology is better at meeting performance goals and constraints [34, Chapter 5].

The same point can be made with an example closer to machine learning. Take the case of BenevolentAI presented earlier, which utilizes *Best First Search* (BFS), a search algorithm highly successful for navigating graphs and trees. The primary goal of BFS is to find the most promising path to a target node based on a given heuristic. In this respect, BFS has proven to be extremely effective for searching suitable drugs within BenevolentAI's knowledge graph [14, 604]. Classified under type<sub>1</sub>-RI *history of (un)successful implementations*, BFS contributes to the reliability of BenevolentAI and the justification of its outputs.

Likewise, past failures must be, and typically are, avoided by competent programmers. The history of computing is littered with cases of failed software that changed specification and coding practices. Therac-25 is one tragic case [52]. As reported, the algorithm used by Therac-25 was not thoroughly validated, and the testing process was insufficient to catch critical bugs that led to radiation overdoses. Furthermore, there was poor error handling and reporting in the software. Error messages were often cryptic, and operators were not adequately trained to understand and respond to them. Finally, there were no redundancy safety mechanisms that could ensure that software failures do not result in such catastrophic outcomes. From the perspective

---

<sup>9</sup> See Chapter 14 by Manganini and Primiero in this volume.

of CR, these all amount to diverse indicators of the unreliability of the algorithm used in Therac-25.

### 3.1.2 Type<sub>2</sub>-RI: Computer-based scientific practice

Assessing the technical performance of algorithms facilitates justification in terms of increasing accuracy, predictive power, low error rates, tolerance to domain change, and the ability to multi-purpose algorithms and data, among other factors. However, this assessment is silent on the adequacy of algorithms for scientific purposes. A reliabilist epistemology must offer standards by which the algorithm used in a scientific context can be warranted to a greater or lesser degree.

Let me illustrate these ideas with a familiar example. Wu and Zhang's automatic facial recognition system exemplifies how accuracy alone does not exhaust the reliability of an algorithm. As mentioned in section 2.1, the AUC-ROC measured 0.9540 predictive accuracy for their CNN. Such tremendous results cemented these researchers' confidence in the scientific merits of the algorithm. However, as discussed, there is no basis for such optimism. Criminality is a socially constructed concept that depends on diverse and sometimes contradictory interpretations of the socio-economic basis of criminality, psychological studies of criminals, and laws that determine when and to what degree someone is considered a criminal. Without reference to some of these concepts and frameworks, the prediction –however accurate– lacks the grounds for legitimate scientific claims.

Under CR, the reliability of algorithms is not exclusively assessed based on high predictive accuracy. Science involves more than just measuring and classifying algorithmic outputs. In this respect, I believe that algorithms cannot and should not operate in isolation from the broader context of scientific undertakings. We need to delve not only into standard non-algorithmic scientific practice, but also into a form of scientific practice that evolves with and heavily depends on algorithms. Type<sub>2</sub>-RI is an attempt to capture the family of token-RI connected to a larger body of scientific theories, beliefs, and practices within which algorithms are specified, coded, utilized, and maintained. In what follows, I lay out two potential candidates.

#### Expert knowledge

*Expert knowledge* is an umbrella term that covers the myriad of background education, knowledge, activities, training, virtues, and skills of researchers that bring to bear a broad range of talents to the specification, coding, use, and maintenance of algorithms in scientific contexts. Understood as a reliability indicator, *expert knowledge* reports on the many ways in which scientific expertise, technical expertise, and general competencies can be implemented into an algorithm.

To best understand this indicator, we must look at its various functions. For starters, it puts forward the algorithm's competencies, scope, and theoretical assumptions as conceptualized by the researchers involved in the specification, coding, maintenance,

and execution of the algorithm. It also accounts for the ability to describe a target system and its conditions for adequacy (e.g., to be applicable in a specific domain, to be representative of a particular condition, to be context-sensitive, to be repurposed). Expert knowledge covers social practices tailored to the development of algorithms, aptitudes to anticipate their merits intelligibly, and abilities of agents to manipulate them. For instance, setting up the variety of initial conditions, datasets, parameters and hyper-parameters (epochs, batch size, number of neurons, number of layers, dropout rate, etc.), all of which are complex yet critical for the performance and scientific merits of the algorithm. Consider determining which parameter to prioritize as a reliability indicator. Their selection and optimization are not trivial and yet fundamental for the general performance of algorithms (convergence of results, accuracy, overall performance) [54]. van Rijn and Hutter have conducted an informative experiment to show that the final performance metrics for deep learning models vary according to how different researchers select and optimize algorithmic parameters and instantiations [38]. Thus understood, experts contribute to the overall reliability of an algorithm by specifying relevant internal data-types, structures, relations, operations, and the like. They also credit reliability (or might identify instances of unreliability) by their pick and choose of datasets, parameters, and other variables.

As a reliability indicator, expert knowledge also attempts to accommodate the complexities of algorithms through the division of cognitive labor. Rather than being developed in isolation, algorithms involve a myriad of direct and indirect stakeholders (e.g., software engineers, physicians and chemists – in the case of BenevolentAI –, biologists – in the case of AlphaFold –, and psychologists and legal officers – in what should have been the case for Wu and Zhang). A core team specifies and codes algorithms utilizing ready-made computer modules others have coded. They employ measuring techniques others have designed, constructed, and calibrated. They analyze data using mathematical and statistical techniques others have validated. They make use of mathematical and computational methods others have devised and tested. There is no development of algorithms in solitude. Teams with diverse cognitive strengths and talents collectively collaborate in a variety of ways. Hence, the success or failure of algorithms is tailored to this collective knowledge, just as much as it depends on individual competencies. What one team member overlooks, another might notice. What one team member forgets, another might foresee. What one team member does not know how to solve, another might be able to teach. Thus diversified, the range of achievable solutions is far greater than what is available in atomized practices.

Interestingly, the role and value of experts are being increasingly recognized in philosophical studies on algorithms. Ratti and Graves [40] argued that documenting developers' motives and the code and specification of an ML are indicators of the reliability of the system. Newman [3] has argued along similar lines with respect to computer simulations. Newman considers the entire practice of software engineering to be at stake, from test plans to selecting programming languages and modeling tools, including configuration management.

While I am sympathetic to these ideas, my interpretation of expert knowledge is somewhat broader. It includes technical personnel with no training in software de-

velopment, practices that exceed software engineering standards, and accommodates the possibility that complete documentation of an algorithm is not always available.

In practice, non-technical personnel are intimately involved in algorithmic development (e.g., physicians and chemists in the case of BenevolentAI, and biologists and chemists in the case of AlphaFold), despite having little to no idea how key features of the system are specified and implemented. Their expertise is, however, crucial for the assessment of the reliability of the algorithm, and thus must be considered. Typically, their role is to inform, supervise, and sometimes even test the specification and coding of algorithms. But of course, these roles and interactions vary among cultures of research [47].

In connection with this, local practices and vernacular terminology often exceed what is captured by software engineering standards. Consider for instance how ML naturalizes or ‘fossilizes’ concepts. Once a concept is coded into the system, it is universally and indistinguishably applied across large and heterogeneous databases with varying degrees of success. Take the concept of ‘health’ as a case in point. One interpretation takes statistical measures and standards of normal biological measurements of someone’s body as the baseline for whether they are healthy. This concept of ‘health’ can be relatively straightforwardly implemented on an algorithm. However, the same concept also allows interpretations tailored to the diverse values of an individual or a community [17]. If a community considers blood transfusion to be harmful, they will treat any members of the community who have received a blood transfusion as unhealthy [49]. Implementing a cogent definition of health is no trivial matter.

Lastly, anyone who has written a piece of code knows all too well that not every line of code is documented. And even if algorithms were fully documented, this is no guarantee of understanding the code and its various functions. Thorough documentation – when it happens – and well-intended software engineering might still fall short of capturing the methodological and epistemological competencies of algorithms. We need to highlight the subtle interpretations, gentle disagreements, and non-verbal practices pervading computational and scientific practice and which make their way into the algorithm.

Let me finish by noticing that this reliability indicator brings about another important aspect of algorithms, namely, that they might only be *locally* reliable. The idiosyncrasies attached to documenting, specifying, coding, executing, and maintaining algorithms might make them only reliable in one context but not necessarily in another. This is, I believe, at the root of IBM’s Watson for Oncology’s difficulties of implementation in South Korea and Denmark, despite its success in the US market [30, 45]. Notoriously, Watson for Oncology was capable of analyzing large amounts of data and multiple variables, rendering accurate diagnoses and treatments for cancer patients in the US. But while IBM presents Watson for Oncology as offering more objective medical decisions and more accurate diagnoses than actual oncologists [28], it has been reported that many of these claims have been aggrandized [39, 37]. When implemented in South Korea and Denmark, only a fraction of the outputs rendered by the algorithm matched – or closely matched – the local clinician’s best diagnosis [35].

### Knowledge-based integration

Scientific results do not come in discrete bits, nor are the objects of scientific inquiry independently sanctioned. Instead, scientific theory and practice constitute a web of mutually supportive claims and commitments that are reached after complex negotiations in complex socio-economic and political environments. However, many studies utilizing algorithms portray a sanitized image of scientific research, where there is privileged access to structured data, undisputed model implementation, and meaningful representations of the world.

Wu and Zhang, for example, state that the quality of their databases and the methods implemented for data analysis prevent “the garbage of human biases from creeping in” [7, 2]. Given “race, gender and age, the faces of [the] general law-abiding public have a greater degree of resemblance compared with the faces of criminals” [7, 2]. It is, however, doubtful whether Wu and Zhang’s CNN has any scientific merits. One reason (to add to those previously mentioned) is that Wu and Zhang’s CNN is largely disconnected from accepted bodies of scientific knowledge. More precisely, the categories their CNN purports to use (i.e., {criminal} and {non-criminal}) are posited in isolation from established evidence, models of criminal psychology, social studies on crime, and the relevant theories on criminality. Thus, the CNN’s outputs are based solely on picking out facial traits from selected photos, rather than being premised on a larger body of knowledge implemented in the algorithm.

I will call approaches that conceive of algorithms as disconnected from the larger body of scientific knowledge *just a bunch of data analysis* (JBDA). By doing this, I intend to emphasize that an algorithm performing mere data analysis, but disconnected from concepts, theories, law-like principles, hypotheses, and other scientific units of analysis, is unlikely to merit scientific credentials, regardless of its predictive accuracy. To my mind, JBDA ignores the ‘bigger picture’ of knowledge integration, interpretation, and operationalization into algorithms. In fact, I consider JBDA as misleadingly portraying algorithms as an objective, unambiguous, and scientifically grounded examination of data that produces scientifically meaningful outputs. Nothing could be further from the truth. Wu and Zhang’s CNN approach is an archetypal JBDA, as it depicts scientific practice as granular, consisting of discrete pieces of information, separately secured and individually sanctioned. To these authors’ minds, “like most technologies, machine learning is neutral” [7, 2]. JBDA approaches advocate a form of scientific practice that is non-perspectival, socially disinterested, and impartial (i.e., epistemically and normatively neutral<sup>10</sup>), and disembodied from a larger corpus of scientific knowledge.

Are there instances where JBDA approaches are scientifically intelligible? I believe so. As suggested, there are indeed cases where mere data analysis renders valuable scientific insight about a subject matter. But for such cases, one needs to provide further justification that relates JBDA with a larger corpus of knowledge. A plausible interpretation of an account of scientific practice with algorithms capable of accommodating cases of JBDA takes the bulk of scientific knowledge and

---

<sup>10</sup> For a critical view on this perspective in the context of algorithms, see [8].

practices in the field under study as background knowledge and as affording sufficient grounds to underwrite particular claims made with the algorithm.<sup>11</sup> To briefly illustrate this idea, as this point encroaches on issues discussed under type<sub>3</sub>-RI (see section 3.1.3), consider BenevolentAI\*, an ML system whose working principles are JBDA. Suppose that BenevolentAI\* puts forward baricitinib\* as a drug with high chances of combating COVID-19 symptoms. Would researchers be justified in believing baricitinib\*? Surely not at face value, but only once it is embedded in a larger body of knowledge about COVID-19 and after some clinical trials show its scientific worth. What the examples of BenevolentAI and BenevolentAI\* show, I believe, is that we might still be justified under JBDA-like algorithms if (a) the algorithm – as a whole or as constituent parts – implements scientific models, theories, principles, categories, and/or other elements purposed in our corpus of scientific knowledge, or (b) its outputs are later assessed by the relevant community and within a corpus of scientific knowledge (more on this in section 3.1.3.)

Let me further illustrate this reliability indicator. Take again BenevolentAI, which utilizes information gathered from scientific research papers, structured and unstructured biomedical data, and drug and pharmaceutical industry data. BenevolentAI also implements knowledge graphs that structure data into causal relationships between known diseases, genes, environmental factors, and approved drugs [23]. Furthermore, BenevolentAI aligns with auxiliary assumptions, theories, and structures of drug molecular profiles, as well as mechanisms integral to the process of damaging healthy cells and tissues. It also incorporates theories about genetics, medical studies of disease, and biological models relating genes to drug effects. Through this knowledge-based integration, outputs produced by BenevolentAI are better justified than those by BenevolentAI\*. Indeed, this JBDA-like version does not implement any accepted model or concept into its algorithm, does not operationalize knowledge graphs, and does not represent mechanisms integral to knowledge about damaging healthy cells and tissues. It is the theoretical rigor, along with a history of (un)successful implementations [4], domain and expert competence, and possibly some skilled insight that leverages the reliability of BenevolentAI over and above BenevolentAI\*.<sup>12</sup>

### 3.1.3 Type<sub>3</sub>-RI: Social construction of reliability

The performance of algorithms (type<sub>1</sub>-RI) is undoubtedly required for justification. But while a necessary condition, it is certainly not sufficient. The paradigmatic example is Wu and Zhang’s CNN. This algorithm leverages a subset of RI<sub>1</sub> (most prominently, validation) but makes claims about an alleged law of facial recogni-

---

<sup>11</sup> Helen Meskhidze makes a similar claim using ML applications in astrophysics. The chapter uses “physics-informed machine learning” where physical laws and domain-specific knowledge implemented in the algorithm are crucial for its success. In terms of an epistemology of algorithms, my approach differs in that Meskhidze is interested in fostering transparency and interpretability (See Chapter 18 by Meskhidze in this volume).

<sup>12</sup> Thanks go to Emanuele Ratti for pressing on clarifying this point.



tion that is hard to accept. Expert and knowledge integration (type<sub>2</sub>-RI) are also fundamental to the reliability of algorithms. But again, necessary but not sufficient. BenevolentAI furnishes a good example. The algorithm is robust and built on a solid scientific basis. However, baricitinib was later flagged as counter-prescribed for immunocompromised patients. So, what is missing? To my mind, the reliability of algorithms must also be assessed within social processes that aim to achieve standards of scientific value and thresholds for acceptance of  $\bar{o}$ . Let me put the same idea in different form. The performance of algorithms along with expert knowledge and knowledge-based integration observe that the relevant scientific structures and processes, commitments and categories, entities and concepts are correctly implemented into, and computed by the algorithm. But justification of  $\bar{o}$  requires something else. We need to further seek for coherence and consistency of  $\bar{o}$  with a larger corpus of knowledge. In this way, we observe that the output is in agreement with accepted scientific commitments, standards of quality, evidence, and relevance, among other scientific qualifications.

Let me quickly illustrate how this reliability indicator would work using the story behind BenevolentAI. After announcing baricitinib, diverse groups within the scientific community began to debate the benefits –and dangers– of this drug. A major concern that emerged was that the mechanisms of action of baricitinib would block JAK-STAT signaling pathways (mainly mediated by JAK1 and JAK2), thus impairing interferon-mediated antiviral responses [21, 1013]. Blocking interferon would allow attack by other viruses (e.g., herpes zoster and herpes simplex) which in some cases may be more harmful than COVID-19. Favalli and colleagues [21] later reported on the potential harms of administering baricitinib to some patients, most importantly immunodeficient patients. As a response, Richardson and colleagues accepted the conditions under which BenevolentAI was a reliable indicator –and reasonably use this debate and incorporate new functions into the algorithm.

As a reliability indicator, this scientific debate regulated on how much evidence was required to accept BenevolentAI's outputs. It draw thresholds of which errors and artifacts can be tolerated, and to what extent. It also determined which assumptions are fit for purpose. Commitments to reliable algorithms are commitments to a network of scientific methodologies, standards, and traditions expressed through scientific debate. As Catherine Elgin points out, this network enables scientists to build on each other's work. They can be confident that justified outputs have the epistemic value their discipline prescribes [5, 77]. Of course, disputes and disagreements among community members are to be expected. There may be conflicts over values, methods, and what constitutes acceptable evidence. Take again Favalli's concerns about administering baricitinib to a specific group of patients. Whereas Richardson largely agreed with Favalli's concerns, research on the drug continued. Further laboratory testing confirmed Richardson's beliefs.

The social formation and justification of beliefs is a complex enterprise that is not always successful. There are situated background assumptions and perspectives. Scientific inquiry is permeated by contextual values and interests. Many concepts built into algorithms are socially constructed, generational, and discipline-idiosyncratic. Take again variations in the definition of 'health' and 'disease' [43, 22]. Each concept

operates under a myriad of cultural, political, economic, and moral values. Caruana and colleagues [12] discuss a neural network specified to predict pneumonia risk scores in patients and their readmission to hospital. Caruana et al. find that asthmatic patients are at low risk and thus less likely to require hospitalization than other patients (with chronic lung disease, for instance). Caruana et al.'s finding is statistically accurate (type<sub>1</sub>-RI and type<sub>2</sub>-RI indicators are present). However, the system is perceived as unreliable given that physicians have different starting assumptions about what a predictive algorithm should provide. According to Theunissen and Browning [?], physicians assume that the algorithm is predicting outcomes according to a shared baseline of care rather than differential care relative to the background of the patient. The ML outputs can be called into question because of oversimplifications in one or more assumptions in the system. This further shows that neither type<sub>1</sub>-RI nor type<sub>2</sub>-RI are individually—or jointly—sufficient for the reliability of algorithms.

CR makes an effort to foster belief-forming methods that accommodate social interventions, scientific scrutiny, and inter-domain justifications. Beliefs are justified in relation to a network of interconnected scientific beliefs. Yet, we cannot expect it to be recurrently successful. Contingent values and interests within the scientific community also find their way into the use and perception of the outputs of algorithms. Just like many other scientific methodologies, entrenching the reliability of algorithms requires a delicate balancing act. We bring together our best technical knowledge, theories, methodologies, and social skills. We do so in an attempt to justify believing that the output of a given ML system can be scientifically valuable, true, or can be epistemically trusted.

## 4 Final thoughts

I have presented CR as a reliabilist epistemology for the justification of algorithms' outputs. I also presented and discussed diverse types and token reliability indicators that form the basis of reliable algorithms. In sum, I set out to defend the following claim: we have—or increasingly have—justification for believing algorithms' outputs when they are rendered by a reliable belief-forming method. To be reliable here means that the algorithm is specified, coded, used, and maintained utilizing specific, tailor-made reliability indicators designed for the purpose at hand.

Admittedly, CR construes justification as inherently provisional. As discussed, reliability indicators might change over time and even be replaced. They might also be hard to measure or resolve conflicts. But I believe this is part of the self-critical and self-correcting endeavors that we find in scientific research. It might also be the best epistemic effort we can offer given a context and our limited resources. These activities constitute our best knowledge, metrics, methodologies, and practices, even if subjected to further scrutiny and revision.

I also listed a few issues with CR that I was unable to address but which might be considered conditional to its acceptability. I can accept that. But CR is a step in the right direction, if not as a more adequate epistemology of algorithms, at least as

an alternative to internalist epistemologies. In this regard, the chapter has achieved its goal of setting up an externalist epistemology of algorithms, a goal that should be appreciated in its own right.

## 5 Acknowledgments

This paper has been in the making for quite some time. This means that many people need to be acknowledged and thanked for their diverse comments and suggestions. Let me begin with my co-editor. Thank you, Giorgia, for the endless close readings of this paper. Thanks also go to Jack Casey for many interesting discussions around these topics and helping me to avoid several mistakes. Many other people deserve recognition: Federica Russo, Emanuele Ratti, Edoardo Datteri, Giuseppe Primiero, Viola Schiaffonati, Rawad El Skaf, Manuel Barrantes, Karin Jongma, Andrea Ferrario, Nico Formanek, Charles Rathkopf, and Atocha Aliseda. Some of them believe in CR, some of them don't. But they all have been supportive and encouraging of my ideas. Naturally, all wrongs—and rights—are mine. Finally, thanks go to Kass and Diego. Kass for being supportive beyond my comprehension. Thank you so much. And Diego, for teaching me an absolute truth: that everything is silly.

## References

1. Jumper, J., Evans, R., Pritzel, A., Green, T., Figurnov, M., Ronneberger, O., Tunyasuvunakool, K., Bates, R., Žídek, A., Potapenko, A., Bridgland, A., Meyer, C., Kohl, S., Ballard, A., Cowie, A., Romera-Paredes, B., Nikolov, S., Ain, R., Adler, J., Back, T., Petersen, S., Reiman, D., Clancy, E., Zielinski, M., Steinegger, M., Pacholska, M., Berghammer, T., Bodenstein, S., Silver, D., Vinyals, O., Senior, A., Kavukcuoglu, K., Kohli, P., Hassabis, D. (2021). Highly accurate protein structure prediction with AlphaFold. *Nature*, 596, 583-589.
2. Durán, J. M. (2020). What is a Simulation Model? *Minds and Machines*, 30, 301-323.
3. Newman, J. (2016). Epistemic Opacity, Confirmation Holism and Technical Debt: Computer Simulation in the Light of Empirical Software Engineering. In F. Gadducci & M. Tavosanis (Eds.), *History and Philosophy of Computing. HaPoC 2015. IFIP Advances in Information and Communication Technology* (Vol. 487, pp. 256-272). Springer.
4. Durán, J. M., & Formanek, N. (2018). Grounds for Trust: Essential Epistemic Opacity and Computational Reliabilism. *Minds and Machines*, 28(4), 645-666.
5. Elgin, C. Z. (1996). *Considered Judgement*. Princeton University Press.
6. Medeiros, J. (2021). How tech is changing healthcare. From rapid development and rollout of the Covid-19 vaccines to the science of isolation, machine-learning-enabled gene editing and digitised medicine. *Wired*. Retrieved from <https://www.wired.co.uk/article/future-health-trends>
7. Wu, X., & Zhang, X. (2017). Responses to Critiques on Machine Learning of Criminality Perceptions (Addendum of arXiv:1611.04135). *arXiv*, arXiv:1611.04135v3. <https://doi.org/10.48550/arXiv.1611.04135>
8. Pozzi, G., & Durán, J. M. (2024). Informativeness and epistemic injustice in explanatory medical machine learning. *AI & Society*.

9. Durán, J. M., & Jongsma, K. R. (2021). Who is afraid of black box algorithms? On the epistemological and ethical basis of trust in medical AI. *Journal of Medical Ethics*, 47(5), 329-335. <https://doi.org/10.1136/medethics-2020-106820>
10. Myszczyńska, M., Ojames, P., Lacoste, A., Neil, D., Saffari, A., Mead, R., Hautbergue, G., Holbrook, J., & Ferraiuolo, L. (2020). Applications of machine learning to diagnosis and treatment of neurodegenerative diseases. *Nature Reviews Neurology*, 16, 440-456. <https://doi.org/10.1038/s41582-020-0377-8>
11. Parker, W. S. (2020). Model Evaluation: An Adequacy-for-Purpose View. *Philosophy of Science*, 87, 457-467.
12. Caruana, R., Lou, Y., Gehrke, J., Koch, P., Sturm, M., & Elhadad, N. (2015). Intelligent Models for HealthCare: Predicting Pneumonia Risk and Hospital 30-Day Readmission. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 1721-1730). Association for Computing Machinery. <https://doi.org/10.1145/2783258.2788613>
13. Durán, J. M., van der Vloed, D., Ruifrok, A., & Ypma, R. J. F. (under review). From understanding to justifying: computational reliabilism for AI-based forensic evidence evaluation. *FSI Synergy*.
14. Segler, M. H. S., Preuss, M., & Waller, M. P. (2018). Planning chemical syntheses with deep neural networks and symbolic AI. *Nature*, 555(7698), 604-610. <https://doi.org/10.1038/nature25978>
15. Durán, J. M. (2021). Dissecting scientific explanation in AI (sXAI): A case for medicine and healthcare. *Artificial Intelligence*, 297, 103498. <https://doi.org/10.1016/j.artint.2021.103498>
16. Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). "Why Should I Trust You?": Explaining the Predictions of Any Classifier. In *KDD '16: Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 1135-1144).
17. Richman, K. A. (2004). *Ethics and the Metaphysics of Medicine. Reflections on Health and Beneficence*. MIT Press.
18. Richardson, P., Griffin, I., Tucker, C., Smith, D., Oechsle, O., Phelan, A., Rawling, M., Savory, E., & Stebbing, J. (2020). Baricitinib as potential treatment for 2019-nCoV acute respiratory disease. *The Lancet*, 395(10223), e30-e31. [https://doi.org/10.1016/S0140-6736\(20\)30304-4](https://doi.org/10.1016/S0140-6736(20)30304-4)
19. Goldman, A. (1979). What Is Justified Belief? *The Justification of Belief*, 105(9), 1-23.
20. Turk, M., & Pentland, A. (1991). Eigenfaces for recognition. *Journal of Cognitive Neuroscience*, 3(1), 71-86.
21. Favalli, E. G., Biggioggero, M., Maioli, G., & Caporali, R. (2020). Baricitinib for COVID-19: a suitable treatment? *The Lancet*, 20, 1012-1013.
22. Sisti, D., & Caplan, A. L. (2017). The concept of disease. In M. Solomon, J. R. Simon, & H. Kincaid (Eds.), *The Routledge Companion to Philosophy of Medicine* (pp. 5-15). Routledge.
23. Smith, D. P., Oechsle, O., Rawling, M. J., Savory, E., Lacoste, A. M. B., & Richardson, P. J. (2021). Expert-Augmented Computational Drug Repurposing Identified Baricitinib as a Treatment for COVID-19. *Frontiers in Pharmacology*, 12, 709856. <https://doi.org/10.3389/fphar.2021.709856>
24. Humphreys, P. (2021). Epistemic Opacity and Epistemic Inaccessibility. *Pre-Print*.
25. Morales-Hernández, A., Van Nieuwenhuysse, I., & Rojas González, S. (2022). A survey on multi-objective hyperparameter optimization algorithms for machine learning. *Artificial Intelligence Review*. <https://doi.org/10.1007/s10462-022-10359-2>
26. Creel, K. A. (2020). Transparency in Complex Computational Systems. *Philosophy of Science*, 87(4), 568-589. <https://doi.org/10.1086/709729>
27. Alvarado, R., & Humphreys, P. (2017). Big Data, Thick Mediation, and Representational Opacity. *New Literary History*, 48(4), 729-749. <https://doi.org/10.1353/nlh.2017.0037>
28. Swelitz, I. (2016). Watson goes to Asia: hospitals use supercomputer for cancer treatment. *Statnews*. Retrieved from <https://www.statnews.com/2016/08/19/ibm-watson-cancer-asia/>
29. Douglas, H. (2004). The irreducible complexity of objectivity. *Synthese*, 138, 453-473.
30. Vulsteke, C., Ortega Arevalo, M., Mouton, C., Stam, K., Goethals, R., Ameye, F., Populaire, C., Peeters, M., & Verdonck, P. (2018). Artificial intelligence for the oncologist: hype, hubris, or reality? *Belgian Journal of Medical Oncology*, 12(7), 330-333.

31. Blanz, V., & Vetter, T. (1999). A morphable model for the synthesis of 3D faces. In *Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques* (pp. 187-194). ACM Press/Addison-Wesley Publishing Co.
32. Fagiolo, G., Moneta, A., & Windrum, P. (2007). A Critical Guide to Empirical Validation of Agent-Based Models in Economics: Methodologies, Procedures, and Open Problems. *Computational Economics*, 30, 195-226.
33. Wachter, S., Mittelstadt, B., & Russell, C. (2018). Counterfactual explanations without opening the black box: automated decisions and the GDPR. *Harvard Journal of Law and Technology*, 31(2), 841-887.
34. Pfleeger, S. L., & Atlee, J. M. (2009). *Software Engineering: Theory and Practice* (4th ed.). Pearson.
35. Hamilton, J. G., Genoff Garzon, M., Westerman, J. S., Shuk, E., Hay, J. L., Walters, C., Elkin, E., Bertelsen, C., Cho, J., Daly, B., & others. (2019). "A tool, not a crutch": patient perspectives about IBM Watson for oncology trained by Memorial Sloan Kettering. *Journal of Oncology Practice*, 15(4), e277-e288.
36. Fetzer, J. H. (1998). Program Verification: The Very Idea. *Communications of the ACM*, 37(9), 1048-1063.
37. Ross, C., & Swetlitz, I. (2018). IBM's Watson supercomputer recommended 'unsafe and incorrect' cancer treatments, internal documents show. *Statnews*. Retrieved from <https://www.statnews.com/wp-content/uploads/2018/09/IBMs-Watson-recommended-unsafe-and-incorrect-cancer-treatments-STAT.pdf>
38. van Rijn, J. N., & Hutter, F. (2018). Hyperparameter Importance Across Datasets. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining* (pp. 2367-2376). ACM.
39. Ross, C., & Swetlitz, I. (2017). IBM pitched its Watson supercomputer as a revolution in cancer care. It's nowhere close. *Statnews*. Retrieved from <https://www.statnews.com/2017/09/05/watson-ibm-cancer/>
40. Ratti, E., & Graves, M. (2022). Explainable machine learning practices: opening another black box for reliable medical AI. *AI and Ethics*. <https://doi.org/10.1007/s43681-022-00141-z>
41. Lorscheid, I., Heine, B.-O., & Meyer, M. (2012). Opening the 'black box' of simulations: increased transparency and effective communication through the systematic design of experiments. *Computational and Mathematical Organization Theory*, 18, 22-62.
42. Durán, J. M. (2013). Explaining Simulated Phenomena: A defense of the epistemic power of computer simulations (PhD thesis). Institut für Philosophie, Universität Stuttgart. Retrieved from [https://elib.uni-stuttgart.de/bitstream/11682/5409/1/Thesis\\_Duran.pdf](https://elib.uni-stuttgart.de/bitstream/11682/5409/1/Thesis_Duran.pdf)
43. Boorse, C. (2011). Concepts of Health and Disease. In F. Gifford (Ed.), *Handbook of the Philosophy of Science* (pp. 13-64). Elsevier.
44. Wu, X., & Zhang, X. (2016). Automated Inference on Criminality using Face Images. *arXiv*, arXiv: 1611.04135v1. <https://arxiv.org/pdf/1611.04135v1>
45. Emani, S., Rui, A., Rocha, H. A. L., Rizvi, R. F., Juacaba, S. F., Jackson, G. P., & Bates, D. W. (2022). Physicians' Perceptions of and Satisfaction With Artificial Intelligence in Cancer Treatment: A Clinical Decision Support System Experience and Implications for Low-Middle-Income Countries. *JMIR cancer*, 8(2), e31461. <https://doi.org/10.2196/31461>
46. Elgin, C. (2017). *True Enough*. MIT Press.
47. Sundberg, M. (2010). Cultures of simulations vs. cultures of calculations? The development of simulation practices in meteorology and astrophysics. *Studies in History and Philosophy of Modern Physics*, 273-281.
48. Massimi, M., & McCoy, C. D. (Eds.). (2020). *Understanding Perspectivism. Scientific Challenges and Methodological Prospects*. Routledge.
49. Richman, K. A., & Budson, A. E. (2000). Health of organisms and health of persons: An embedded instrumentalist approach. *Theoretical Medicine and Bioethics*, 21(4), 339-352.
50. Goldman, A. I. (2012). *Reliabilism and Contemporary Epistemology*. Oxford University Press.
51. Humphreys, P. W. (2009). The Philosophical Novelty of Computer Simulation Methods. *Synthese*, 169(3), 615-626.

52. Leveson, N. G., & Turner, C. S. (1993). An investigation of the Therac-25 accidents. *Computer*, 26(7), 18-41.
53. Durán, J. M. (2018). *Computer simulations in science and engineering. Concepts - Practices - Perspectives*. Springer.
54. Hutter, F., Hoos, H., & Leyton-Brown. (2014). An efficient approach for assessing hyperparameter importance. In E. P. Xing & T. Jebara (Eds.), *Proceedings of the 31st International Conference on Machine Learning* (Vol. 32(1), pp. 754–762). PMLR.
55. Humphreys, P. W. (2004). *Extending Ourselves: Computational Science, Empiricism, and Scientific Method*. Oxford University Press.
56. Clark, A. (1987). The kludge in the machine. *Mind and Language*, 2(4), 277-300.
57. Philip Kitcher, *The Advancement of Science: Science Without Legend, Objectivity Without Illusions*, Oxford University Press, New York, 1993.
58. Oberkamp, W. L. and Roy, C. J. (2010) *Verification and validation in scientific computing*. Cambridge University Press.
59. Alston, William P. (1995) How to Think About Reliability. *Philosophical Topics*, 23(1):1–29.