



Models of Medical Reasoning

7

Margherita Benzi

*“You will not apply my precept,”
he said, shaking his head
“How many times have I said to you that
when you have eliminated the impossible
whatever remains, however improbable,
must be the truth?”*

Arthur Conan Doyle

Summary

In this paper, I focus on reasoning about individual cases in medicine. The relevance of the individual case has always been present both in medical and clinical research and has acquired further significance in recent times. Here, after recalling the importance of reasoning about individual cases, I introduce the distinction between general and singular causation and explain why medical reasoning needs the latter. I then consider three areas in which this necessity is most apparent: diagnosis, early phase clinical trials in oncology, and forensic medicine. I argue that the best approach to this form of causal reasoning is the counterfactual approach provided by structural equation models. I conclude by presenting some considerations on the relation between the building of structural models for singular causation and the contemporary “data deluge.” The aim of this paper is mainly philosophical: it is an attempt to interpret some actual medical approaches in a new perspective and at the same time connect some aspects of contemporary trends of research in medicine and bioinformatics with

M. Benzi (✉)


Department of Law and Political, Economic and Social Sciences, University of Piemonte Orientale, Alessandria, Italy
e-mail: margherita.benzi@uniupo.it

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2023
N. Rezaei (ed.), *Brain, Decision Making and Mental Health*, Integrated Science 12,
https://doi.org/10.1007/978-3-031-15959-6_7

115

some classical philosophical questions. I do not give practical suggestions, but I am also persuaded that philosophical work on kinds of reasoning and on singular causation can be of interest for practicing physicians and experts in other fields.

Graphical Abstract/Art Performance



Digital Twins.
(Reproduced with permission from Getty Images).

Keywords

Causal reasoning • Diagnosis • Drug testing • Forensic medicine • Singular causation • Structural equation models

QR Code



Scanning the QR code directs you to the word cloud of the chapter that is made up of the words we discussed throughout the whole book in relation to the chapter's keyword, which is intentionally not included in the word cloud. Find the keyword and put it in place in the puzzle according to the clues provided in Chap. 32. The mystery hidden in the puzzle is the quote of *Brain, Decision-making, and Mental Health*.

1 Introduction

Medical science differs from most other sciences for the importance assigned to the individual case. According to Gorovitz and MacIntyre [1], medical science is a “science of particulars.” In their account, the need for a particularistic approach derives from the fact that some particular entities, like the higher primates, cannot be described solely in terms of the relevant physical and chemical mechanisms. How those mechanisms operate results from the “particular history” of each individual and from the circumstances that accompany it. It is the importance of the particular history of each individual that makes it difficult to draw a line between a rule and its exceptions; in medicine, it can be said that exceptions are the rule.

The relevant changes in biomedical reasoning and decision-making linked to the advances in information technology seem to point towards a kind of medical research and practice even more “individual-centered.” The *Precision Medicine Initiative* [2], a research program launched in 2015 by former United States President Barack Obama, is based on the very notion of “tailored health care.”

Until now, most medical treatments have been designed for the “average patient.” As a result of this “one-size-fits-all” approach, treatments can be very successful for some patients but not for others. Precision medicine, on the other hand, is an innovative approach that takes into account individual differences in people's genes, environments, and lifestyles.

The general aim of precision medicine, also known as personalized medicine, is that of finding prevention and treatment strategies respectful of—or even based on—individual traits [3]. In describing the program, Collins and Varmus [4] remark that these ideas are not new, but their clarification and the very possibility of realization is due to:

the recent development of large-scale biologic database (such as the human genome sequence), powerful methods for characterizing patients (such as proteomics, metabolomics, genomics, diverse cellular assays, and even mobile health technology), and computational tools for analyzing large sets of data.

This new trend of research and healthcare is, therefore, based on the possibility of assembling and analyzing huge amounts of data, including behavioral, physiological, and environmental parameters.

In the precision medicine program, both the representation of individual patients and their treatment are obtained through large databases, where data can be both relative to the individuals and very large classes of patients. The epistemological implications and some of the perils of “the data turn” have been tackled by scientists and philosophers of science. Here I do not intend to elaborate a direct approach to these kinds of questions but to introduce a way of representing individual cases and reasoning about them that *is* not data-driven, even if it can point to a principled way of *using* large databases. The proposed approach concerns a class of problems that are extremely important for decision-making in medicine: problems of singular causation.

2 General and Singular Causation

Philosophers distinguish between *general* (or *type*) causal statements, as in “Smoking causes lung cancer,” and *singular* (or *single-case*, *token*, or *actual*) causal statements,¹ as in “The sore throat of this patient has been provoked by an infection of the bacterium *Streptococcus pyogenes*,” or also “If this patient had not smoked, the probability of her getting lung cancer would have been lower.” The relation between the two levels of causation is controversial, both from the epistemological and the ontological point of view.

According to generalist theories, single-case causal claims are inferred from general causal laws or regularities, which can be deterministic or probabilistic. From a metaphysical point of view, individual cases of causation should be seen as instances of causal relations holding between types of events. From this perspective, when we say that a particular episode of the sore throat condition has been caused by *Streptococcus pyogenes* (*S. pyogenes*), this happens because we know that there is a general law linking *S. pyogenes* infection and sore throat, and we are observing a particular exemplification of that general law. General laws and regularities can be deterministic or probabilistic.

On the other hand, according to singularist theories of causation, knowledge about the truth of general causal claims is obtained by generalizing from individual cases of causation. *Metaphysical* singularist theories of causation see singular causal relations as primitive and general causal relations as generalizations from the former. Singular causal relations can hold even if they are not instances of regularities.

A third philosophical approach to causation [5, p. 6] considers singular causation and general causation as independent:

The problems of type and token probabilistic causation turn out to be quite distinct. Causal claims made on one of the two levels of causation turn out to be quite independent of claims

¹ Philosophers sometimes distinguish between *singular* and *actual*. For our purpose here we may set aside the distinction.

made on the other and the two levels of probabilistic causation require quite different kinds of theories.

Many contemporary theories of singular causation are based on counterfactual definitions of causation, the most basic of which is the following:

If c and e are actual distinct events, then c causes e if and only if, c had not been, then e would never have occurred.

Decades of philosophical research on the counterfactual interpretation of causation have highlighted several problems deriving from this definition and produced noticeably sophisticated versions of it, but for the moment, we adopt it as a working definition. Counterfactual definitions, as introduced by Lewis [6], help us to see causes as “difference-makers,” where the difference made by a cause is “a difference from what would have happened without it” (see also Menzies and Beebe [7]).

Besides distinguishing between singular and general causation, philosophers have identified and analyzed different notions of cause. Although these are of considerable importance in the philosophy of medicine today, I will not discuss them here in detail—but see [8, 9]. However, it should be recalled that many authors defend a pluralistic position, recognizing that different theories of causation are acceptable in different contexts of application.

3 Singular Causation in Medicine: Why Statistical Probability Might not Be the Right Choice

Much causal research in the field of medicine concerns the discovery of relations of general causation: they investigate, for instance, the problem of whether the exposure to a certain substance or chemical could cause a certain condition or whether a particular substance may contribute to preventing a certain disease. The inference from causal regularities to singular statements is not direct and not problem-free. Moreover, the kind of causal queries mentioned above concern the search for *the effects of causes*: they have the form “Does hemlock cause death?” The forms of causal reasoning we are examining in this section are different because:

- i. they are about causal relations holding between specific events; and
- ii. the causal query is about *the causes of effects*, as in “Was Socrates death caused by his drinking hemlock?”

Dawid et al. [10] analyze the gap between assessing the effects of causes and assessing the causes of effects:

While much of science is concerned with the effects of causes (EoC), relying upon evidence accumulated from randomized controlled experiments and observational studies, the problem of inferring the causes of effects (CoE) requires its own framing and possibly different data [...] The statistical literature is only of limited help here as well, focusing largely on the traditional problem of the EoC.

According to the authors, assessing the CoE is the typical inferential activity of legal courtrooms. In the following, I show how their considerations also apply to medical diagnosis, early-phase cancer clinical trials, and causal inference in legal medicine. Our point is that:

- i. in these forms of reasoning, what we are looking for is not the most (statistically) probable result, but the best hypothesis that explains the situation under study; and
- ii. when the inquiry has a causal character, it involves the notion of singular causation.

3.1 Diagnosis

Typically, the activity of obtaining a diagnosis is often a multi-step procedure, which involves a plurality of inferential techniques. Once a diagnosis is obtained, however, it can be sketchily described as the relation between a set of data relative to a given patient (such as clinical signs, reported symptoms, and laboratory findings) and a possible syndrome, disease, or other pathological condition. Many authors [11–14] agree that this relation between the data relative to a given patient and the diagnostic response is an *explanatory* one: typically, diagnoses link the patient’s data with the abnormal situation that could explain them. Not all diagnoses are *causal* explanations—many of them are merely classificatory, and it is controversial whether classifications are explanations—but here, we will focus on those that have a causal character and on the inferential procedures on which they are based.

In the classical account of explanation in terms of Hempel’s ‘covering law models’ [15], the explanation is viewed as a deductive inference from deterministic laws and initial conditions to the phenomenon to be explained—called *explanandum*—or as the inductive inference from statistical laws and initial conditions to a highly plausible hypothesis. Given that, except for pathognomonic signs and symptoms, the relation between signs and symptoms and the disease that could have caused them is *probabilistic*, we may be tempted to view causal diagnoses as inferences of the second kind, namely as statistical-inductive explanations. The example of sore throat can be represented in a standard way as follows:

- Most observed patients with (a certain kind of) sore throat had been infected by *Streptococcus Pyogenes*
 - The patient x has (a certain kind of) sore throat
- [makes very likely]
- x has been infected by *Streptococcus Pyogenes*

In this approach, the procedure of formulation of diagnosis could be seen as the procedure of finding the ‘right’ statistical laws that confer a high inductive

probability to a certain diagnostic hypothesis. However, to model the diagnosis as a statistical-inductive explanation is highly problematic. It is not difficult to find counterexamples that conform to the model but are not genuinely explanatory (and have a poor diagnostic value). Maung [13, p. 47] analyzes the example of a patient with right-sided paralysis due to cerebral palsy who subsequently has a left hemispheric stroke:

In this case, the diagnosis of left hemispheric stroke is explanatorily irrelevant to the patient's right-sided paralysis. Nonetheless, according to the covering law account, the patient's right-sided paralysis would still be explained by his or her diagnosis of left hemispheric stroke, along with the statistical generalization that a large proportion of patients diagnosed with left hemispheric stroke present with right-sided paralysis.

This counterexample is meant to show that the explanations provided by diagnoses should be, in many cases, *explicitly causal* explanations: at least in some cases, causality is necessary for the correctness of the diagnosis.² But this requirement is not enough. Given two alternative diagnostic hypotheses, H_1 and H_2 , that provide a causal explanation of a given set of signs and symptoms, the choice of the most probable one does not always guarantee that the better choice has been made: it is always possible that the patient under examination is an instance of the less probable alternative. The old recommendation to medical students [16]: "When you hear hoof-beats, think horses, not zebras" loses its utility when the approaching quadruped *is* a zebra.

A possible defense of the regularistic approach would consist in the suggestion that a suitable specification of the patient's reference class would lead to the "right" causal law. It might well be that H_2 is the less probable hypothesis for the whole population, but the more probable for the subpopulation of patients relevantly similar to the one under examination. In this case, adding qualifications such as age, gender, ethnicity, lifestyle, etc., should bring about, in the end, the right reference class and the right causal regularity. However, as Maung [13] remarks, this strategy raises two counter-objections of considerable impact [17, 18]. The first is that a description so detailed as to identify the right reference class could be impossible since there could be unknown contingent factors, which we cannot account for due to our ignorance. The second is that even we were able to provide a fully detailed description of the relevant traits of the clinical case, it would require placing it in a reference class so detailed as to make the appeal to general laws inappropriate.

The above considerations support the idea that regularities and laws (deterministic or probabilistic), although important, do not do all the explanatory work in causal, explanatory diagnosis [11]. At the end of the day, what matters is the *true* causal history of the patient's condition, not the causal explanation to which the statistical data assigns the greatest probability. Moreover, there are some cases in which statistical data are scarce, as in rare diseases or non-existent, as in "new"

² This requirement should protect diagnoses also from other counterexamples raised by spurious correlation.

diseases. The problem of the scarcity of data is even more important in the area of research we are going to consider now.³

3.2 Early Phase Clinical Trials

For more than a decade, philosophers of science have extensively investigated the epistemological foundations of clinical trials. One of the major issues at stake is whether the experimental clinical trials are a reliable tool for the ascertainment of the causal link between the administration of a drug and its intended therapeutic effects [19]. Another field that revealed itself of significant interest for the philosophy of causality is the study of possible adverse drug reactions [20]. This area of inquiry is often cited as one of those which will benefit the most from the data revolution, given that technological devices will allow better monitoring of the effects of drugs after they have been commercialized and, therefore, a better knowledge of the adverse effects associated to their use.

However, this is far from obvious for some kinds of clinical trials. An example is phase 1 clinical trials for testing oncology medication. These trials are usually performed on relatively small groups of cancer patients whose disease has become resistant to standard treatment. The trials aim to establish, among other things, the safety of a new drug. When an adverse event occurs, one of the main problems that experimenters face consists in ascertaining whether the event has been produced by the new drug. In phase 1 oncology clinical trials, the detection of the causal history of a reported adverse reaction is made particularly challenging by the comorbidity often associated with cancer, the heterogeneity of tumors, and the fact that often patients take many other drugs. Given the particular setting of early phase trials, it is unavoidable to make relevant causal judgments when the number of events is relatively low and when we have no access to control data.

The search for an objective, standardized procedure to evaluate the causal role of the drug for observed adverse reactions has produced a plurality of methods, none of which is entirely satisfactory [21]. Aronson and Hauben [22] show that in some cases, anecdotal evidence can be extremely reliable and defend the legitimacy of considering single case reports in assessing adverse drug reactions. They note that, although in general anecdotal reports are not considered good quality evidence, in some cases they provide definitive evidence of the fact that a certain drug caused an adverse event. The problem is finding general criteria to evaluate this type of evidence.

As in the case of diagnosis examined before, the actual causal history of the adverse reaction seems to be more valuable than the appeal to causal laws or regularities. Furthermore, it can be said that in this scenario, there are no “causal regularities” to appeal to because the possible causal nexus between the drug and the adverse event is the object of inquiry. The point I would like to stress here is that, although general causal knowledge is very important in this, as in every other

³ I am grateful to Mattia Andreoletti for suggestions and information on this topic.

scientific context, we should see the problem of causal assessment of adverse drug reactions in phase 1 oncology clinical trials from a singularist perspective. By this, I mean to say that, as in the case of clinical diagnosis, in this kind of clinical trial, the individual traits of the patients are crucial: evidence about putative adverse effects in just one patient can be extremely relevant, as it could reveal the fact that the drug is not safe. Given the presence of an adverse event e for a certain patient, what we want to know is whether taking the drug, alone or in conjunction with other factors, *actually* caused e . Our ideal goal is to reconstruct the true “causal history” that brought about e .

This reconstruction of the causal history of the adverse event can be highly uncertain, and we may end up with only a probable answer. However, it should be remarked that in this phase of the trial, our primary (ideal) target is *not* to determine *the probability* that the drug brings about some adverse reaction. Rather, what we want to discover is instead if there is a deterministic causal history in which the drug is the ‘culprit.’ And this brings us to the third area: forensic medicine.

3.3 Forensic Medicine

Singular causation is extremely important in forensic medicine, particularly in determining causes of death and in assessing legal responsibility. The highly individual character of forensic medicine is also well described by Russo and Williamson [23, p. 56] with the following considerations on forensic autopsy:

In determining the cause of death, the pathologist is also asked to determine the *manner* in which death occurred by identifying the series of events which led to death. Most jurisdictions include the following ‘manners of death’: homicide, accidental, natural, suicide, and undetermined. In forensic autopsies, examination may also serve to collect further evidence relevant to the crime investigated [...].

Even where there is no crime involved, legal reasoning shows a high degree of individualization. In the following real-world example in which the actual cause is not the most frequent one, Messina [24] reports on the death of a fourteen-year-old girl during an asthma attack. Here, the most probable cause of death was the asthma attack itself, even if this event is not frequent in childhood and adolescence. However, the death appeared somewhat dubious to the coroner, and an autopsy was required. The autopsy produced a totally different explanation for the death of the girl: even if the death occurred during an asthma attack, it was, with all probability, not caused by it. The autopsy showed that the girl had “silent” heart disease. This fact, and the suddenness of her death, was compatible with the fact that the death was due to a cardiovascular side effect of inhaling a bronchodilator medicine during the asthma attack (a rare adverse effect of β_2 -agonists).

To select this explanation, the coroner had to exclude many alternative possibilities: his final verdict was obtained through an inference to the best explanation, given the individual characteristics of the patient, the modality of death, and the exclusion of many other possible causal chains.

4 Bottom-Up Reasoning

Our remarks are in no way intended to deny the importance of causal regularities in diagnostic causal explanations; rather, the problem is how to use them when such explanations are not statistical (and are not based on a progressive restriction of the reference classes). One possible solution to this problem is in terms of mechanistic explanation [13, 23].⁴ Here we follow an alternative—but not incompatible—approach based on causal models. Our choice is based on some common characteristics of the kind of reasoning under discussion.

4.1 Individual Causal Explanations as “Stories”

A common feature of the kinds of reasoning we have presented is that explanation requires the reconstruction of the causal chain leading to the event under causal inquiry. This feature of explanation of individual cases has been stressed both by authors working in medicine and philosophy of science. In the introduction, we mentioned the central place assigned by Gorovitz and MacIntyre [1] to “particular histories” in medical science. Rizzi and Pedersen [25] see a causal medical explanation as to the identification of the “train of events” involving the singular case, while Thagard [26, p. 74] describes medical explanation as “a kind of narrative explanation of why a person gets sick.” Ankeny [27, p. 1006] gives a detailed account of how cases provide evidence of causal relations in clinical medicine and focuses on case reports, which typically have a highly narrative structure. The causal explanation of the conditions of a certain patient is seen as the “development of a detailed narrative that outlines various putative causes and systematically excludes certain ones as irrelevant based on additional information.” However, the analogy between stories and individual explanations is not only due to their narrative structure but also to the kind of reasoning involved. A short reminder on this point will help present a traditional connection between causal explanation and abduction.

4.2 Detective Stories and (Holmesian) Abduction

Parallel to the idea that causal reasoning about single cases makes substantial use of stories is the view that these stories are of a well-defined type: detective stories. Often, to explain the actual condition of a particular “case,” we need to ascertain the true chain of events that has determined it. This can be obtained by formulating a set of alternative hypothetical “stories,” rejecting them as new data, and refuting them until only one is selected. The reference to detective stories is explicit in

⁴ “The mechanistic malfunction conception involves lying out the details of a normal physiological mechanism and depicting the pathology as an impairment of this normal mechanism” [13, p. 55].

Ankeny [27]. More generally, Thagard [26, pp. 71–72] points out the analogy between medical reasoning about cases and legal reasoning:

Inference to medical causes is similar to legal inference concerning responsibility for crimes. In a murder case, for example, the acceptability of the hypothesis that someone is the murderer depends on how well that hypothesis explains the evidence, and on the presence of a motive that would provide a higher explanation of why the accused committed the murder.

Aronson and Hauben [22], in their defense of the importance of the single case reports as evidence in early phase clinical trials, elaborate on what they call “the crime scene metaphor.” The parallel between the search for the murderer and the search for the cause of an adverse reaction is developed by means of a list of *topoi* from detective fiction. They see an analogy with the *culprit caught at the scene of the crime* when:

objective physiochemical testing shows that a pathological lesion is composed of the drug or metabolite [...] and the event must not have been possible in the absence of the drug [22, p. 1267].

In such cases, where there are no possible confounders, we can immediately identify the drug as the cause, or at least as one of the main contributing causes of the lesion. The crime scene analogy is the *culprit seen committing the crime* when we can specify the pattern of a lesion or its location in a sufficiently precise way to decide, without further analysis, that the drug has been the cause. We can see an analogy with the *culprit incriminated by recreating the scene of the crime* when confirmatory tests are “ethically and scientifically feasible” [22, p. 1268]. Finally, the authors cite the “adverse drug reactions related to infections” as a case of “*Culprit DNA found at the scene of crime*” [ibid.].

The kind of reasoning examined in the above cases is often characterized as abductive reasoning: based on the available evidence, a hypothesis is inferred that, if true, would explain the evidence [28, 29]. The more appropriate characterization is that of *Holmesian inference* [30], where, in the ideal case, we consider a plurality of hypotheses, and all the alternative hypotheses are discarded but the true one. The application of ‘legal’ reasoning should be apparent in the field of forensic medicine.

5 Network Representations

Several authors believe that networks provide the best graphical representation of the multiplicity of possible causal chains containing the one that explains the event under inquiry. However, the choice of networks as tools for representing causal relations is linked to different theoretical views of causality and causal inference. Rizzi and Pedersen [25, p. 238] argue that the:

search for causal factors in the singular case is performed by using, to a certain extent, knowledge of the general causation. One applies or compares the body of knowledge of general causation to the situation where the singular case is assessed.

The knowledge of general causal relations can be represented by a network model, showing the multiple causal pathways that could have produced the particular condition of a given patient. The network model represents, therefore,

(all) possible causal factors which are known or believed to be potentially able to lead to the effect.

It is also conceived as a useful instrument for reasoning about singular causation: causal reasoning about a singular case can be performed by comparing pieces of the historical evidence relative to that case with factors contained in the reference case:

when seeking the actual causes in a singular case, backtracking in these trees eventually results in the selection of one strand – the train of events regarded as valid causal complexes for the singular case.

A further network model is presented by Thagard [26, p. 61]. Analogously to Rizzi and Pedersen, Thagard proposes an account of medical explanation

as a causal network instantiation, where a causal network describes the interactions among multiple factors, and instantiation consists of the observational or hypothetical assignment of factors to the patient whose disease is being explained.

The former proposals of network models suggest how to use background knowledge (such as pathophysiological knowledge, knowledge of general laws and mechanisms, etc.) to represent a multiplicity of possible causal pathways within which to trace individual causal paths.

Among the methods using network representation of causal structure for causal reasoning, the more well-known are causal Bayesian networks [31] and the structural equation models (SEM) approach [32–35]. As it is well known, the two approaches are strictly related and often collected under the label of graphical causal models. Both can be used to represent general or singular causal relations. First computational approaches using graphical models (DAGs) for medical diagnosis were based on Bayesian networks [36, 37]. Today, Bayesian networks are an important area of machine learning dedicated to the discovery of causal structures from observational data.

SEM approach offers a powerful framework for solving problems of actual causation, as it allows modeling counterfactual reasoning, allows a detailed representation of very specific contexts of occurrence,⁵ and facilitates reasoning from effects to causes.⁶ These reasons make the SEM approach particularly valuable for the kind of problems we are treating here [38, 39]. In the next section, we give a sketchy introduction to the use of formalism applied to the problems of adverse reactions.

⁵ “The basic idea is to extend the basic notion of counterfactual dependency to allow ‘contingent dependency’. In other words, while effects may not always counterfactually depend on their causes in the actual situation, they do depend on them under certain contingencies” [33, p. 844].

⁶ “Our notion is more appropriate for a *retrospective* notion of causality: given all the information relevant to a given scenario, was $X = x$ the actual cause of $Y = y$ in that scenario?” [33, p. 846].

6 Structural Equation Models

An SEM is a formal representation of the domain with respect to which we ask causal queries. The components of an SEM are a set V of *variables* and a set S of *structural equations* [35], specifying the relations of functional dependence among variables. *Variables* represent causal relata and can take any value from a specified range of *binary variables* that range over the values $[0, 1]$, where the two values represent, respectively, the occurrence or non-occurrence of events. The set of *equations* in *SEM* represents the causal structure of the model: linking events represented by variables on the left-hand to their direct causes (called “parents”), represented on the right side of the equation.

The formalism of structural equations enables reasoning about what would have happened if the parent of a certain variable had been assigned a different value. For example, asking whether the administration of a certain new drug to a certain patient caused the adverse effect E is to ask whether if *the drug* had not been administered, E would not have occurred or would have occurred in a completely different way. We should therefore ask the question “What would have happened if the situation had been different?” and reason counterfactually.

To assess whether a counterfactual statement is true in a model, we consider a modified version of the model in which:

- i. we ‘cut’ all the links to the variables in the antecedent of the counterfactual; and
- ii. we assign the values specified in the antecedent to these variables. The counterfactual is true if and only if the consequent is true in the modified model.

We can graphically represent the salient causal relations among variables employing DAGs, *directed acyclic graphs*, whose nodes represent variables, and edges represent direct causal dependencies of each variable from its parents as stated by the structural equations. In the following example, we present the occurrence of an adverse event (E) in a certain patient simultaneously taking two drugs B and D , each of which is individually sufficient to bring about the adverse reaction of that patient.

Variables: E (Patient experiences the adverse event), B (Patient takes the drug B), D (Patient takes the drug D),

Events:

$E = 1$ if the patient experiences the adverse event, $E = 0$ otherwise.

$B = 1$ if the patient takes the drug B , $B = 0$ otherwise.

$D = 1$ if the patient takes the drug D , $D = 0$ otherwise.

Structural equations:

(1) B .

(2) D .

(3) $E = \max (B, D)$.

The corresponding graph is shown in Fig. 1. If both B and D take the value 1, we incur in a typical situation of symmetrical overdetermination, which causes problems for the ‘plain’ counterfactual definition of causation: had drug B not been taken by the patient, the adverse event would still have occurred due to the assumption of drug D, and vice versa. Therefore, to assess the causal influence of each drug on the occurrence of the adverse event, we must evaluate our counterfactuals in modified models where only one drug is administered. The need to establish which modifications of the original models are legitimate has enriched the initial definition of actual causation with a series of the further specification [33, 34, 40], which is not the case to treat here.

In causal reasoning with structural equation models, the reliability of the answers obtained for our causal queries much depends on how we build the model, and the construction of a model is highly problematic. However, some general advantages can be listed:


- the possibility of building a model representing a causal structure and to assess the plausibility of each component of a series of alternative causal routes would reduce the extreme *subjectivity* of judgments that still affects clinical diagnostic judgment;
- causal models guarantee *more specific* causal inference, e.g., in clinical trials, the building of a model tailored to an individual patient would force an explicit representation of the comorbidity and polypharmacy of that patient; and
- given that in SEMs, all the hypothetical possible causal chains to the effects are made explicit, the inferential procedure exhibits a remarkable *explicative capacity*.

Using structural models does not in itself guarantee infallibility. In the SEM approach, the answer to causal queries depends on what is represented in the model, and this seems to require a large amount of general knowledge in building the model and choosing the right variables. In the following section, we offer some suggestions to alleviate the problem.

7 Normality and Similarity

A further important problem that the theory of structural models must face is the problem of so-called “structural isomorphism.” The problem arises when two causal ‘scenarios’ are represented by means of causal models having identical structures, but the causal responses licensed by the isomorph models are acceptable with respect to one scenario and counterintuitive with respect to the other. This suggested that the theory was incomplete and that it should have been enriched with a distinction between normal or *default* values and abnormal values of a variable [32, 41, 42]: causes involve changes or deviations from normality; therefore, a verdict of causation is legitimate only when variables representing alleged effects take an abnormal value. Normality should relate to different systems of norms, like

Fig. 1 DAG associated to the example



statistical, moral, social norms, and “norms of proper functioning that apply to artifacts and biological organisms” [43, p. 598].

If we turn to the kinds of problems treated in this paper, we realize that the definition of normality as ‘proper functioning’ is the most interesting one, provided that it is not taken too literally. In the first steps of a clinical diagnosis, a ‘normal’ person is a person without the signs and symptoms recognized by his doctor; a normal oncologic patient in a clinical trial is a person who has cancer and has never been administered the drug under study; in a forensic context, normality is even more context-dependent: in some cases, it may consist in the survival of the alleged victim, in others in his death from natural causes. In general, our intuition is that for each individual under some form of causal inquiry, the standard of normality should be given by her ‘*normal counterpart*.’

How could the normal counterpart of an individual be obtained, or at least be approximated? The idea of “normal counterpart,” i.e., a counterfactual representation or a virtual model of the individual case under exam, is not new. On the one hand, it borrows from the analysis of the use of exemplars and prototypes in biomedical sciences [44, 45]. On the other hand, it draws inspiration from computer science, which has developed normative theories of similarity analysis, case-based reasoning (CBR),⁷ and, more recently, digital twins.⁸

⁷ CBR is a methodology of reasoning widely used to solve problems in many medical areas. Briefly, CBR fulfils two main tasks: (1) the *retrieval* of cases similar to the one under study, and (2) the *adaptation* of former solutions to similar cases [46]. CBR requires the construction of a database containing cases with which the new cases will be compared. Each new case will therefore be compared to an *index case*, which is a case similar to the new one according to relevant features; the index case has a paradigmatic role [45] and can be seen as a model of the patient, representing its ‘normal’ counterpart.

⁸ “Digital Twins stand for a specific engineering paradigm, where individual physical artifacts are paired with digital models that dynamically reflects the status of those artifacts. When applied to persons, Digital Twins are an emerging technology that builds on in silico representations of an individual that dynamically reflect molecular status, physiological status and life-style over time. We use Digital Twins as the hypothesis that one would be in the possession of very detailed bio-physical and lifestyle information of a person over time. This perspective redefines the concept of ‘normality’ or ‘health,’ as a set of patterns that are regular for a particular individual, against the backdrop of patterns observed in the population” [47].

The building of a virtual representation of the individual should support singular causal reasoning by allowing the identification of what counts as a ‘deviation from the norm’ for a particular individual. The comparison with the virtual counterpart would perform two functions, the first in providing a ‘customized’ meaning of normality, and the second in choosing the right variables for the structural model representing the ‘history’ of the patient. For example, in the case of an early-phase clinical trial, given that a patient suffered an adverse event e , the structural equation model should represent all possible causal routes to e in that patient, while the virtual counterpart of the patient should provide information about what is normal for her, and what her condition would probably have been if she had not received the treatment under study.

The robustness of this kind of causal reasoning analysis coincides with the “degree” or measure of similarity: the more the model is similar, the better the causal judgment will be. From a philosophical point of view, a well-defined notion of similarity to a model can provide the groundwork/basis to grasp the concept of causality as ‘deviation from normality.’

Several measures of similarity have been proposed, giving rise to an entirely new line of research, which cannot discuss here in-depth. However, it seems appropriate to quote here Brown [48, pp. 1–2], who calls *similaromics* “the generation or identification of patients similar to an index patient” and remarks:

Although patient similarity is in its early stages, ultimately information about diseases, risk factors, lifestyle habits, medication use, comorbidities, molecular and histopathological information, hospitalizations, or death are compared with laboratory investigations, imaging, and other clinical data assessing medical evidence of human behavior [...] patient similarity represents a paradigm shift that introduces disruptive innovation to optimize personalization of patient care.

8 Conclusion

The emphatic tone of the quotation above should not induce over-optimism. Inferences based on virtual counterparts, as every inference based on analogies and similarity, is riddled with difficulties; so is the use of massive databases that should be employed for the construction of the models. It is highly improbable that the method I have proposed here will not inherit most of the problems linked to both analogical reasoning and big data. However, it presents some advantages:

- it points to a link between the most advanced theoretical research on individual causation in some areas of medical reasoning and current research in the field of bioinformatics causality;
- it offers a rigorous theoretical ground for causal reasoning, a kind of reasoning that cannot be eliminated by health sciences;
- it suggests solutions that make use of big data but are not exclusively data-driven, given the role of structural equation models; and

- it avoids problems of explainability, given that each step in the construction of a structural equation model is visible and justifiable.

If I had to hazard a prediction on the state of the research in the next thirty years, I think that many problems on biases and non-explainability of algorithms will be partly solved, and this also holds true in the field of causal reasoning. I also foresee that the research on methods of causal inference, far from disappearing, will reach high levels of precision and sophistication. Vice versa, I have strong doubts about the possibility of solving or dissolving in the next thirty years the fundamental problems of causality. But the task of philosophy is not solving problems but letting problems emerge.

Core Messages

- Diagnosis, forensic reasoning, and early clinical trials require reasoning on singular causes.
- Reasoning on singular cases in medicine is analogous to abductive reasoning used in detective stories: find the culprit!
- Structural equation models can be one of the best tools for medical reasoning on singular causes.
- An alternative and promising way to make causal inferences about singular cases can be based on digital twins.

Acknowledgements This research has a financial support of the Università del Piemonte Orientale [FAR 17].

References

1. Gorovitz S, MacIntyre A (1975) Toward a theory of medical fallibility. The Hastings Center 17 Report, V:13–23 (then in *J Med Philos* I, 1, 1976:51–71)
2. The Precision Medicine Initiative. <https://obamawhitehousearchives.gov/precision-medicine>
3. Boniolo G, Campaner R (2019) Causal reasoning and clinical practice: challenges from molecular biology. *Topoi* 38(2):423–435
4. Collins FS, Varmus H (2015) A new initiative on precision medicine. *N Engl J Med* 372(9):793–795
5. Eells E (1991) Probabilistic causality. Cambridge University Press, Cambridge
6. Lewis D (1973) Causation. *J Philos* 70(17):556–567
7. Menzies P, Beebe H (2020) Counterfactual theories of causation. *Stanford Encyclopedia of Philosophy*, Summer 2020 edition, EN Zalta (ed)
8. Gillies D (2019) Causality, probability, and medicine. Routledge, London
9. Illari P, Russo F (2014) Causality: philosophical theory meets scientific practice. OUP, Oxford
10. Dawid AP, Fagman DL, Fienberg SE (2014) Fitting science into legal contexts: assessing effects of causes or causes of effects? *Sociol Methods Res* 43(3):359–390

11. Benzi M (2011) Medical diagnosis and actual causation. *L&PS-Logic Philos Sci* 9:365–372
12. Cournoyea M, Kennedy AG (2014) Causal explanatory pluralism and medically unexplained physical symptoms. *J Eval Clin Pract* 20(6):928–933
13. Maung HH (2017) The causal explanatory functions of medical diagnoses. *Theor Med Bioeth* 38(1):41–59
14. Sadegh-Zadeh K (2012) *Handbook of analytic philosophy of medicine*. Springer, Dordrecht
15. Hempel CG (1965) *Aspects of scientific explanation and other essays in the philosophy of science*. Free Press, New York
16. Hunter K (1996) “Don’t think zebras” Uncertainty, interpretation, and the place of paradox in clinical education. *Theor Med Bioeth* 17(3):225–241
17. Cartwright N (2005) How can we know what made the Ratman sick? Singular causes and population probabilities. In: Jokić A (ed) *Philosophy of religion, physics, and psychology: essays in honour of Adolf Grünbaum*. Prometheus Books, Amherst NY, pp 377–390
18. Dragulinescu S (2012) On anti-Humeanism and medical singular causation. *Acta Anal* 27(3):265–292
19. Cartwright N (2011) A philosopher’s view of the long road from RCTs to effectiveness. *The Lancet* 377(9775):1400–1401
20. Osimani B (2013) Until RCT proven? On the asymmetry of evidence requirements for risk assessment. *J Eval Clin Pract* 19(3):454–462
21. Agbabiaka TB, Savović J, Ernst E (2008) Methods for causality assessment of adverse drug reactions: a systematic review. *Drug Saf* 31(1):21–37
22. Aronson JK, Hauben M (2006) Anecdotes that provide definitive evidence. *BMJ* 333(7581):1267–1269. <https://doi.org/10.1136/bmj.39036.666389.94>
23. Russo F, Williamson J (2011) Generic versus single-case causality: the case of autopsy. *Eur J Philos Sci* 1(1):47–69
24. Messina R, Consonno S, Ferrari L, Valente B (2003) Accesso asmatico e morte improvvisa in soggetto quattordicenne: ammissibile concorso di un farmaco β_2 - stimolante. *Minerva med* 123(2):77–84
25. Rizzi DA, Pederssen SA (1992) Causality in medicine: towards a theory and terminology. *Theor Med* 13(3):233–254
26. Thagard P (1998) Explaining disease: correlations, causes, and mechanisms. *Minds Mach* 8(1):61–78
27. Ankeny RA (2014) The overlooked role of cases in causal attribution in medicine. *Philos Sci* 81(5):999–1011
28. Lipton P (2004) *Inference to the best explanation*, 2nd edn. Routledge, London
29. Nathan MJ (2016) Counterfactual reasoning in molecular medicine. In: Boniolo G, Nathan MJ (eds) *Philosophy of molecular medicine: foundational issues in research and practice*. Routledge, New York, pp 192–214
30. Bird A (2010) Eliminative abduction: examples from medicine. *Stud Hist Philos Sci A* 41(4):345–352
31. Spirtes P, Glymour CN, Scheines R, Heckerman D (2009) *Causation, prediction and search*. MIT Press, Cambridge
32. Halpern JY, Hitchcock C (2011) Actual causation and the art of modeling. In: Dechter R, Geffner H, Halpern J (eds) *Heuristics, probability, and causality: a tribute to Judea Pearl*. College Publications, London, pp 383–406
33. Halpern JY, Pearl J (2005) Causes and explanations: a structural-model approach. Part I: causes. *BJPS* 56(4):843–887
34. Hitchcock C (2009) Causal modelling. In: Beebe H, Hitchcock C, Menzies P (eds) *The Oxford handbook of causation*. Oxford University Press, New York, pp 299–314
35. Pearl J (2009) *Causality: models, reasoning, and inference*. Cambridge University Press, New York

36. Cooper G (1999) An overview of the representation and discovery of causal relationships using Bayesian networks. In: Glymour CN, Cooper G (eds) *Computation, causation, and discovery*. AAAI/MIT Press, Menlo Park, pp 4–62
37. Long W (1989) Medical diagnosis using a probabilistic causal network. *Appl Artif Intell* 3(2–3):367–383
38. Richens JG, Lee CM, Johri S (2020) Improving the accuracy of medical diagnosis with causal machine learning. *Nat Commun* 11(1):1–9
39. Shpitser I, Pearl J (2012) Effects of treatment on the treated: identification and generalization. arXiv preprint [arXiv:1205.2615](https://arxiv.org/abs/1205.2615)
40. Hall N (2007) Structural equations and causation. *Philos Stud* 132(1):109–136
41. Halpern JY (2016) Appropriate causal models and the stability of causation. *Rev Symb Log* 9(01):76–102. <https://doi.org/10.1017/S1755020315000246>
42. Blanchard T, Schaffer J (2017) Cause without default. In: Beebe H, Hitchcock C, Price H (eds) *Making a difference*. Oxford University Press, Oxford, pp 175–214
43. Hitchcock C, Knobe J (2009) Cause and norm. *J Philos* 106(11):587–612
44. Schaffner KF (1986) Exemplar reasoning about biological models and diseases: a relation between the philosophy of medicine and philosophy of science. *J Med Philos* 11(1):63–80
45. Ankeny RA (2007) Wormy logic: model organisms as case-based reasoning. In: Creager ANH, Lunbeck E, Wise MN (eds) *Science without laws: models, systems, cases, exemplary narratives*. Duke University Press, Durham NC, pp 46–58
46. Schmidt R, Montani S, Bellazzi R et al (2001) Case-based reasoning for medical knowledge-based systems. *Int J Med Inform* 64(2):355–367
47. Bruynseels K, Santoni de Sio F, van den Hoven J (2018) Digital twins in health care: ethical implications of an emerging engineering paradigm. *Front Genet* 9:31
48. Brown SA (2016) Patient similarity: emerging concepts in systems and precision medicine. *Front Physiol* 7:561



Margherita Benzi Ph.D. in Philosophy of Science at the University of Genoa, teaches Logic and Philosophy of Science at the University of Eastern Piedmont (UPO). She published some books on causality in Italian (*Uncertain Reasoning*, 1997; *Discovering causes*, 2003; *Singular Causes*; 2020) and a number of papers in *Historia Mathematica*, *International Statistical Review*, *Logic and Philosophy of Science*, *Topoi*, *Disputatio*, and other journals. She is head of the UPO section of *PhilHeaD*, an interuniversity center for the philosophy of medicine, and a member of the roundtable of philosophy of medicine. She is on the board of teachers of the Northwestern Italian Doctorate in Philosophy (also enclosing Universities of Genoa, Turin, and Pavia), for which she teaches critical thinking and inductive fallacies. Her main interests are causal reasoning, philosophy of medicine, and argumentation theory.