

Warranted Diagnosis

David Limbaugh^{a,b,c}, David Kasmier^b, Werner Ceusters^{b,d}, Barry Smith^{a,b,d}

^a Department of Philosophy, University at Buffalo, Buffalo NY, USA

^b National Center for Ontological Research, University at Buffalo, Buffalo NY, USA

^c CUBRC, Buffalo NY, USA

^d Departments of Biomedical Informatics and Psychiatry, University at Buffalo, Buffalo NY, USA

Abstract

A diagnostic process is an investigative process that takes a clinical picture as input and outputs a diagnosis. We propose a method for distinguishing diagnoses that are warranted from those that are not, based on the cognitive processes of which they are the outputs. Processes designed and vetted to reliably produce correct diagnoses will output what we shall call ‘warranted diagnoses’. The latter are diagnoses that should be trusted even if they later turn out to have been wrong. Our work is based on the recently developed Cognitive Process Ontology and further develops the Ontology of General Medical Science. It also has applications in fields such as intelligence, forensics, and predictive maintenance, all of which rely on vetted processes designed to secure the reliability of their outputs.

Keywords:

Biomedical Ontology, Diagnosis, Warrant

Introduction

Hogan and Ceusters (1) define a diagnostic process as follows:

[DP1] Diagnostic Process =def. An interpretive process that has as inputs 1) a clinical picture of a given patient, 2) an aggregate of representations of 2a) at least one type of disease and 2b) at least one type of phenotype whose instances are associated with instances of that disease, and has as output 3) an assertion to the effect that the patient has a disease of a certain type.

[DP1] was proposed as a means of excluding from the realm of diagnoses those cases where a clinician (or fortune-teller) arrives at what we might outwardly think of as a diagnosis though what is in fact a lucky guess or a matter of mere hearsay – f.i. merely reading about a diagnosis and agreeing with it.

[DP1] rules out such cases because lucky guesses, hearsay, and so forth are products of processes that rely on sources of information unrelated to the formulation of a diagnosis, for instance horoscopes or rumors. More specifically, such cases are *not* produced by processes that input: 1) a *representation of phenotypes* that are *clinically abnormal* (called a ‘clinical picture’), 2a) an aggregate of representations of at least one type of *disease* and 2b) at least one type of phenotype whose instances are associated with instances of that disease.

When taken in conjunction with the further definitions provided in Table 3., [DP1] allows us to rule out lucky guesses and hearsay on the part of non-clinicians from counting as diagnoses. But they are only the first step in addressing luck in medical diagnoses. A goal of medicine is to develop processes that *reliably* produce diagnoses *that can be trusted*. While a trusted – or *trustworthy* – diagnosis may in fact be incorrect, it is still reasonable and appropriate to assume at the time when it is produced that it is correct. Two types of diagnoses therefore need to be distinguished, namely: those that were, and those that were not the output of a process that was successfully designed or vetted to reliably produce correct diagnoses. [DP1] does not yet draw this distinction.

The aim of this paper is threefold. It is, first, to revise [DP1] by expanding what counts as an output, and thereby further refining the definition of diagnosis. Second, to incorporate the Cognitive Process Ontology (CPO) (2), an extension of the Mental Functioning Ontology (MF) (3), into our treatment of diagnostic processes. This will result in a new extension ontology of the CPO, namely the Medical Cognitive Process Ontology (MCPO) and will allow us to introduce a new term for those diagnoses that should be trusted, which we shall henceforth refer to as ‘warranted diagnoses’. ‘Warrant’ is used here in the Plantingan sense (4) to mean roughly *trustworthy because of how it was produced*. This in turn will allow data about diagnoses to be tracked along a new and important dimension.

Methods

To show that [DP1] is insufficient to distinguish warranted from unwarranted diagnoses, we apply it to two clinical scenarios and observe the results in regards to warrant and luck. We then examine whether, by exploiting the principles of referent tracking, we can create a definition that will allow us to account for warrant and luck.

Materials

We input key terms from the Ontology of General Medical Sciences (OGMS) (5) and the term ‘representation’ from Smith and Ceusters (6). We also apply the principles of referent tracking as spelled out in (7–10).

Referent tracking (RT) is a strategy for organizing data that uses unambiguous names (for instance alphanumeric strings), called ‘instance unique identifiers’, to refer to entities, i.e. particulars, in the world. In addition, RT provides a system for creating metadata to track different portions of reality (PORs) and the relations between them as the world changes. Importantly, we understand the world as including also the referent tracking system itself and the data that the system organizes. These data, too, can be assigned instance unique identifiers when needed.

Of interest to us here is the way in which a referent tracking system (RTS) categorizes errors. (An RTS is an implementation of referent tracking.) Errors, too, are PORs, and so errors can be tracked. The question raised in this communication is whether tracking errors using an RTS can help us to categorize errors in a way that takes warrant into account. The types of errors we pay attention to are represented by specific error codes (listed in Tables 1 and 2 below), and we will use these to document our analysis.

For the sake of continuity we use two scenarios from Hogan and Ceusters (1). Both scenarios involve the same patient, Mr. Jones. In each case, Mr. Jones has the disease *type 2 diabetes mellitus*. The cases differ in regard to who is the treating physician and whether that physician gave a correct diagnosis. The scenarios, and the details added thereto, are as follows.

Scenario 1: Correct diagnosis by physician

Dr. Anne Smith sees Mr. Jones in her office. She takes a history and physical, performs certain laboratory tests, and on the basis of her analysis of the findings, she correctly concludes that Mr. Jones has type 2 diabetes mellitus. She subsequently writes her diagnosis in the patient’s medical record.

Scenario 2: Incorrect diagnosis by physician at a later time

Mr. Jones is traveling on vacation, when he falls ill. He sees Dr. Jane Miller who does not have any of his past records available, and thus is not aware of the previous diagnoses made by Dr. Smith. Dr. Miller infers a new clinical picture of Mr. Jones, and incorrectly concludes on its basis that Mr. Jones has *type 1 diabetes mellitus*. She accordingly records a diagnosis of type 1 diabetes mellitus in her medical record for Mr. Jones.

Details Added to the Scenarios

Scenario 1: The laboratory testing performed by Dr. Smith is unknowingly completed using unreliable equipment.

Scenario 2: The new clinical picture formed, and the subsequent diagnosis asserted, by Dr. Miller was the output of a type of diagnostic process that, according to peer review, is highly reliable at correctly diagnosing type 1 diabetes mellitus. The diagnostic process was carried out properly.

Analysis

Scenario 1: A lucky diagnosis

A diagnostic process that depends on unreliable equipment cannot be trusted to produce a correct diagnosis. This is because the fidelity of a clinical picture of the sort used by Dr. Smith is subject to the reliability of the equipment used in testing. Unreliable equipment results in an unreliable clinical picture, and an unreliable clinical picture results in an unreliable diagnosis. (‘Reliable’ here means: has a high likelihood of being correct.)

We have stipulated that Dr. Smith’s diagnosis was in fact correct; however its correctness is, because of the use of bad equipment, not the product of a procedure which is of a sort that has been vetted for use in the given environment. Furthermore, the fact that Dr. Smith was not aware that the equipment was unreliable changes nothing regarding the warrant of the diagnosis. The trustworthiness of a diagnosis is a function of the reliability of the process in the circumstance used to produce the diagnosis. Dr. Smith’s diagnosis fails to be warranted on the account we are here proposing.

Scenario 2: An unlucky diagnosis

Dr. Miller’s diagnosis is wrong. However, in order to diagnose Mr. Jones, she properly used, in an environment for which it was vetted, a peer reviewed and highly reliable diagnostic process. Thus, although Dr. Miller’s diagnosis is wrong, until the time she comes to believe that it *is* wrong, her diagnosis is warranted and should be trusted. This is because we should trust the outputs of highly reliable processes unless that output is reasonably called into question.

Towards a Revision of the Definition of ‘Diagnostic Process’

[DP1] discriminates between diagnoses and non-diagnoses only on the grounds of whether or not the process inputs a clinical picture and an aggregate of representations (of at least one type of disease and at least one type of phenotype whose instances are associated with instances of that disease) and outputs an assertion to the effect that the patient has a disease of a certain type. Thus, the assertions in both scenarios qualify as diagnoses. The only difference between them, according to [DP1], is that Dr. Smith’s diagnosis is correct and Dr. Miller’s diagnosis is incorrect.

This means that there is no discrimination along the dimensions of reliability or warrant. Incidentally, among the correct diagnoses Dr. Miller could have made is that Mr. Jones *does not* have type 1 diabetes. Though this appears to be a legitimate diagnosis, [DP1] does not allow it to be classified as such. It also does not allow disease course, disorder, or some combination of these (including combinations also involving disease) to be the subject of a diagnosis.

Referent Tracking

RT, too, did not until now have the resources to account for the distinction between warrant and luck. RT asks the following six questions in order to establish whether the author made a mistake when adding an assertion (such as ‘Mr. Jones has diabetes’) to an RTS:

1. Does the POR represented by the assertion objectively exist?
2. Is the represented POR objectively relevant?
3. Does the author believe that the represented POR exists?
4. Does the author believe that the POR represented by the assertion is relevant?
5. Is the assertion in the RTS the assertion intended by the author to represent the POR?
6. In what way does the assertion in the RTS refer?

Note: an answer of ‘No’ to question’s 4 or 5 means the author believes the POR does *not* exist or is *not* relevant.)

Answer configurations P+1 and A+1 through A+4 (Table 1) indicate no error, while all other configurations (Table 2) indicate some error in the RTS.

Table 1 – Referent Tracking Codes that Indicate No-Error

'P' = 'present', 'A' = 'absent', '+' = no error, 'N/A' = 'not applicable', and 'N/C' = 'not considered'							
Questions							
Configurations		Q1	Q2	Q3	Q4	Q5	Q6
	P+1	Yes	Yes	Yes	Yes	Yes	Refers Correctly
	A+1	No	N/A	No	N/A	N/A	N/A
	A+2	Yes	No	Yes	No	N/A	N/A
	A+3	Yes	No	N/C	N/A	N/A	N/A
	A+4	No	N/A	N/C	N/A	N/A	N/A

Table 2 – Referent Tracking Codes that Indicate Error

'P' = 'present', 'A' = 'absent', '+' = no error, 'N/A' = 'not applicable', 'N/C' = 'not considered', 'Conf.' = 'configuration'							
Questions							
Conf.	Q1	Q2	Q3	Q4	Q5	Q6	
P-1	No	N/A	Yes	Yes	Yes	No	No Referent
P-2	No	N/A	No	N/A	N/A	No	No Referent
P-3	No	N/A	Yes	No	N/A	No	Refers Inaccurately
P-4	Yes	No	N/C	N/A	N/A	No	No Referent
P-5	Yes	N/A	N/C	N/A	N/A	No	No Referent
P-6	Yes	No	Yes	Yes	Yes	Yes	Refers Correctly
P-7	Yes	No	Yes	Yes	No	No	No Referent
P-8	Yes	No	Yes	Yes	No	No	Refers Inaccurately
P-9	Yes	Yes	Yes	Yes	Yes	Yes	Redundant Reference
P-10	Yes	No	Yes	Yes	Yes	Yes	Redundant Reference
P-11	Yes	Yes	Yes	Yes	Yes	Yes	Ambiguous Reference
P-12	No	N/A	Yes	Yes	Yes	Yes	Ambiguous Reference
A-1	Yes	Yes	Yes	Yes	N/A	N/A	N/A
A-2	Yes	Yes	No	N/A	N/A	N/A	N/A
A-3	No	N/A	Yes	No	N/A	N/A	N/A
A-4	Yes	No	No	N/A	N/A	N/A	N/A
A-5	Yes	Yes	N/C	N/A	N/A	N/A	N/A

Table 3 – Key Terms from OGMS used in the Analysis

Term with Definition / Elucidation
Clinical Phenotype: A clinically abnormal phenotype (1).
Clinical Picture: A representation of a clinical phenotype that is inferred from a combination of, for example, diagnoses and laboratory, image, and clinical findings about a given patient (1).
Disease: A disposition (i) to undergo pathological processes that (ii) exists in an organism because of one or more disorders in that organism (5).
Diagnostic Process [DP1]: An interpretive process that has as inputs 1) a clinical picture of a given patient 2) an aggregate of representations of 2a) at least one type of disease and 2b) at least one type of phenotype whose instances are associated with instances of that disease, and has as output 3) an assertion to the effect that the patient has a disease of a certain type (1).
Diagnosis: A conclusion of an interpretive process that has as input a clinical picture of a given patient and as output an assertion to the effect that the patient has a disease of such and such a type (1).
Disease Course: The totality of all processes through which a given disease instance is realized (5).
Disorder: A causally relatively isolated combination of physical components that is (a) clinically abnormal and (b) maximal, in the sense that it is not a part of some larger such combination (5).
Phenotype: A (combination of) bodily feature(s) of an organism determined by the interaction of its genetic make-up and environment (5).
Representation: A quality which is about or is intended to be about a portion of reality (6).
<i>x</i> is a Clinically Abnormal Phenotype: <i>x</i> is a non-canonical phenotype of an organism and <i>x</i> increases the organism's risk of being harmed (5).
<i>x</i> is a Portion of Reality: <i>x</i> exists or is a configuration of existents (6).

According to this strategy, Dr. Smith's diagnostic assertion is *without* error and should receive a code of P+1 – indicating an answer of 'Yes' to questions 1 through 5 and an answer of 'successfully refers' to answer 6. Dr. Miller's diagnostic assertion would be considered *with* error and should receive a code of P-1, which indicates 'Yes' to questions 3-5 and 'No' to question 1, 'not applicable' to question 2, and 'no referent' to question 6. As for [DP1] so also for the error configurations in RT, the only discrimination allowed between Dr. Smith's and Dr. Miller's assertions is that Dr. Smith's is correct and Dr. Miller's incorrect. No further dimension of assessment is available.

Table 4 – ‘BE’ means ‘Believes Exists’, ‘BR’ means ‘Believes is Relevant’

Believes Exists	Believes Relevant	Warranted BE	Warranted BR	Warrant Code
Yes	Yes	Yes	Yes	w2
		Yes	No	w1-2
		No	No	w0
No	N/A	Yes	N/A	w1
		No	N/A	w0
Yes	No	Yes	Yes	w0
		Yes	No	w1-2
		No	No	w0
N/C	N/A	N/A	N/A	wna

Results

The result of our analysis is that a notion of ‘warrant’ is required, both in OGMS and in RT, to discriminate between diagnoses that are warranted and those that are merely lucky.

To this end we import into OGMS the term ‘warranted assertion’ from CPO, along with the required dependencies. This allows us to define ‘warranted diagnosis’ and ‘proper diagnostic functioning’ (See Table 5). This introduces a normative aspect to the treatment of both ‘diagnosis’ and ‘diagnostic process’. These new terms now belong to both CPO and the Medical Cognitive Process Ontology (MCPO).

The definition of ‘diagnostic process’ also needs to be revised to allow both for negative diagnoses and for those cases where the output of a diagnostic process involves a representation of a disorder, disease course, or some combination of these (including combinations involving disease).

To add warrant to RT, we import the CPO term ‘representation that is believed’ (RTB) and its subclass ‘representation that is warranted’ (RTW). This allows us to add an additional question to RT’s error-checking questionnaire: “Is what the author believes about the existence and relevance of portions of reality warranted?”

The answer to this additional question depends on questions 3 and 4, both which ask what the author believes. Warrant is only applicable when there is a representation that is believed (RTB) to apply warrant to. And when warrant is applied, this makes the representation a representation that is warranted (RTW). Note that just because warrant is applied this does not mean that the RTB is actually warranted only that it has been assessed as such. Furthermore, it is probable that not all truly warranted RTBs will be tagged as warranted. There is room for error in the application of warrant and in the determination of whether warrant has been applied at all.

Additionally, in all cases, if the RTB of “Believes POR exists?” (question 3) is unwarranted – that is, if it is not an RTW – then the RTB of “Believes POR relevant?” (question 4) is also unwarranted. This is because a positive answer to “Believes POR exists?” is required for an RTB about that POR’s relevance to be warranted. I cannot be warranted in believing that the Jabberwocky is *relevant* without also believing that the Jabberwocky exists. Finally, if answers to question 3 or 4 are

not considered or *not applicable*, then warrant is not applicable to those answers either.

The representation of warrant in an RTS takes the following forms (see also Table 4):

1. ‘w1’ the author’s RTB about a POR’s existence is warranted and warrant is not applicable to the author’s RTB about the POR’s relevance.
2. ‘w1-2’ the author’s RTB about a POR’s existence is warranted but what the author believes about the POR’s relevance is unwarranted.
3. ‘w2’ the author’s RTBs about a POR’s existence and relevance are both warranted.
4. ‘w0’ the author’s RTB about a POR’s existence is unwarranted and warrant is not applicable to the author’s RTB about the POR’s relevance.
5. ‘wna’ warrant is not applicable to the author’s RTBs about the existence or relevance of a POR.

The codes are intended to be appended to the current RT error codes; for instance ‘P+1w2’ would indicate: an answer of ‘yes’ to every question, that all assertions refer correctly, and that both RTBs are warranted.

Discussion

Revising the definition of ‘Diagnostic Process’

[DP1] needs to be revised, first, so that it will refer not simply to diagnostic assertions about a patient’s having a certain disease, but rather to the patient’s *either* having a certain disease, *or* participating in a certain *disease course*, *or* having a certain *disorder*, *or* having none, some, or all of these.

These additions are important because it may be, for example, that a patient is first diagnosed as participating in a disease course, for instance manifesting *cortisol deficiency*. The latter might then be a realization of a number of different diseases and is only later diagnosed as a case of the specific disease of *congenital adrenal insufficiency*. The assertion that the patient has cortisol deficiency is no less a diagnosis than is the assertion that the patient has congenital adrenal insufficiency, even though only the latter refers to a disease. Each of these assertions can be arrived at through a combination of a clinical picture and an aggregate of relevant representations; each is a proper subject of medical concern; and each calls for a treatment plan. The same can be said generally of disease, disease course, disorder, and any and all combinations of these. Each is what, for convenience in this paper, we shall call a ‘condition’ (compare (11,12)). Each is properly asserted during or as an output of a diagnostic process, as something had by a patient. (Note that ‘had’, here, is shorthand for either ‘has a part that participates in’ (in the case of a disease course), ‘has a part that has part’ (in case of a disorder), or ‘has a part that bears’ (in other cases).) There are successful diagnoses and there are failed diagnoses. The family of successful diagnoses includes assertions of the form ‘has condition’; however, it also includes assertion of the form ‘has no condition’ (or ‘is healthy’) (13). A diagnostic process is an investigative process – an inquiry – into the health of a patient and as long as that inquiry concludes with an assertion of one or other of the forms ‘has condition’ or ‘has no condition,’ then the goals of the inquiry are met.

Table 5 – Terms from CPO with draft definitions, together with MCPO to be integrated with OGMS

Term	Definition
Assertion	An information quality entity that is the concretization of a descriptive information content entity that is expressible by means of a sentence (CPO).
Clinical Picture that is Warranted	A clinical picture that is a representation that is warranted (MCPO).
Cognitive Process	A mental process that creates, modifies or has as participant some cognitive representation (CPO).
Cognitive System	A system all of whose parts are also parts of a single organism and which realizes mental dispositions (CPO).
Cognitive Representation	A mental representation that has a mind-to-world direction of fit (CPO).
Confidence Value	A mental quality that, when fused with a cognitive representation CR, determines the extent to which a cognitive system operates as if CR is veridical (CPO).
Descriptive Information Content Entity	An information content entity that describes some portion of reality.
Diagnostic Process (Revised) [DP2]	An investigative process that has as inputs: 1) a clinical picture of a given patient, 2) an aggregate of representations of 2a) at least one type of disease, disease course, or disorder and 2b) at least one type of phenotype whose instances are associated with instances of that disease, disease course, disorder, or combination thereof, and has as output(3) an assertion based on 1) and 2) to the effect that the patient does or does not have a disease, disease course, disorder, or combination thereof of a certain type (MCPO).
Information Quality Entity	A quality that is the concretization of some information content entity (6).
Investigative Process	A cognitive process whose agent intends to establish or confirm that some portion of reality exists or does not exist (CPO).
Mental Quality	A quality which specifically depends on an anatomical structure in the cognitive system of an organism and is experiential (compare with (6)) (CPO).
Mental Representation	A representation which is a mental quality (6).
Process of Proper Cognitive Functioning	A process of cognitive functioning that has been successfully vetted or designed to reliably form veridical cognitive representations in environments of given types (CPO).
Processes of Proper Diagnostic Functioning	A diagnostic process that inputs representations that are warranted, including a clinical picture that is warranted, and, based on these inputs, outputs a warranted assertion to the effect that the patient does or does not have a disease, disease progression, disorder, or combination thereof of a certain type (MCPO).
Representation that is Believed	A cognitive representation that is fused with a positive confidence value (CPO).
Representation that is Warranted	A representation that is believed formed through proper cognitive functioning in a vetted- or designed-for environment (CPO).
System	A material entity including as parts multiple objects that are causally integrated (16).
Warranted Assertion	An assertion that is based on a representation that is warranted (CPO).
Warranted Diagnosis	A warranted assertion to the effect that the patient does or does not have a disease, disease progression, disorder, or combination thereof of a certain type and that is the output of proper diagnostic functioning (MCPO).

This allows for ordinary diagnostic practices like asserting that a person is cancer free after treatment. An example of a failed diagnostic process would be a concluding assertion such as: ‘inconclusive regarding the presence of a condition’. (We leave open here the question of how specific a diagnostic assertion must be to qualify as a diagnosis.)

To apply these improvements to the definition of ‘diagnostic process’ we replace the (undefined) term ‘interpretative process’ with the term ‘investigative process’ from CPO, which represents a subclass of what the Mental Functioning Ontology (MF) terms a ‘cognitive process’:

Cognitive Process =def. A mental process that creates, modifies or has as participant some cognitive representation (MF, CPO).

Investigative Process =def. A cognitive process whose agent intends to establish or confirm that some portion of reality exists or does not exist (CPO).

An investigative process can be as simple as glancing upwards to confirm the position of the hands of a clock and as complex as an international terrorist hunt.

Putting the above together, ‘diagnostic process’ can now be defined as follows:

Diagnostic Process [DP2] =def. an **investigative** process that has as inputs: 1) a clinical picture of a given patient, 2) an aggregate of representations of 2a) at least one type of disease, **disease course, or disorder** and 2b) at least one type of phenotype whose instances are associated with instances of that disease, **disease course, disorder, or combination thereof**, and has as output 3) an assertion based on 1) and 2) to the effect that the patient **does or does not have** a disease, **disease course, disorder, or combination thereof** of a certain type. (Revisions are in bold)

A diagnostic process is aimed at establishing or confirming the presence of a condition in a patient. Investigations unfold as an agent follows *indicators*, which are portions of reality that affect that agent’s estimation that some other portion of reality exists. Practically anything (real) can be a portion of reality. So not only are universals and instances portions of reality (and potential indicators), but so also are combinations of these, such as a patient in Tucson, Arizona having a stage four carcinoma in his lung at 12pm MST on October 12, 1972 (14). A clinical picture, because of what it represents, is the key type of indicator for clinical diagnostics.

Mental Representation

Our task here and onward is to examine the systems in which warrant plays a role, and specifically to explore the types of cognitive processes which provide its substrate.

‘Mental quality’ is a key term in the CPO and is a subclass of BFO:quality (15). ‘x is a Mental Quality’ means, provisionally, that 1) x is a quality which specifically depends on an anatomical structure in the cognitive system of an organism (compare with (6)) and 2) x is, for lack of a better word, experiential. The term ‘experiential’ is meant to distinguish mental qualities from other qualities that inhere in an anatomical structure in a cognitive system, like the shape of Broca’s area. We remain agnostic as to what a mental quality’s physical basis might be, that is, what sort of independent continuant it inheres in.

Mental qualities are either representational or they are not. Non-representational mental qualities include those that are responsible for giving emotional and sensational processes their characteristic feel. For example, the process of experiencing pain *hurts* because of the mental qualities involved in that process, and similarly for experiences of sorrow or joy.

In formulating the above elucidation, we are agnostic as to which parts of an organism constitute its cognitive system. We do however assume that it includes parts of the brain. The term ‘structure’ should also be understood in a very general sense,

including for instance areas of the brain with particularly dense neuronal connections specialized to specific sorts of mental functioning. Broca’s area is a structure in this broad sense.

System =def. A system is a material entity including as parts multiple objects that are causally integrated (16).

Cognitive System =def. A system all of whose parts are also parts of a single organism and which realizes mental dispositions (CPO).

The definitions of ‘system’ and ‘cognitive system’ presented here are also provisional, and should be read in conjunction with the proposed definition of ‘bodily system’ found in (17).

Aboutness

Some mental qualities are representations defined as follows:

Representation =def. A quality which is about, or is intended to be about, a portion of reality (6).

Mental Representation =def. A representation which is a mental quality (6).

Mental representations are responsible for the intentionality (directedness or aboutness) found in a cognitive process (henceforth just ‘about’ or ‘aboutness’). When asked, “What are you thinking about?” the answer is dependent on your mental representations. We address ‘is about’ first and then discuss ‘intended to be about’.

Is About

For the purpose of this discussion we distinguish two kinds of entities that have aboutness: mental representations and information content entities (ICEs).

ICEs are BFO:generically dependent continuants (GDCs), which means that an instance of an ICE can have multiple concretizations (15). For example, the particular instance of an ICE that is *Grey’s Anatomy* – also an instance of the subtype *textbook* – not only exists as concretized by the pattern of qualities inhering in the physical book (made of ink, glue, and paper) on your shelf, but also in the physical books on the shelf in the library and at the bookstore. It also exists as concretized by the electromagnetic patterns that form the pdf file in your laptop. *Grey’s Anatomy* is concretized in each case (they are all distinct copies of the same textbook). It is concretized by distinct instances of complex quality patterns inhering in different individual books or digital files. *Grey’s Anatomy* then depends generically on each and every book (or file) that concretizes it, and each and every book (or file) would have to be destroyed to successfully destroy *Grey’s Anatomy* itself.

Mental representations, in contrast to ICEs, are BFO: *specifically dependent continuants*. Thus, an instance of a mental representation specifically depends on part of a cognitive system and is only located where its bearer is located.

Two Types of Aboutness

We distinguish two types of aboutness: *original* and *derived* (18). This distinction mirrors that between bona fide and fiat boundaries; both types of boundaries exist and are genuine, but the former are associated with ‘natural demarcations’, such as walls and rivers, while the latter only come into existence through the intentional actions of agents such as the signing of a legal document that specifies a property line (19). Even fiat property lines are then parts of reality, and have legal significance.

Entities with Original Aboutness

Entities that have original aboutness are various types of mental representations, qualities of parts of a cognitive system that are about this or that, for example when I see an apple before me. The aboutness here is *original* because, like bona fide boundaries, the *is about* relation between a mental representation and its referent is not derived from the intentions of agents in any way analogous to the way in which fiat boundaries come into existence. We suspect there is no original aboutness outside of mental representations.

Entities with Derived Aboutness

Entities with derived aboutness are ICEs concretized in symbols (quality-patterns) such as ‘dog’ or ‘ π ’, either spoken, written, or otherwise instantiated (for example on a memory stick). Symbols have their ICEs imparted upon them by the intentionality of agents. Thus, it was only after an act of naming that ‘ π ’ became one way of expressing the ICE otherwise expressed as ‘pi’. Following Chisholm’s doctrine of the primacy of the mental, the derived aboutness of an ICE, is *explained* in terms of some original aboutness (20). The reason why the ink or pixel pattern ‘ π ’ is associated with the ICE “the ratio of a circle’s circumference to its diameter” is, first, because of the original aboutness *in the mind of* William Jones, who first introduced that symbol to carry this ICE, and, later, because of the original aboutness *in the minds of* nearly every student who learned the language of mathematics. Derived aboutness thus obtains for every ‘ π ’ appearing in books about geometry on library shelves. In the cases treated by Chisholm, original aboutness always precedes derived aboutness. We reserve for later a discussion of cases, like automated surveillance, where the temporal ordering is reversed.

Merely Intended To Be About

Representations can exist without an aboutness relation (6,21). For example, the symbol (pattern of ink or pixels) ‘ π ’ would still *exist* even without an aboutness relation in a world in which the pattern was never associated with any ICE. And similarly, mental representations can also exist without an aboutness relation. While there is no *miasma*, there are mental representations in the minds of some medical historians that are suited to be about *miasma*. These mental representations would be about *miasma* were such an entity to exist.

Though a mental representation can fail to have an aboutness, mental representations are always *intended to be about* something. This primitive notion ‘intended to be about’ describes a *suitedness* to be about something or other. A mental representation can be suited to be about x and at the same time fail to be about x because x does not exist. As an analogy, a key is suited to open a particular type of lock, but that does not imply that the key is ever used to open such a lock nor even that instances of locks, of the type that the key is suited to open, do in fact exist.

Cognitive Representation

‘Cognitive representation’ is a subtype of ‘mental representation’. The distinguishing feature of a cognitive representation is what Searle called its ‘mind-to-world direction of fit’ (22,23). Cognitive representations can be more or less accurate. If a cognitive representation is inaccurate, then the error is in the cognitive representation and not elsewhere; the

cognitive representation aims to fit what it is intended to be about *in the world* and not vice versa. ‘ x is a Cognitive Representation’ means x is a mental representation that has a mind-to-world direction of fit (CPO).

Contrast this with a type of mental representation that would be associated with a desire; a desire demands that the world fit it and not *vice versa*; it has a world-to-mind direction of fit.

Veridicality

A distinction should be drawn between a representation’s being correct and its degree of correctness. A representation is correct (henceforth; ‘veridical’) when it is about *the* portion of reality that it is intended to be about (6). This is not the same as only the representation’s constituent representational units being about what they are intended to be about (6). Take for example a cognitive representation CR_1 “Samuel Albert is in my living room.” CR_1 has some degree of accuracy because “Samuel Albert” successfully refers to *Samuel Albert* and “my living room” successfully refers to *my living room* – fulfilling the cognitive representation’s implication that each exists. However, CR_1 is not veridical because Samuel Albert and my living room aren’t in the right configuration because Samuel Albert is not actually in my living room. To be veridical is not a matter of accuracy simpliciter but rather only of accuracy to the degree of detail that is appropriate given the level of granularity of the representation in question (24). When I assay the color of my wallpaper, I do not concern myself with the molecules on its surface.

Representation that is Believed

Some cognitive representations are taken by the agent to be veridical. These are what we referred to above with the term ‘representation that is believed’ (RTB).

An RTB is treated by the agent (or by his cognitive system) as being actually true, even though it may not in fact be actually true. More specifically, what distinguishes an RTB from a mere cognitive representation is that the latter is fused with a positive confidence value (Compare what Meinong has to say about *Ernstgeföhle* or, more generally, about *serious* (or *earnest*) *mental phenomena* in (25).)

‘Fusion’ is a term adapted from Husserl (26) (who in turn takes it from Stumpf (27)) and is a primitive relationship that obtains between multiple quality instances when they are so closely related that an additional quality instance seems to emerge from them.

Take for example what appears to be a solid green image displayed on a television screen, which upon very close inspection is revealed to have a color made out of tiny yellow and blue squares (or pixels), which give a green appearance to the naked eye. The pixels are bearers of many instances of yellow and blue, and these instances appear to have fused into an additional instance of greenness. Similarly, when an instance of a cognitive representation and an instance of positive confidence value are fused together in a cognitive system there seems to be an additional quality instance: an instance of an RTB.

A *confidence value* is a non-representational mental quality that, when fused with a cognitive representation, determines how that cognitive representation is utilized by a cognitive system. The result is that the cognitive system operates with that

cognitive representation as if it is veridical. If cognitive representation CR₂ – that “My coffee is still too hot to drink” – is fused with a positive confidence value, then CR₂ might be taken as input by the agent’s cognitive system when making a decision as to when to take a sip of the coffee.

Importantly, a fused confidence value should not be confused with second-order cognitive representations, such as a representation about the likelihood of another representation’s being veridical (as for example when you are asked: “Are you sure?”). Such second-order mental representations are distinct from the pre-introspective and non-representational confidence that we find fused with those cognitive representations which are RTBs. Here a second-order mental representation is an evaluation of the confidence value fused with a mental representation. We leave the full explication of this distinction for future work.

Confidence Value =def. A mental quality that, when fused with a cognitive representation CR, determines the extent to which a cognitive system operates as if CR is veridical (CPO).

With this in mind we can now define ‘representation that is believed’ as follows:

Representation that is Believed (RTB) =def. A cognitive representation that is fused with a positive confidence value (CPO).

Representation that is Warranted

Following Plantinga (4), a representation that is warranted (RTW) is an RTB which holds an epistemically privileged place in a cognitive system because it was produced by some designed or vetted process so that, when in an environment of the sort that it was designed or vetted for, it reliably outputs veridical cognitive representations. As such, the produced RTB is not just *de facto* fused with a positive confidence but also is such that it *should* be fused with a positive confidence. Instances of such processes are instances of ‘proper cognitive functioning’, and the cognitive representations formed by such processes are *warranted*:

Process of Proper Cognitive Functioning =def. A cognitive process that has been successfully vetted or designed to reliably form veridical cognitive representations in environments of given types (CPO).

Representation that is Warranted (RTW) =def. A representation that is believed formed through proper cognitive functioning in a vetted- or designed-for environment (CPO).

The privilege of an RTW is not that it is in every case correct. (“Reliably” does not imply: *in every case*.) Rather it is that it can justifiably be used in a cognitive process without further scrutiny.

Expanding on Warrant

A paradigm example of an RTW is one formed on the basis of sense perception when in the appropriate environment (for instance in an otherwise quiet room with good acoustics for hearing, or a clean palate for tasting). Consider the following: if you perceive that a ball is red while in a well-lit room standing in close proximity to the ball, then a cognitive representation that “The ball is red” formed on the basis of this experience is

warranted and can be used without further scrutiny to form other cognitive representations, like “The ball is my favorite color.”

It is assumed that the neurocognitive structures underlying normal perceptual processes – such as processes of forming a representation that the ball is red based on seeing that the ball is red – were in a sense vetted in the course of evolution. The ability to form representations that are believed based on encounters with external reality was part of what kept our ancestors alive and able to reproduce.

Domain Specific Warrant

Evolution is not the only way in which processes come to be vetted, however. There are other classes of vetted cognitive processes associated with domains of inquiry involved, for example, in predictive maintenance, medicine, and intelligence analysis. For example, an RTB about how much oil is left in an engine, formed on the basis of the perceptual/kinesthetic experience of a properly used dipstick, should be assumed to be veridical; as should a cognitive representation formed on the basis of a document about the results of a blood panel provided by an experienced laboratory technician. Furthermore, a process of intelligence gathering that relies on signals from an array of active and passive sensors to locate a satellite can also produce, in the minds of suitably qualified analysts, representations that are warranted.

Warrant in Medicine

In medicine, a process of proper cognitive functioning is designed and then vetted through peer review and empirical testing. For example, Marzolf et al. (28) describes a procedure using the *Spot Vision Screener* to screen for ophthalmological pathology. This procedure had already been vetted for certain pediatric cases (29), but it is there further vetted for cases where the patient is a child with developmental disabilities.

How reliable a process must be at producing veridical cognitive representations so as to be considered a process of proper cognitive functioning is in part a matter for peer review. The notion of ‘can be reasonably used without scrutiny’ which is at the heart of warrant will differ from field to field and, at least to some extent, be related to the question of when a cognitive representation is actionable. Whether or not a cognitive representation is actionable is determined on the basis of the general goals of experts in the associated field. In medicine the relevant cognitive representations are in many case not first-order representations (such as “The patient has condition *x*”) but second-order representations (“There is a 0.6 likelihood that the assertion ‘the patient has condition *x*’ is *veridical*”). If a medical provider is unwarranted in regards to the first cognitive representation but warranted in regards to the second, then, depending on factors such as the severity of condition *x* and the cost of treatment, the medical provider should recommend that treatment begin. She will do so, for instance, when the condition is elevated risk of stroke, say from silent atrial fibrillation (30), and the treatment is a daily regimen of aspirin (31).

The Relationship Between Warrant and Veridicality

Not all veridical cognitive representations are warranted. Sometimes we get lucky. Here is an example. Taking a patient’s temperature is an instance of proper cognitive functioning; it is a successfully designed and vetted process that is reliable at

forming veridical cognitive representations when in the right environment. Part of being *in the right environment* includes using a properly calibrated thermometer.

Importantly, cognitive representations that are warranted are not always veridical. As we saw in part one, a provider who exercises proper cognitive functioning to screen for a condition will obviously still sometimes form non-veridical cognitive representations that are warranted (32). Though proper cognitive functioning is reliable (in the vetted- or designed-for environment) it is not infallible. Even if a process of proper cognitive functioning were to identify a condition with 0.99999 reliability there would still be room for error.

Warrant and Proper Diagnostic Functioning

A clinician must assert her findings to create a diagnosis. An assertion is not a cognitive representation.

Assertion =def. An information quality entity that is the concretization of a descriptive information content entity that is expressible by means of a sentence (CPO).

An information quality entity may be a check mark in a medical form on paper or on a screen. Or it may be an entire clinical note. ‘*x* is a Descriptive Information Content entity’ means *x* is an information content entity that describes some portion of reality.

As such, diagnoses cannot be formed through proper cognitive functioning: only the cognitive representations that diagnoses are based on can be.

Warranted Assertion

A diagnosis is an *assertion* that is the output of a diagnostic process. An assertion inherits the same warrant status as the cognitive representation it is based on. (By ‘based on’, here, we mean the relationship that holds between a mental representation and an assertion during a normal act of communication.) As such, if a cognitive representation is warranted, then an assertion based on that cognitive representation is warranted also. This provides us with a notion of ‘warranted assertion’:

Warranted Assertion =def. An assertion that is based on a representation that is warranted.

The importance of warranted assertions cannot be understated. They are the means by which we form cognitive representations about entities that we can only learn about through communicating with others; for example, it is through a warranted assertion that a patient forms a cognitive representation about her own diagnosis. It is also likely the means by which you know where you were born.

Warranted Clinical Picture

Furthermore, an instance of an RTW that qualifies as a clinical picture can be termed a ‘clinical picture that is warranted’.

Clinical Picture that is Warranted =def. A clinical picture that is a Representation that is Warranted (MCPO).

An instance of a ‘clinical picture that is warranted’ should be assumed as veridical in a diagnostic process.

Proper Diagnostic Functioning

We can use ‘warranted assertion’ and ‘clinical picture that is warranted’ to build the following definition of ‘proper diagnostic functioning’:

Proper Diagnostic Functioning =def. A diagnostic process that inputs representations that are warranted, including a clinical picture that is warranted, and, based on these inputs, outputs a warranted assertion to the effect that the patient does or does not have a disease, disease progression, disorder, or combination thereof of a certain type (MCPO).

Warranted Diagnosis

This allows us to define ‘warranted diagnosis’:

Warranted Diagnosis =def. A warranted assertion to the effect that the patient does or does not have a disease, disease progression, disorder, or combination thereof of a certain type and that is the output of proper diagnostic functioning (MCPO).

Future Work and Limitations

Creating definitions that allow data about the reliability of medical processes and diagnoses to be tracked paves the way for more and we believe better research about, among other things, patient safety. Warrant provides both a new dimension along which to collect data about patient safety and a direction to pursue in the forming of metrics for the quality of patient care.

That being said, the applications of warrant, especially in medicine, may be limited because of the uncontrolled environments that many medical processes unfold in. Without a controlled environment it is difficult to vet a process for a specific circumstance. This is something that is easier to achieve, for example, in an area like industrial design (where prototypes can be tested) than in medicine. Thus, it is still unclear how warrant would be applied in many clinical scenarios.

There are however clear applications for warrant in any field that relies on investigative processes to achieve goals. Intelligence, forensics, finance, and predictive maintenance are all areas where it is important to distinguish a warranted assertion from mere luck. We expect research on warrant to continue to be pursued in these fields in addition to being further developed for medicine.

Conclusion

The addition of MCPO to OGMS allows for warrant and luck to be distinguished, both in terms of diagnoses and in referent tracking. Scenarios hitherto distinguishable only by the veridicality of their diagnoses can now be distinguished in terms of whether each doctor’s assertion was or was not warranted. It is also now possible to track data in an RTS, not only in terms of its fidelity, but also in terms of whether or not an assertion should be (or should have been) trusted.

CPO itself is a new addition to the growing suite of ontologies that are compliant with Basic Formal Ontology (BFO). It represents the kinds of mental processes that are relevant to acts of reasoning, analysis, and investigation as they occur not only in medicine but also in other areas. It thus has applications to science in general, to intelligence analysis, finance, forensics, industrial design, preventive maintenance, software debugging, and many more (33). We believe that it will have applications also in Artificial Intelligence (AI) research, specifically in relation to the problem of what is called “Explainable AI”, by

providing a vehicle for incorporating explanation-related data into the training sets used for machine learning (34).

Address for correspondence

David Gordon Limbaugh – dglimbau@buffalo.edu

Bibliography

1. Hogan WR, Ceusters W. Diagnosis, misdiagnosis, lucky guess, hearsay, and more: An ontological analysis. *J Biomed Semantics* [Internet]. 2016;7(1):1–15. Available from: <http://dx.doi.org/10.1186/s13326-016-0098-5>
2. Limbaugh DG. Introducing the Cognitive Process Ontology [Internet]. Presented at the Seventh Clinical and Translational Science Ontology Workshop. 2019 [cited 2019 Apr 6]. Available from: <https://buffalo.app.box.com/s/yz5hbri1iir0rlku8lro53qqjrw8d12>
3. Hastings J, Ceusters W, Jensen M, Mulligan K, Smith B. Representing Mental Functioning: Ontologies for Mental Health and Disease. *ICBO 2012 3rd Int Conf Biomed Ontol* [Internet]. 2012;1–5. Available from: http://ontology.buffalo.edu/smith/articles/ICBO2012/MFO_Hastings.pdf
4. Plantinga A. *Warrant and Proper Function*. New York, New York: Oxford University Press; 1993.
5. Scheuermann RH, Ceusters W, Smith B. Toward an ontological treatment of disease and diagnosis. *AMIA Summit Transl Bioinforma* [Internet]. 2009;(i):116–20. Available from: <http://www.ncbi.nlm.nih.gov/pmc/articles/PMC3041577/>
6. Smith B, Ceusters W. Aboutness: Towards foundations for the information artifact ontology. *CEUR Workshop Proc*. 2015;1515:1–5.
7. Ceusters W, Manzoor S. How to track absolutely everything. In: Obrst L, Janssen T, W. C, editors. *Ontologies and Semantic Technologies for the Intelligence Community Frontiers in Artificial Intelligence and Applications*. Amsterdam: IOS Press; 2010. p. 13–36.
8. Ceusters W, Smith B. A Realism-Based Approach to the Evolution of Biomedical Ontologies. *AMIA 2006 Symp Proc*. 2006;121–5.
9. Ceusters W. Dealing with mistakes in a referent tracking system. In: KS H, editor. *Proceedings of Ontology for the Intelligence Community 2007 (OIC-2007)*. Columbia, MA; 2007. p. 5–8.
10. Seppälä S, Smith B, Ceusters W. Applying the realism-based ontology-versioning method for tracking changes in the basic formal ontology. *Front Artif Intell Appl*. 2014;267:227–40.
11. Limbaugh DG. The harm of medical disorder as harm in the damage sense. *Theor Med Bioeth* [Internet]. 2019;40(1):1–19. Available from: <https://doi.org/10.1007/s11017-019-09483-y>
12. Wakefield JC. The concept of mental disorder: On the boundary between biological facts and social values. *Am Psychol* [Internet]. 1992 Mar 1;47(3):373–88. Available from: <http://psycnet.apa.org/journals/amp/47/3/373.html>
13. Ceusters W, Elkin P, Smith B. Negative findings in electronic health records and biomedical ontologies: a realist approach. *Elsevier*. 2007;76:S326–33.
14. Bittner T, Smith B. A Theory of Granular Partitions. In: *Applied Ontology*. 2008.
15. Smith B. *Basic Formal Ontology 2.0* [Internet]. [github.com](https://github.com/BFO-ontology/BFO). 2015 [cited 2019 Apr 6]. Available from: <https://github.com/BFO-ontology/BFO>
16. Mungall C. *Relations Ontology: System* [Internet]. *Ontobee* [cited 2019 Apr 10]. Available from: http://purl.obolibrary.org/obo/RO_0002577
17. Smith B, Munn K, Papakin I. Bodily systems and the spatial-functional structure of the human body. *Stud Health Technol Inform* [Internet]. 2004;102:39–63. Available from: <http://www.ncbi.nlm.nih.gov/pubmed/15853263>
18. Searle JR. *Intentionality: An essay on the philosophy of mind*. Cambridge, UK: Cambridge University Press; 1983.
19. Smith B. *Fiat Objects*. *Topoi* [Internet]. 2001;20(2):131–48. Available from: <http://dx.doi.org/10.1023/A:1017948522031>
20. Chisholm RM. *The Primacy of the Intentional*. *Synthese*. 1984;61(1):89–109.
21. Kasmier D, Limbaugh DG, Smith B. Foundation for a Realist Ontology of Cognitive Processes. In: *Proceedings of the 10th International Conference on Biomedical Ontology (ICBO 2019)*. Buffalo, NY: University at Buffalo; 2019.
22. Searle JR. *Rationality in Action*. Cambridge, Mass., {USA}: The MIT Press; 2001.
23. Anscombe GEM. *Intention*. Ithaca, NY: Cornell University Press; 1963.
24. Smith B, Brogaard B. Quantum Mereotopology. *Ann Math Artif Intell*. 2002;36(1–2):153–75.
25. Meinong A. *Über die stellung der gegenstandstheorie im system der wissenschaften*. Leipzig: R. Voigtländers verlag; 1907.
26. Husserl E. *Logical Investigatinos*. Translated by J. N. Findlay from the second German Edition. New York: Humanities Press; 1970.
27. Fissette D. *Carl Stumpf* [Internet]. *Stanford Encyclopedia of Philosophy*. 2019. Available from: <https://plato.stanford.edu/archives/spr2019/entries/stumpf>
28. Marzolf AL, Peterseim MM, Forcina BD, Papa C, Wilson ME, Cheeseman EW, et al. Use of the Spot Vision Screener for patients with developmental disability. *J AAPOS* [Internet]. 2017;21(4):313-315.e1. Available from: <http://dx.doi.org/10.1016/j.jaapos.2017.04.008>
29. Peterseim MMW, Papa CE, Wilson ME, Davidson JD, Shtessel M, Husain M, et al. The effectiveness of the Spot Vision Screener in detecting amblyopia risk factors. *J AAPOS* [Internet]. 2014;18(6):539–42. Available from: <http://dx.doi.org/10.1016/j.jaapos.2014.07.176>
30. Page RL, Tilsch TW, Connolly SJ, Schnell DJ, Marcello SR, Wilkinson WE, et al. Asymptomatic or "silent" atrial fibrillation: frequency in untreated patients and patients receiving azimilide. *Circulation*. 2003;107(8):1141–5.
31. Zinkstok S, Vermeulen M, Stam J, de Haan R, Roos Y. A randomised controlled trial of antiplatelet therapy in combination with Rt-PA thrombolysis in ischemic stroke: rationale and design of the ARTIS-Trial. *Trials*. 2010;11(1):1–7.
32. Reus NJ, Lemij HG. Diagnostic Accuracy of the GDx VCC for Glaucoma. *Ophthalmology*. 2004;111(10):1860–5.
33. Marrin S, Clemente JD. Improving Intelligence Analysis by Looking to the Medical Profession. *Int J Intell CounterIntelligence* [Internet]. 2005;18(4):707–29. Available from: <http://www.tandfonline.com/doi/abs/10.1080/08850600590945434>
34. Landgrebe J, Smith B. Making AI meaningful again. *Synthese*. 2019;1–23.