Biomedical Ontologies

Barry Smith

Departments of Philosophy and Biomedical Informatics, University at Buffalo, Buffalo NY, USA

We begin at the beginning, with an outline of Aristotle's views on ontology and with a discussion of the influence of these views on Linnaeus. We move from there to consider the data standardization initiatives launched in the 19th century, and then turn to investigate how the idea of computational ontologies developed in the AI and knowledge representation communities in the closing decades of the 20th century, showing also how this idea has influenced SNOMED CT and other medical terminologies. Against this background we then show how the Foundational Model of Anatomy, the Gene Ontology, Basic Formal Ontology and other OBO Foundry ontologies came into existence and discuss their role in the development of contemporary biomedical informatics.

Keywords: biomedical ontology, Aristotle, Linnaeus, SNOMED CT, Foundational Model of Anatomy, Gene Ontology, Basic Formal Ontology

1   From Aristotle to Linnaeus

The term 'ontology' (*ontologia*) is a neologism, coined in the 17th century as an alternative to the Greek 'metaphysics', referring to the branch of philosophy that is engaged in the study of what exists or of 'what it is to be' in the most general sense. An ontology in the modern sense – which includes the biomedical ontologies discussed in what follows – is a taxonomy-like artifact that is structured in such a way as to be useful not only to humans but also to computers. The connection between the two senses of 'ontology' turns on the fact that the foundation of any taxonomy is the idea of what is general, and this idea lies at the center of both ancient metaphysics and contemporary ontology.

We are all implicitly familiar with the idea of what is general through our use of general terms (nouns and noun phrases) in both everyday and scientific contexts. In a tradition which goes back to Plato and Aristotle, such general terms are said to refer to *universals*, where each universal is associated with the collection of those particulars in reality which are its instances.[1]

A universal, sometimes also called a *kind* or *type*, is in some sense that which all its instances have in common. Aristotle held that we acquire knowledge of universals by observing the particulars that instantiate them. He himself displayed a lifelong interest in descriptive biology, and his metaphysics is

---

[1] 'Particular' means: an entity that exists in some unique region of space and time. Where universals are repeatable in indefinitely many instances; Particulars are unrepeatable.

rooted in his experience of the world as a place populated by instances of universals in the biological domain (Grene 1963, Lennox 1984, Leroi, 2014). Science, from Aristotle's point of view – which he saw as 'natural philosophy' – is thereby focused on what is qualitative in nature, which is to say on describing of how specific instances of universals such as *human* or *bird* are at specific times *running* or *sitting* or *perspiring* or *rotting*.

The process of mathematicization of science initiated by Galileo has of course diminished the attraction of qualitative, descriptive views of science of the sort embraced by Aristotle. But elements of such views still survive today, not least in the world of biomedicine. Even the idea that we gain knowledge of universals by examining their instances in reality is still very much a part of science. This is because the terms scientists use in formulating scientific laws are precisely terms representing universals. Laws link universals. And we gain knowledge of such laws by performing experiments in which we engage with the instances of the universals linked.

1.1   Science and common sense

Aristotle himself studied anatomy, astronomy, embryology, geography, geology, meteorology, physics and zoology.[2] His starting point for exploring such domains scientifically consists in establishing what kinds of entities they contain, thus in mapping the corresponding domain universals. This implies also creating a terminology consisting of the general terms with which to represent those universals. Understanding 'what it is to be' for the entities in a domain means to acquire knowledge of what the particulars instantiating any given universal in that domain have in common.

Aristotle assumed that human beings are in harmony with the world we find around us in the sense that, when we observe reality, we are able to grasp in our minds the universals instantiated by the things we see and touch. He does not seek deeper theories of what lies 'behind' or 'beyond' appearances, because to seek such theories would be to assume that the world is not as it appears to be. Certainly, there is room for error on the Aristotelian view. This relates, however, to particular perceptions only; it leaves the general features of perceptual knowledge untouched. Thus the Aristotelian – like the commonsensical – way of thinking about the world

---

[2] He also wrote on aesthetics, ethics, logic, metaphysics, government, politics, economics, psychology, rhetoric and theology (Sallam 2010). In his article on Aristotle's biology in the *Stanford Encyclopledia of Philosophy*, Lennox (2021) reports that in 1837 the anatomist Richard Owen 'introduced a survey of Aristotle's zoological studies 'by declaring that "Zoological Science sprang from [Aristotle's] labours, we may almost say, like Minerva from the Head of Jove, in a state of noble and splendid maturity."'

will never concede that it is false throughout. Error is a *local phenomenon*, it does not distort our *entire outlook*. Modern science, on the other hand (and the Platonic and Democritian philosophies it absorbed) postulated just such *global distortions*. (Feyerabend 1978, p. 148)

We return to this way of thinking in section 4.1, when we investigate the idea of a *canonical ontology* and thereby explore the reasons why biomedical ontologies have evolved in such a way as to include elements deriving both from Aristotelian thinking and from modern, data-driven science.

1.2   Aristotle on definitions

The treatises compiled by Aristotle's students and promulgated in his name constitute a virtual encyclopedia of Greek knowledge. Typically, each treatise begins with a review of earlier contributions to the study of the relevant domain and then presents Aristotle's own view with the aid of definitions of one or more central terms. Book II of the *Physics*, for example, begins with the definition of 'nature'; Book II of *On the Soul* with the definition of 'soul'.

The term 'definition' is used in a variety of ways in the Aristotelian corpus, but the core of his account rests on the idea that there is a hierarchy (or multiple hierarchies) of more and less general universals, with each universal in the hierarchy relating as *species* to its *genus* (in other words to its immediate parent in the hierarchy). A species is then defined by stipulating what is specific about the instances of its genus which makes them instances of the species.

The way this works is illustrated Figure 1, which is based on the tree-form representation of universals attributed to Porphyry, the author of an influential introduction to Aristotle's *Categories* from the 3rd century (see Barnes, ed. 2006), which was also the standard textbook in logic for at least a millennium after his death.

Take the full-size image from p. 5 of the pdf attachment here: https://classic.europeana.eu/portal/en/record/2048128/39246
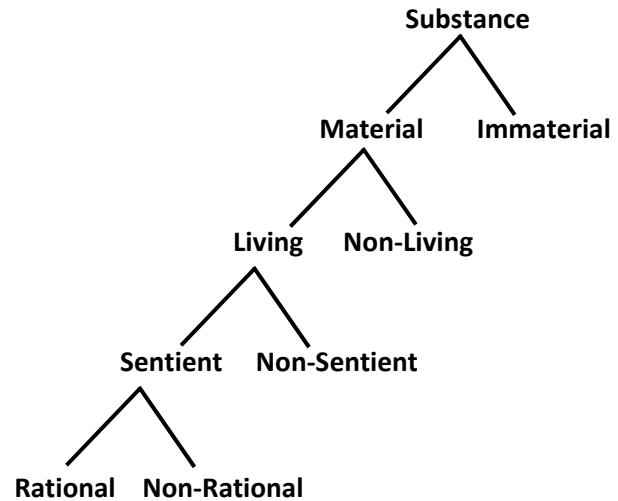

Figure 1: Tree of Porphyry (from Brunczwik 1748)

At the apex of the tree is the universal Substance, which divides into the branches of Material and Immaterial. The former is then divided into Living and Non-Living, and Living Substance is divided further into the Sentient and the Non-Sentient. Sentient Substance, finally, is divided into the Rational and the Irrational, which then provides the means to define a Human Being as a Rational Sentient Living Material Substance or – taking advantage of the way in which definitions at lower levels incorporate logically the contents of definitions at higher levels – that a Human Being is a Rational Animal.[3] This feature of incorporation allows successively ever more complex thoughts to be expressed by means of expressions that remain relatively compact at each stage. It is exploited in mathematics and in all the sciences to enable human beings to deploy ever more complex ideas in still understandable ways.

We can now understand why the Aristotelian rule for creating definitions is referred to by Aristotle's interpreters as definition *per genus et differentiam.* What this means is that a definition of a species consists of two parts specifying, respectively, its genus, and an associated specific difference. The instances of the species that is being defined are then all and only those instances of the genus that satisfy the specific difference. We shall return to this rule for formulating definitions in section 4.3 below.

1.3   Aristotle's table of categories

Aristotle created a number of more or less fragmentary examples of what we would nowadays call domain

---

[3] In Aristotle's *Posterior Analytics*, we find a definition of human being as 'animal, mortal, footed, biped, wingless' (92al-2) and alternatively as 'animal, tame, biped' (96b31); in *Metaphysics* Z 'biped' and 'animal' are said to 'constitute the definition of man' (1038a3). In his *Politics* (1.1253a) Aristotle's defines an animal of the human kind as a *zoon politikon*, or in other words as a social animal, an animal that naturally lives in a *polis*.

ontologies. At the same time he worked also on creating a top-level ontology, which would bind those domain ontologies together within a single logically and ontologically coherent framework by providing a common starting point – a common set of highest *genera* – for the definitions of their respective terms. The treatise compiled by Aristotle's students under the title *Metaphysics* addresses issues arising at this highest level of generality.

Aristotle's ideas on such matters are, as we shall see, still of considerable influence today. But he was of course not the only philosopher in the ancient world to have developed an influential approach to metaphysics. His competitors in this respect included his own teacher Plato, who conceived universals as inhabiting a timeless, ideal world, remote from the world of things we encounter in our everyday experience. They included also Democritus, who propagated the doctrine according to which all that exists are 'atoms and the void'; and Heraclitus, who embraced a process philosophy according to which 'everything is flux'. In the battle of grand theories, however, it is clear that Aristotle was – at least until the 17th century – the overwhelming victor. This was in part a result of the fact that Aristotelian metaphysics formed the foundations of Christian theology. But it also reflects the relative faithfulness of many of Aristotle's ideas, when compared with those embraced by his competitors, to much of what is accepted by common sense: that human beings (for example) exist, that they grow and develop through time, that they have qualities, habits, dispositions of various sorts, stand to each in various relations, occupy places, and so forth.

As will already be clear, one main starting point of Aristotle's metaphysics is the term 'substance'. At *Metaphysics* 1030b6–12, for example, he asserts that 'a proper definition states the essence of an entity, by which is meant "substance"'. It would be a difficult task to provide here an account of what Aristotle understood by this term. Suffice it to state here that in almost all the contexts where he provides examples of substances he refers to organisms (in particular to humans and horses).

Substance is in Aristotle's terminology one of the *categories*, which means that it is one of those most general universals which go to form his top-level ontology. They are what we call 'primitives' in virtue of the fact that they cannot be *defined* by the method of genus and specific difference because, lying at the very top, they lack a genus.

In some places Aristotle suggests that *substance* is the *only* category. There is, however, another strain in Aristotle's thinking according to which there are multiple universals on this highest level. First is the category of *substance*, which is marked by the fact that it and its subordinate universals are instantiated in

every case as a matter of necessity. What this means is that if a being is, for example, a horse at any time in its existence, then it is a horse at all times in its existence.

Second is a collection of *accident* categories, which are those highest-level universals which hold of their bearer not essentially but as a matter of accident. This means that they can be gained and lost during the course of the bearer's existence.
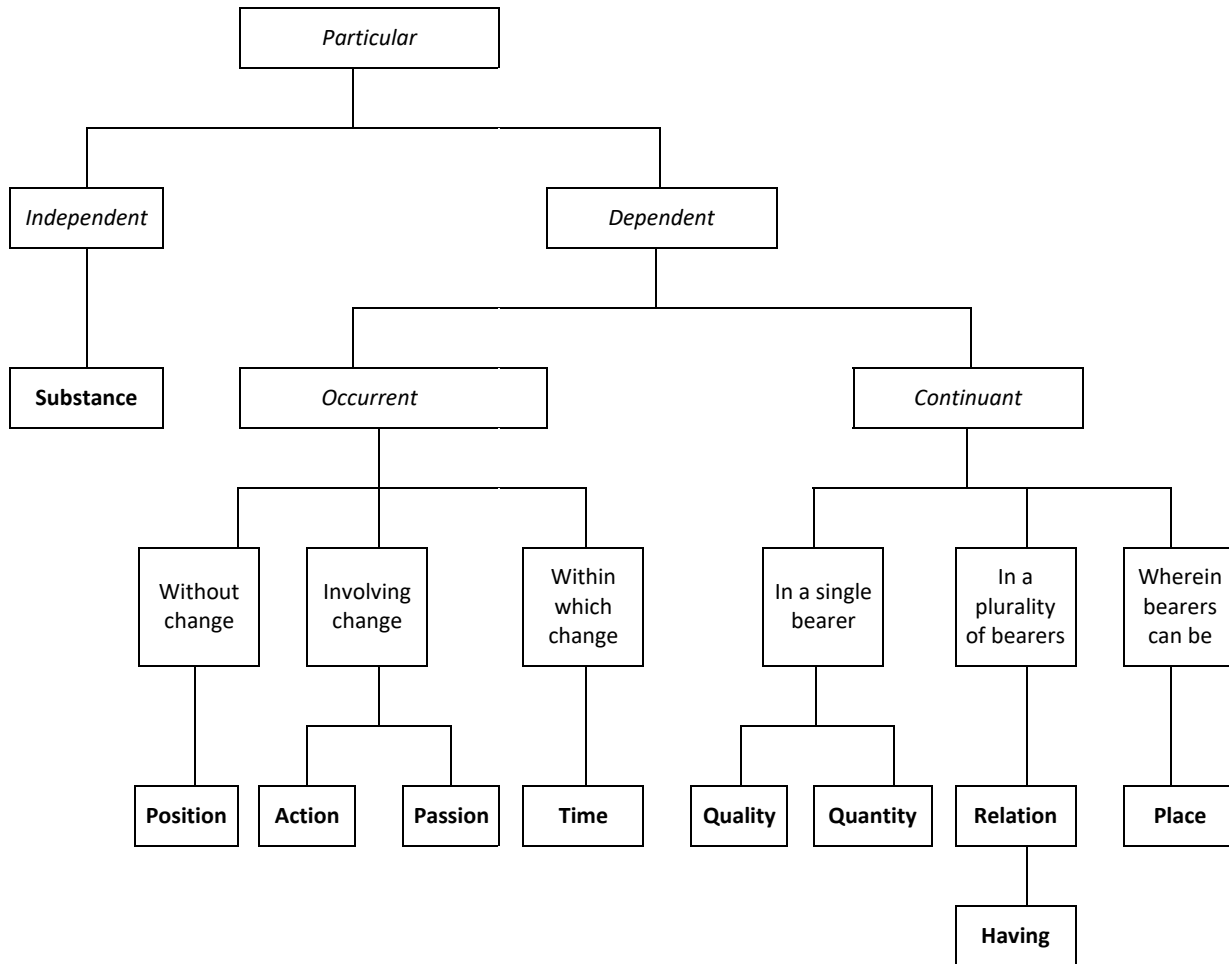

Figure 2: Aristotle's table of categories in the arrangement proposed in (Jansen, 2007). Categories are picked out in **bold.** The significance of the divisions marked by the italicized labels added by Jansen will become clear in section 6.3.

Aristotle presents different versions of his list of categories. Figure 2 presents the version described at *Cat.*, 2a34-35, 2b3-5, 2b15-17 in which, along with *substance*, nine *accident* categories are distinguished, examples for each of which are (from *Cat.*, 1b25-2a4):

Position: is lying, is sitting	**Time:** yesterday, last year	**Relation:** double, half, larger
Action: cutting, burning	**Quality:** white, grammatical	**Having;** has shoes on, has armor on

Passion: being cut, being burned   **Quantity**: four foot, five foot   **Place:** in the Lyceum, in the market-place

This table should not be seen as being complete. Further categories can be added to make more explicit. the categorical structure of the world. Indeed, as Jansen (2007) points out, the project presented by Aristotle in the *Categories*, 'seems to be rather a working report on an ongoing research project than something ultimate and completed.' One candidate additional category might be that of *hole*, an entity which can potentially be occupied or filled by something material, as for example a womb may be occupied by a fetus, or a trench be filled with water. In his treatise *On the parts of animals* Aristotle refers to the esophagus as 'the channel through which food is conveyed to the stomach' (III, 3). At the same time he refers to it as being 'of a flesh-like character'. It is then the flesh-like entity, rather than the channel, which forms a *part* in the sense relevant to this treatise. This turns on the fact that Aristotle's system of categories has no room for *places* or *locations* which are not occupied, nor for *channels*, or *cavities*, or *voids* – or, more generally, for *holes which are not filled* (see Casati and Varzi 1994 and 4.2 below).

The table of categories in Figure 1 may be extended also by recognizing universals in addition to those which have *particulars* as their instances. On some views, for example, there can be higher-level universals, which themselves have universals at lower levels as their instances. One example of such a higher-level universal would be the universal *universal* (Botti Benevides et al. 2019).

1.4   Linnaeus's *Scala Naturæ* and *Genera Morborum*

We can imagine Aristotle's ten categories as forming the top of a much larger hierarchy formed by universals belonging to successively more specific orders of being at lower levels. This idea was elaborated in the doctrines of the *Scala Naturæ* created by medieval philosophers, where each kind of entity is slotted into its own proper place in what is seen as a Great Chain of Being. This in turn formed one starting point for what we understand today as the tree of life, first systematically documented in Linnaeus's *Systema Naturæ*, the title of the 10th edition of which is:

> *System of nature through the three kingdoms of nature, according to classes, orders, genera and species, with characters, differences, synonyms, places.*

Here the 'three kingdoms' are *minerals*, *plants* and *animals*. In his *Genera Morborum*, Linnaeus applied this taxonomical method to the realm of diseases, distinguishing eleven classes, thirty-seven orders, and three hundred and twenty-five species of human disease, a small selection of which is presented in Figure 3.

Morbi

Nervii

Mentales
(Disturbance in the mental functions)

Ideales
(Of the judgment principally)

DELIRIUM. Being temporarily of acutely unsound mind, while having a fever

MANIA. Being chronically of universally unsound mind, in a furious manner

DAEMONIA. Being chronically of partially unsound mind in a demonic, devilish, furious, frightful, and terrible manner

MELANCHOLIA. Being chronically of partially unsound mind in an earnestly sad or sorrowful manner

ImaginarlI
(Imagination principally affected)

VERTIGO. False perception of dizziness, a circular whirling of the head

HYDROPHOBIA. Horror of drinking. Aversion to drink or water

CACOSITIA. Aversion to food

ANXIETAS. Aversion to worldly things, especially those that cause grief or a broken heart

SOMNAMBULISMUS. Sleepwalking

Pathetici
(Disorders full of pathos)

BULIMUS. Insatiable longing for food

POLYDIPSIA. Continuous longing for drink. Constant or abnormal thirst

EROTOMANIA. Longing for chaste lovers
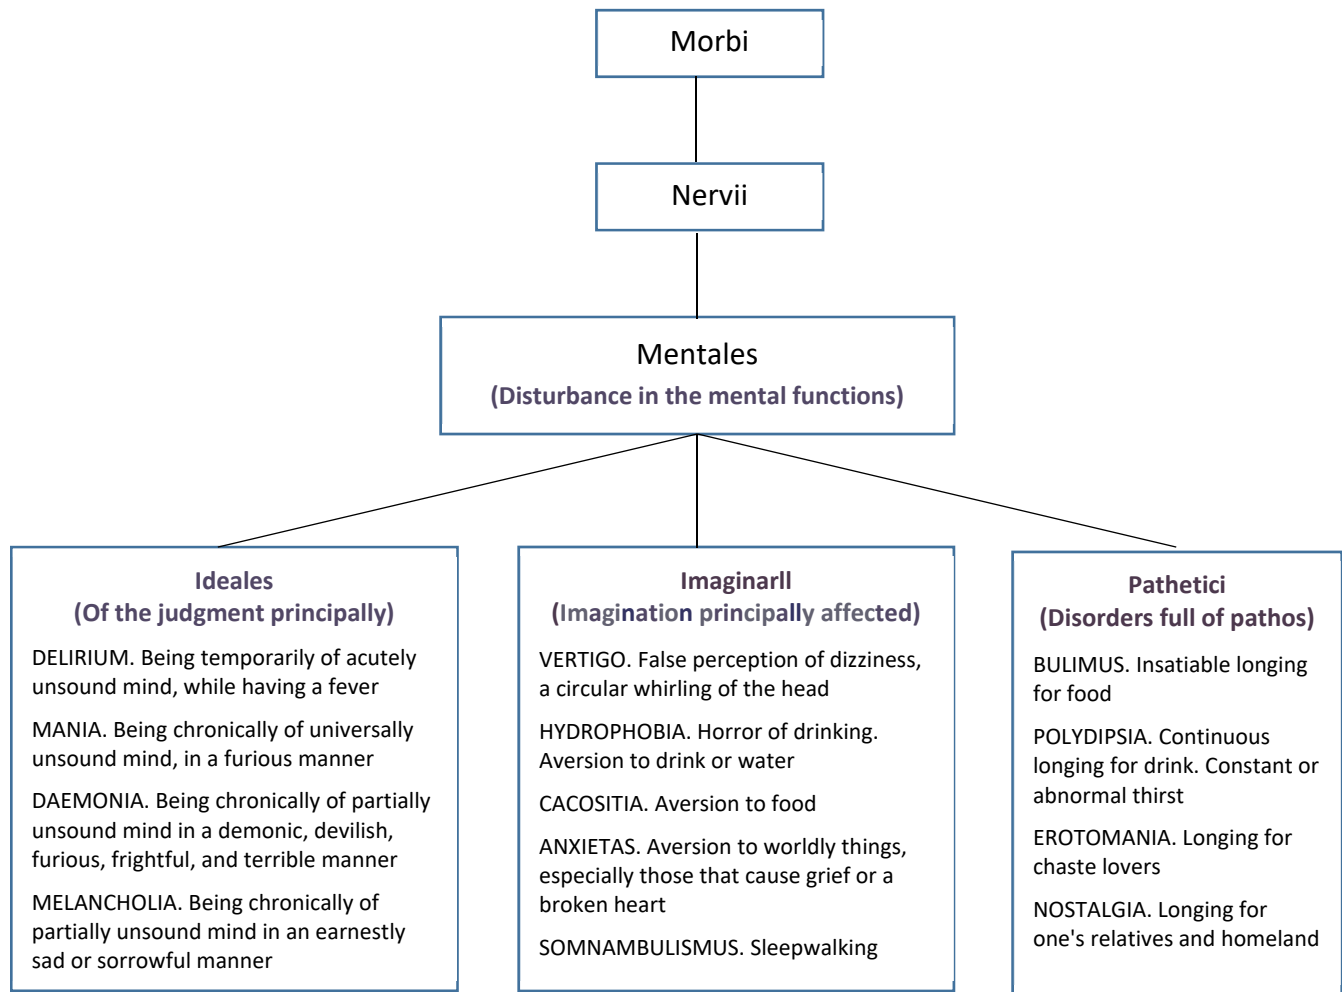
NOSTALGIA. Longing for one's relatives and homeland

Figure 3. Fragment from Linnaeus, *Genera Morborum* (1759), showing parts of Orders 1-3 of the 5th Class (Mental diseases), extracted from (Munsche and Whitaker 2012); see also (Egdahl 1907)

1.5 Physics

For some 2,000 years after its basic ideas were first set forth by Aristotle and his early commentators, the discipline of metaphysics advanced hardly at all, to the degree that it formed the central part of what was habitually referred to as *philosophia perennis*. The preeminent role of (Aristotelian) metaphysics in the pantheon of philosophical disciplines began to be challenged, however, from around the time of Descartes and Kant, who awarded pride of place to epistemology, which deals not with *being* but rather with *knowledge*. At the same time the preeminent role of philosophy itself among the disciplines was challenged by the rise of the empirical sciences.

During the 18th century, physics, in particular, evolved from its status as a qualitative and primarily descriptive discipline to become a quantitative and by degrees predictive discipline rooted in the mathematics-based study of observational and experimental results. This led, however, to an increasingly more urgent practical need for a standardized terminology for the communication of such results in a way that would allow international scientific and technological collaboration. Initiatives to address this need culminated in 1875 with the ratification of the Metre Convention, a landmark treaty which created the International Bureau of Weights and Measures (BIPM). This led in turn to the SI international standard system of units, which has served since 1960 as the universally accepted specification of the units of measure for physical quantities.

The SI standard incorporates controlled vocabularies not only for the representation of the physical units of measure but also for the kinds (universals) of physical magnitudes for whose measurement these units are employed. Both units and the corresponding magnitudes are divided into two classes of *base* and *derived* (see Figure 4), where the latter are defined in terms of the former. Force, for example, is defined as *mass times acceleration*, and a newton, the unit of force, is defined as *the force needed to accelerate one kilogram of mass at the rate of one metre per second squared in the direction of the force applied*.[4]

---

[4] New definitions of the base units were proposed in (BIPM 2019). Johansson points out in his (2021) that these new definitions appear to be circular, given that symbols for what is to be defined appear also in the defining expressions. He is indeed able to show that there are substantially non-circular definitions underneath these circularities, but he also identifies certain problems that still remain.
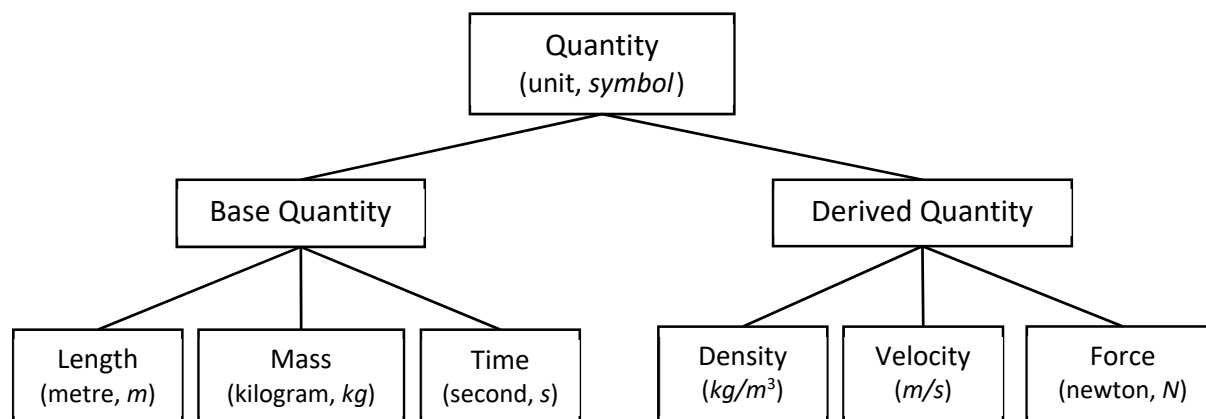
Figure 4: Examples of base and derived quantities and units in the SI system of units

The success of the SI system is one important reason why the need for ontology has made itself felt hardly at all in the domain of physics. But there is a second, and no less important reason, which turns on the fact that the kinds of entities and relationships with which physics is concerned are rigorously defined using mathematical equations.[5] The language of mathematics thereby serves in the physical domain as the lingua franca for the communication of scientific knowledge. These two factors together effectively ensure the mutual exploitability of both theories and results across all physical subdisciplines and all application areas, a feature of the physical domain that has shown itself to be indispensable to the success of all modern technology.

Similar advances in standardization of language and taxonomy were made also in chemistry, beginning with Mendeleev's Periodic Table in 1869 and culminating in the work of the International Union of Pure and Applied Chemistry (IUPAC). The latter has defined rules for naming and classifying organic and inorganic compounds and created thereby a similar level of mutual exploitability of chemical knowledge across all chemical disciplines, both pure and applied. Comparable advances in biological standardization in the wake of Linnaeus were primarily in the anatomical domain, with the *Nomina Anatomica*, dating from 1895, replaced by the *Terminologia Anatomica* in 1998, though it was clear already in 2001 that the latter left much to be desired from the perspective of the new, information-driven approaches in medical science (Rosse 2001).

---

[5] Less well defined are the terms used to represent both the different types of magnitude and the ontological relations between (a) these magnitudes themselves, (b) the symbols appearing in mathematical equations, and (c) the measurement results formulated in terms of SI units. For a treatment of these matters see (Landgrebe and Smith, in preparation).

2   Ontology, logic and artificial intelligence

2.1   Ontological commitment

Before moving to the special case of biomedical ontologies, and to the revolution in the handling of biomedical data which has come in the wake of the human and other 'model organism' genome projects, we need to take account of the rise of ontology in the computer science disciplines which occurred in the closing decades of the last century. The beginnings of this episode in the history of ontology can be traced to the 1948 paper "On what there is" by the prominent philosopher-logician Willard Van Orman Quine. This paper advances a new conception of the proper method of ontology, which at the same time gave new respectability to the term 'ontology' itself, not least through its influence on the work of John McCarthy, creator of the term 'artificial intelligence'.

According to Quine, the ontologist's task is to establish not *what there is* but rather *what are the kinds of entities to which scientists are committed in their theorizing*. The ontologist studies the world, on this conception, by drawing conclusions from the theories developed in the natural sciences.

Each natural science has its own preferred repertoire of entities to the existence of which it is committed, a repertoire which is revealed in its vocabulary. In applying his method for identifying ontological commitments, however, Quine turns not to the controlled vocabularies for physics or chemistry established by organizations such as the BIPM or IUPAC. Rather, he looks instead to a fictional future state of science, in which scientific theories would have been subjected to a quite different sort of regimentation, resting not on consistency in the use of *terms* but rather on the formalization of scientific *propositions* using the language of First-Order Logic (FOL).

2.2   First-order logic (FOL)

FOL is the logical framework of choice that is used by philosophers and others in the formalization of many different sorts of theories. It grew out of the attempts by Frege – the founder of modern logic – and then by Whitehead, Russell and others, to use logic to formalize the whole of mathematics. Building on the success of this work, Rudolf Carnap and other members of the Vienna circle conceived the project of a 'unified science' (Neurath, Carnap and Morris 1938-1968), based – in one version at least – on the idea of a FOL-based axiomatization of all scientific theories. In his *The Logical Structure of the World* (1928), Carnap used logic as a vehicle for the design of 'linguistic frameworks supplying all of the (names for) objects and concepts required by science' (Gottleb and Carus 2021).

The use of FOL brings great benefits, including:

- It allows us to capture in a single formal system many features of our reasoning not only in science and mathematics but also in our everyday affairs.

- It has a mature and sophisticated model-based semantics, which is used in all contemporary ontology applications.

- It exists in a number of semantically equivalent varieties of standardized syntax optimized for specific uses, including computational use in support of ontologies (ISO/IEC 24707, 2018).

When it comes to computer applications, however, FOL has the shortcoming that it is not decidable, which means there is no effective procedure for determining, given a consistent set T of FOL formulas – which might, for example, be the set of axioms of an ontology – whether an additional FOL formula A can be added to T in such a way as to preserve consistency, As we shall see in 2.5, it was the attempt to rectify this shortcoming which led to the development of OWL, the Web Ontology Language.

2.3   Ontological commitment again

At the heart of FOL is the idea of quantified statements, for example of the form $\exists xPx$ and $\forall yQy$, which mean, respectively: some value of the variable x satisfies the predicate P, and: every value of the variable y satisfies the predicate Q. Thus '$\exists$' stands for 'for some', and '$\forall$' for 'for all'.

To determine the ontological commitments of a scientific theory formalized in FOL, on Quine's approach, means to determine which entities belong to the ranges of those variables over which the formulas of the theory quantify – an idea that is captured in Quine's maxim: 'To be is the value of a variable'.

Imagine, for example, that we wish to formalize in FOL the sentence:

(1) Teco is a bonobo,

as part of the evidence base for a scientific theory. A typical FOL rendering of this sentence is:

(2) $\exists x$(is-a-bonobo(x) & x = Teco),

Or, translated back into English:

(3) there is some x which is a bonobo and which is identical to Teco.

It is here Teco alone which is the value of a variable in either (2) or (3). This means that it is to Teco alone that we are ontologically committed in making either of these assertions.

But could we not reformulate (1) in such a way that, say, bonobohood, would serve as value of a variable, for example by writing:

(4) $\exists P$(instantiates(P, Teco) & P = bonobohood)?

The problem here is that (4) is standardly interpreted as belonging not to first- but rather to higher-order logic, which is defined precisely by the fact that it allows quantification over predicates.[6] Quine's use of FOL to determine ontological commitment thus leads to an ontology in which only particulars exist – in other words to a *nominalist* doctrine, according to which *particulars* belong to the realm of *what exists*, but *generals* (universals) belong only to the realm of *what can be said.*

For Aristotle, in contrast, as for his successors in the camp of what we shall henceforth call 'ontological realism' (Smith and Ceusters 2010), there is in addition to Teco a second something that exists, and that contributes to making true the sentence 'Teco is a bonobo', namely some *feature* or *way of being*, some *species* or *natural kind* to which Teco belongs, or some *structure* or *pattern* of DNA in Teco's genome. From the ontologically realist perspective (1) then asserts a relation between Teco and this second something.

2.4   The Vienna circle project to unify science

In a series of groundbreaking contributions around the turn of the last century, Frege and others demonstrated that at least a large portion of *mathematics* could be unified by showing that the corresponding mathematical truths are ultimately truths of logic. The Vienna circle project was much less successful (Dahms 2018). But its underlying idea was influential nonetheless, above all through its effect on the work of John McCarthy and others in the field of artificial intelligence.

McCarthy was a leading figure in the first, logicist or 'symbolic' (also called Good Old Fashioned AI or 'GOFAI'[7]) wave of AI research, contributing *inter alia* to that strand in GOFAI which sought to use FOL-based approaches (including modal logic, the situation calculus, and so on) in order to capture in a formal

---

[6] Serious problems, such as Russell's paradox (Moore 1988), arise in a logic which allows unrestricted quantification over predicates. Smith (2005) describes a simple paradox-free alternative (non-standard) FOL reading of sentences like (4), which involves quantification not over *predicates* but over *universals*.
[7] A term coined by Haugeland in his (1985).

way information about the world, for example to support the building of an intelligent robot programmed with the ontology of common sense that is used by humans.

It was in this context that McCarthy recognized the overlap between work done in philosophical ontology and the activity of building logical theories for AI systems. McCarthy affirmed already in 1980 that builders of logic-based intelligent systems must first 'list everything that exists, building an ontology of our world' (McCarthy 1980). This view, inspired by McCarthy's reading of Quine[8], was advanced also by McCarthy's collaborator Patrick Hayes in his "Naive physics: ontology for liquids" (1978), the first work to use in its title the word 'ontology' in the new sense of the term that is aligned to the use of computers.

As Hayes writes, looking back on the question of early uses of ontology in AI:

> As far as I recall, my use in the title of the 1978 paper was original. I used it deliberately to suggest/imply that the KR problem in AI was connected with philosophical ontology. The background to this was my reading Carnap's *Logical Structure of the World* as an undergraduate, probably some time in 1964. Reading this blew my mind and first got me excited about the idea of using logic to describe the real world. When I got into AI and read McCarthy's "Situations, actions and causal laws" … I was immediately struck by the similarity both of goals and even in places of formal (what would now be called 'ontological') techniques. (Hayes 2013)

2.5   The Web Ontology Language (OWL)

The sort of ontology practiced by Hayes was of considerable influence, as is illustrated for example in the Hobbs and Moore collection entitled *Formal theories of the commonsense world* published in 1985. But work of this sort in AI has of course has been eclipsed in recent years by an approach centered around deep neural networks and related stochastic approaches.[9] In the world outside AI, however, the work of McCarthy and Hayes was just one initial strand in the burgeoning of work in ontology and knowledge representation (or 'KR') that took place from the 1980s onwards, in a movement which received considerable further impetus from the release, in 1999, of Protégé 1.0, a freely available software tool for

---

[8] To quote from McCarthy (2000): 'In philosophy, ontology is the branch that studies what things exist. W.V.O. Quine's view is that the ontology is what the variables range over. Ontology has been used variously in AI, but I think Quine's usage is best for AI.'

[9] Something like GOFAI may be enjoying a mild reawakening, for example, in the recent book on general AI by Marcus and Davis (2019), where a logic-based framework of common-sense AI is seen as a possible avenue to allow the gluing together of various narrow stochastic AIs within a single general framework.

the building of ontologies (and applied not least in the biomedical realm).

At about this time, the drive to find a computationally tractable language for the purpose of developing formal ontologies led to the exploration of subsets of FOL especially in the family of so-called Description Logics (Baader et al. 2008). This culminated in the standardization by the World Wide Web Consortium (W3C) of the Web Ontology Language or 'OWL' in 2004, which is currently the most widely used logical framework for ontology development (https://www.w3.org/OWL/).

The new ontology languages were optimized for computer use, though unfortunately this came at the price of sacrifices in expressiveness (Schulz et al. 2009, Ceusters and Smith 2003). One result of the new ease with which ontologies could be built led accordingly to an upswell of overlapping and often mutually inconsistent efforts, as different groups sought in different ways to overcome the barriers of low expressivity. The results are illustrated for example by the way in which the fashion for agent-based modelling around the turn of the millennium led to the development of some 30 'agent ontologies', under headings such as *action*, *actions*, *activity*, *agent*, *agents*, *agent architecture*, *agent communication*, *agent framework*, and so forth.[10]

Many in the KR community seem to have assumed that the development of many, many ontologies is something positive. It is necessary only that each of the ontologies developed should be associated in the minds of its developers with some potential use case – an idea promulgated for example by Noy and McGuinness in their influential ontology manual (2001), for example in their assertion that 'Deciding whether a particular concept is a class in an ontology or an individual instance depends on what the potential applications of the ontology are.'

The multiplication of ontologies derived also from the fact that, during the period in question grant funding was available for the development only of *novel* ontologies. Efforts to establish the sorts of principles of best practice that might point ontology in a more scientific direction were, on the other hand, neglected(. Ontology development in this period, not surprisingly, gained a bad reputation – the results, it was said, were 'brittle', 'unsustainable', 'unscalable', and rested on over-simplified (and thus often unscientific) models of the relevant subject-matters.

---

[10] These are listed in the catalog of ontologies developed using DAML, the DARPA Agent Markup Language, (http://www.daml.org/ontologies/, last accessed July 30, 2021), which was one of the precursors of the OWL language.

3   The concept orientation

The idea that we should seek to focus on the development of well-grounded reference ontologies was rejected, in many circles, since it was taken to imply that the authors of such ontologies would aspire to the possession of some kind of God's eye perspective. Since such a perspective is unavailable, it was assumed that the best we can achieve is ontologies based on the ontological commitments of specific languages, theories, systems of beliefs, or what we shall encounter below as 'conceptualizations'.

The discipline of knowledge representation has to deal, after all, not with reality, but rather with the knowledge (and thus the concepts) in people's minds. It is thus very likely to lead to the situation in which there is a plurality of ontologies covering the same topic, since there is plurality of knowers whose knowledge is being captured and a plurality of uses to which this knowledge is being put.

KR researches who invested in trying to develop more generally applicable methodologies focused their efforts especially on approaches to ontology development based on model-theoretic semantics, making important contributions not least to the development of the ideas underlying OWL (Baader et al. 2008). There, too, however, there often reigned a presupposition to the effect that we can never understand what a given language or theory is really about – thus we can never compare an ontology to any independent reality beyond. We can, though, build abstract (set-theoretic) models, which we can usefully manipulate, for example in checking the consistency of a set of definitions and axioms.

From around 1990 there were, however, some few who acknowledged the need for a common framework of high-generality terms, axioms, and definitions which would promote ontology reusability by building ontologies in a way that would ensure correspondence to the things and processes in reality they were designed to represent. Thus they embraced the need for just one *agent/action* ontology, just one *software ontology*, and so forth, and they started to ask questions like: 'What is an object/process/attribute/relation? What is a transaction, a person, an organization? How do they depend on each other? How are they related?' (Smith and Welty, 2001)

It was in this way that early examples of top-level ontologies began to be developed with the goal of unifying and systematizing development of domain ontologies at lower levels, though at first with little traction in the broader world of KR scientists in general or of ontology developers in particular.

In the year 1993 Tom Gruber published the first credible attempt at defining what an ontology – in the new, computer science sense of this term – is:

An ontology is an explicit specification of a conceptualization. The term is borrowed from philosophy, where an ontology is a systematic account of Existence. For knowledge-based systems, what "exists" is exactly that which can be represented. When the knowledge of a domain is represented in a declarative formalism, the set of objects that can be represented is called the universe of discourse. This set of objects, and the describable relationships among them, are reflected in the representational vocabulary with which a knowledge-based program represents knowledge. (Gruber 1993)

Gruber's definition was rapidly adopted in multiple sub-fields of computer science. The definition still, of course, leaves open what a 'conceptualization' might be, and how the crucial sentence 'For knowledge-based systems, what "exists" is exactly that which can be represented' is to be interpreted.[11] What can be said with certainty, however, is that the vast majority of those following in Gruber's footsteps assumed that his definition meant that an ontology is a representation not of *entities in reality* but of *concepts*. represented

In the same year, however, there was held the first International Workshop on Formal Ontology in Conceptual Analysis and Knowledge Representation in Padua, which brought together ontologists such as Gruber from the computer science field with philosophers embracing a more traditional approach to the understanding of the meaning of "exists".[12]

3.1   The case of SNOMED

The dominance of what is sometimes referred to in terminology circles as the 'concept orientation' made itself manifest especially in the rapidly expanding field of medical terminology research, which is also the field in which the problems associated with concept-based approaches have been deliberated upon most persistently. The immediate need of medical terminologists was to address the problems raised by the huge numbers of synonymous and quasi-synonymous terms in the various medical disciplines. This problem is of minor importance where only humans are involved in the application of medical terms. Human experts in any given field know full well how to handle synonymy. When computers enter the

---

[11] This sentence is, for someone with a background in philosophical ontology, deeply problematic.
[12] This event served as the launchpad for the subsequent FOIS (Formal Ontology in Information Systems, https://iaoa.org/index.php/fois/fois-history/) conference series, which remains the premier event in ontology (science). Its organizers, Nicola Guarino and Roberto Poli, describe the meeting as 'probably the first interdisciplinary initiative in this area, aiming to explore the connections between philosophers belonging to the tradition of Brentano and Husserl, philosophers of language, and people working on principles of knowledge representation and engineering.' (Guarino and Poli 1995, p. 624). On Guarino's unique intermediating role in these developments see his (2017).

picture, however, matters are different.

The (for a long time) most influential solution to the problem of how computers are to handle synonymy was presented by Cimino in his classic paper on "Desiderata for Controlled Medical Vocabularies in the Twenty-First Century" (1998). Cimino's thesis is that, in the medical domain,

> most systems that report using controlled vocabulary are actually dealing with the notion of concepts. Authors are becoming more explicit now in stating that they need vocabularies in which the unit of symbolic processing is the concept – an embodiment of a particular meaning. Concept orientation means that terms must correspond to at least one meaning ("nonvagueness") and no more than one meaning ("nonambiguity"), and that meanings correspond to no more than one term ("nonredundancy").

The definition of 'concept' here provided – as 'an embodiment of a particular meaning' – is, however, difficult to parse. Embodied in what? Moreover, to understand this definition, we would need to understand already the meaning of the term 'meaning', which philosophers have long recognized as a difficult nut to crack.

The central problem not addressed by Cimino, however, is that the meaning of the word 'concept' itself varies drastically from one community to the next (and sometimes from one paragraph to the next). Certainly, there are acceptable and scientifically well-defined uses of this term, for instance in the study of conceptual change in developmental psychology (Carey 2009). But constructing a medical terminology – for example a gargantuan terminology such as SNOMED CT[13] – is not an exercise in empirical psychology.

Unfortunately, SNOMED CT is itself pervaded (quite literally from top to bottom) by the concept orientation, as is seen in the fact that SNOMED CT Concept is the topmost node of the entire SNOMED CT taxonomical hierarchy and thus subsumes the entire SNOMED CT universe. Thus we have:

> clinical finding *is_a* SNOMED CT Concept,

---

[13] Formerly known as the 'Systematized Nomenclature for Medicine', though this label has been dropped since the SNOMED International Organization no longer sees SNOMED as a nomenclature. 'CT' stands for 'clinical terms'. The first versions of SNOMED were developed by Roger A. Côté, originally under the name 'Systematized Nomenclature of Pathology' (SNOP). When Côté visited the Vatican to present a copy of the four volumes of what was then SNOMED in book form to the Vatican Cardinal Librarian Luigi Poggi, he also met Pope John Paul II, who is said by some to have remarked, 'Do you realize that this spells "DEMONS" backwards?'

environment or geographical location *is_a* SNOMED CT Concept,

pharmaceutical product *is_a* SNOMED CT Concept,

social context *is_a* SNOMED CT Concept,

and many more. Given, therefore, the standard reading of *is_a* – the reading accepted in other concepts by SNOMED CT itself – it follows that when you are suffering from a headache than what you are suffering from *is_a* (is a) SNOMED CT Concept.

The SNOMED CT community does indeed go to great lengths to explain what it means by 'concept', for example in its Editorial Guide from 2018[14], where we read:

(5) SNOMED CT concepts should name classes of things.

(6) A concept is defined as a clinical idea to which a unique concept identifier has been assigned. Concepts are associated with descriptions that contain human-readable terms describing the concept.

(7) A term is defined as a human-readable phrase that names or describes a concept.

We note that in (5) concepts are identified as *names*, while in (6) they are identified as *clinical ideas* and explicitly distinguished from *terms*. In (7), terms are identified as names or descriptions of concepts.[15] Thus the term 'Myocardial Infarction' (for example) describes not myocardial infarction (the clinical phenomenon that appears in certain patients) but rather the *concept Myocardial Infarction*. Terms in medical terminologies à la Cimino, are *about* Cimino's 'embodied meanings'.

Already in 2010 SNOMED CT had responded to criticisms concerning the problems created by such ambiguities in its use of the term, by publishing as the Glossary entry for 'Concept' in its Technical Reference Guide the following warning:

Concept

An ambiguous term. Depending on the context, it may refer to:

---

[14] https://confluence.ihtsdotools.org/download/attachments/75337342/doc_EditorialGuide_Current-en-US_INT_20180731.pdf/.

[15] Friends of the concept orientation can of course criticize those who use the term 'term' as an alternative to 'concept' by pointing out that the term 'term', too, can be misused in a way that involves a blurring of the distinction between 'entities of the domain' and 'entities of language'. This occurs, for example, when (Bauer et al. 2008) describe a new software tool called Ontologizer as 'a Java application that can be used to perform statistical analysis for overrepresentation of Gene Ontology (GO) terms in sets of genes or proteins derived from an experiment.'

- a clinical idea to which a unique ConceptId has been assigned;
- the ConceptId itself, which is the key of the Concepts Table (in this case it is less ambiguous to use the term "concept code");
- the real-world referent(s) of the ConceptId, that is, the class of entities in reality which the ConceptId represents (in this case it is less ambiguous to use the term "meaning" or "code meaning")

But as (Ceusters 2011) notes, merely pointing out this problem does not imply that the problem has been solved. Indeed, the very same Glossary still contains, for example, an assertion to the effect that a SNOMED CT term is '*a text string that represents the Concept*'.[16]

> So what is it then that is represented by a term: (1) the clinical idea, (2) less likely, but nevertheless in line with the expressed ambiguity – the ConceptID, or (3) the real-world referent(s)? The same question must then be asked for the several hundred occurrences of the word '*concept*' throughout the SNOMED CT documentation. In some cases, readers can infer from the context which meaning is intended, but in most cases, only the SNOMED CT authors can provide the answer by rewriting the entire documentation. (Ceusters 2011; see also section 2.1 of Ceusters 2021)

The most recent SNOMED CT documentation reveals no advance on this front, suggesting now that 'concept' and 'term' are interchangeable, and introducing an entirely new characterization of concepts as 'clinically relevant thoughts':

> Concepts, or terms, are represented by unique codes and human readable descriptions. Each concept is a unique clinically relevant thought, across a wide range like *abscess*, *zygote*, *measurement procedure*, or *substance*, as examples. (https://www.imohealth.com/ideas/article/snomed-ct-101-a-guide-to-the-international-terminology-system/)

Curing an abscess, then, means curing a clinically relevant thought.

Attempts have been made to define the term 'concept' (and related terms such as 'conceptualization' or 'conceptual entity') also in other medical terminology circles. But these attempts, too, have been unconvincing. Yet in spite of all the problems,[17] there is still, after more than 20 years of concept-oriented ontology and terminology research in medicine and in other fields, no definition of 'concept' that can be recommended for use. The safest course, therefore, is to recommend in ontology circles that 'concept' not

---

[16] https://confluence.ihtsdotools.org/download/attachments/18780048/TechnicalReferenceGuide_Current-en-US_INT_20100131-significant-changes.pdf.

[17] See for example (Ceusters, Smith, Kumar, Dhaen, 2004 and 2004a), (Bodenreider et al. 2007), (Bona et al. 2019), (Ceusters, Mullin 2019), (Guarino et al. 2009).

be used at all.

3.2   Defending the concept orientation

What, then, is to be said in *favor* of Cimino's idea that vocabularies should be development in which the unit of symbolic processing is the concept? A number of arguments to support this view have been proposed, and these arguments are discussed extensively in (Smith, Kusnierczyk and Ceusters 2006). Here we offer some examples of these arguments in abbreviated form.

First is what we might call the *argument from intellectual modesty*, which can be summarized as follows: It is medical domain experts who must answer for the truth of whatever theories the medical terminology is intended to mirror. Since domain experts themselves will sometimes disagree, any given terminology should embrace no claims as to what the world is like, but reflect, rather, some abstract conceptual substitute derived, somehow, from the different concepts used by different experts.

Against this, however, it can be pointed out that communities of experts working on common domains in the medical as in other scientific fields in fact accept a massive and ever-growing body of consensus truths about the entities in these domains. Where conflicts do arise in the course of scientific development, these are highly localized, and pertain in medicine primarily to specific mechanisms, for example of drug action or disease development. But the latter can serve as the targets of conflicting beliefs only against the background of a large body of shared presuppositions.

Moreover, we can think of no scenario under which it would make sense to postulate special entities called 'concepts' as the entities to which terms subject to scientific dispute would refer. For either, for any such term, the dispute is eventually resolved in favour of one side or the other, and then it is the corresponding real-world entity that has served as its referent all along. Or it is ultimately established that the term in question is non-designating, and then this term is no longer a candidate for inclusion in whatever is the active version of the relevant terminology.

The proposal from SNOMED and other defenders of the concept approach, however, is much more radical. It is that we provide guaranteed referents called 'concepts' not only for terms identified as problematic but for every single term in the terminology. The realist alternative solution is in contrast more modest. It is simply to treat any terminology as subject to a process of evolution (Smith 2006 and 2008). Even terms still subject to dispute can be incorporated into the terminology alongside other terms already accepted as referring to real-world entities, but in such a way that they are marked as being still

subject to dispute and thereby treated logically as unavailable for use in certain sorts of inferences.[18]

Another argument in favor of the concept orientation is the *argument from negative findings*. Consider, for example, the case where a clinician reports a finding of 'absent nipple'. The defender of the concept orientation will argue that there is no real-world entity denoted by this expression, and therefore that the expression must refer to something like a concept. Certainly clinicians need to record such findings. But from the realist point of view, their findings are precisely that a nipple is absent; not that a special kind of ('absent', conceptual) nipple is present (Ceusters, Elkin, Smith 2007).

Next is the *argument from hypertension*. The subject-matters of biology and medicine are, it is held, replete with entities which do not exist in reality but are rather convenient fictions, as in the case of the entities designated by expressions such as 'hypertension' or 'obesity' or 'abnormal curvature of spine'. Such abstractions are, it is held, 'mere concepts', since they reflect not joints in reality but rather certain more or less arbitrary human decisions (which may indeed vary over time).

From the realist point of view, in contrast, such terms are analogous to, for example, 'Poland' or 'the middle ages'. That is, they represent full-fledged entities in the real world, but they are entities whose *boundaries* are precisely the results of decisions made by human beings. The metre, the kilogram, and the second, too, are the results of fiat demarcations of this sort, and so also is hypertension, which rests on a (periodically readjusted) fiat threshold established by consensus among physicians.

Finally, we can mention what we might call the *argument from administration*, which asserts that, for many of the purposes for which medical terminologies are devised, a focus on something like Aristotelian universals would be far too restrictive. Consider the ICD (International Classification of Diseases) term:

> (8) Tuberculosis of adrenal glands, tubercle bacilli not found (in sputum) by microscopy, but found by bacterial culture.

There is no *ontological* difference between tuberculosis diagnosed by microscopy and tuberculosis diagnosed by bacterial cultural, any more than there is such a difference between tuberculosis diagnosed on a Wednesday and tuberculosis diagnosed on a Thursday, or while wearing socks. For the administrative purposes of the ICD and its many users, however, it is important that differences such as those expressed in (8) should be accounted for terminologically.

---

[18] Compare the strategy based on the Modal Relation Ontology outlined in (Rudnicki 2019).

Perhaps, then, a term like (8) should be acknowledged as representing a *concept*? But no, and yet again no. (8) is about tuberculosis; indeed it is about *tuberculosis of adrenal glands* (and thus it is also about *glands*) and similarly it is also about *sputum* (Ceusters and Smith 2015). Wherever (8) occurs in any document prepared by some clinician user of ICD, we can be sure that the author of this document is quite clear in her mind that that is what this term is about. She is not using this term to refer to, for example, someone's clinical thought.

Combination terms like (8) involve the mixing together of properly ontological terms (representing universals in the domains of disease, anatomy, and species taxonomy) with epistemological terms relating to how particular instances of a disease were discovered to exist, a matter of how, in this case, 'reality is perceived, measured, and understood by health professionals.' (Bodenreider et al. 2004) Manipulation of such combinations is an indispensable part of information driven medical research and so there is certainly no objection to developing ontologies whose terms would capture distinctions such as that between a *bacterial culture test* and a *microscopy assay.* Such ontologies are indeed already being developed (see for example (Bandrowski 2017, Gurcan 2017)).

Needed, too, are ontological resources which allow the representation of what we might think of as administrative aspects of medical or scientific discourse. Consider a term such as:

(9) subject in clinical trial SwEaTB for Diagnosing of Acute Tuberculosis

Here we have a term that is not intended to represent a universal or the extensions of a universal (in anything like the Aristotelian sense). Rather, it is intended to capture what we can think of as a convenience combination (also called 'defined classes' (Arp et al. 2015)). We then need to distinguish two kinds of ontologies; what we might call 'reference ontologies', on the one hand (dealt with in sections 4–6 below), which are designed to be of global reach and application neutral and thus to capture universals, together with, on the other hand, 'application ontologies', which result from the combination of terms from reference ontologies together with terms such as (9) developed for local, application-specific purposes (Shaw et al. 2008). Building this sort of bridge between application ontologies and reference ontologies is by no means a trivial matter (Schulz et al. 2021). Experience strongly suggests, however that it is the only course that will avoid the sort of destructive proliferation witnessed in the ontology field in the 1990s.

4 The Foundational Model of Anatomy

From around 2013 there has been occurring a paradigm shift in biomedical terminology and ontology development circles (Schulz et al. 2013) away from the concept orientation. Attempts have since then been made by for example to create an ontologically robust upper-level structure for SNOMED CT (Schulz et al. 2015). The first biomedical ontology to be developed in the spirit of ontological realism, however, came much earlier. This was the Foundational Model of Anatomy (Rosse and Mejino 2003), which addresses the need for a generalizable anatomy ontology that could be used and adapted by any computer-based application that requires anatomical information. The FMA is a domain ontology that represents a coherent body of explicit declarative knowledge about human anatomy. It has the potential for enabling many digital applications involving reference to and manipulation of information about anatomical entities, for instance in educational applications, particularly in the domain of distance learning, and as the basis for computer models, for example in the area of human anatomical development. Its ontological framework can be applied and extended to all other species, and it provides the template for CARO (the Common Anatomy Reference Ontology) (Haendel et al. 2008) and much of the content for the UBERON integrated cross-species anatomy ontology (Mungall et al. 2012).

4.1 The FMA as a canonical ontology

The FMA is very large, comprising some 120,000 terms and over 2.1 million assertions of relationships between the entities represented by these terms. Yet for all its size, it addresses only what we can provisionally think of as the normal healthy human being. This is because the attempt to do justice to, for example, all possible types of variants and pathologies affecting human anatomy would lead to an explosion in size which would make the result unmanageable and probably also of little utility.

Rather, the strategy of the FMA is to constitute a canonical ontology, ranging over types (universals) which are in a sense idealizations of the human organism's body and of its component parts. More precisely, the FMA represents all material objects, all portions of substance, and all spaces that result from the coordinated expression of the structural genes of the human organism (in a good approximation: all parts of a normally developed human body, from the macromolecular to macroscopic levels of granularity).

Canonical anatomy is thus distinct from *instantiated anatomy*, which comprises anatomical data about individual organisms. though it does not itself comprise such data, the FMA serves as a valuable framework for capturing and storing instance-level anatomical data in computable form, by providing the

vocabulary for describing those ways in which instantiated anatomical structures can depart from what is canonical (Rosse and Mejino 2007).
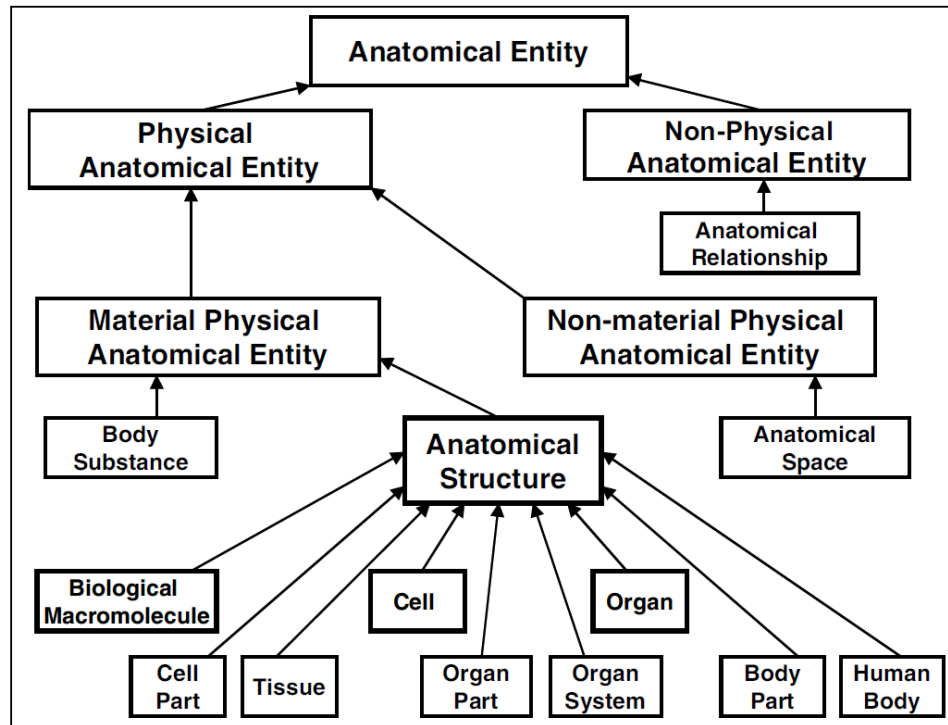

Figure 5: Upper-level structure of the Foundational Model of Anatomy

4.2   Canonical relations

To capture the meanings of its terms in a computer parsable form, the FMA ontology, like the other biomedical ontologies which have followed in its wake, consists primarily of statements of the form 'A rel B', where 'rel' stands in for a relational expression such as '*constitutional_part_of*', '*has_regional_part*', '*is_member_of*', '*is_tributary_of*', and most importantly '*is_a*' (meaning either: is a subtype of, or is a subclass of), which is the relation used to determine the backbone taxonomy of every ontology. The upper part of the FMA backbone taxonomy is represented in Figure 5. ('Anatomical Space', here, refers to the sorts of channels and cavities referred to in section 1.3 above.)

The now standard way of defining *part_of* and other such relations between types in ontologies is by reference to the relations that hold between the corresponding instances of these types, and using the FOL device of quantification. Two major types of definitions are then required, for relations between types of processes, on the one hand, and between objects and their parts and aggregates, on the other. For the

former we have:

> X *has_part* Y =def. for any instance *x* of the process type X there is some instance *y* of the process type Y which is such that *y* **instance-level-part-of** *x*.

> Example: Development of Spleen *has_part* Development of Splenic Lobules.

For the latter, however, we need to take account of time, in order to do justice to the fact that objects can gain and lose parts while preserving their identity:

> X *has_part* Y =def. for any time *t* and for any instance *x* of the object type X at *t*, there is some instance *y* of the object type Y at *t* which is such that *y* **instance-level-part-of** *x* **at** *t*.

> Example: Set of Teeth *has_part* Left Maxillary Dentition.

On the basis of a set of definitions modelled on the above, a group of leaders of different groups of biomedical ontology developers, including not only the FMA and Gene Ontologies but also the GALEN group around Alan Rector, developed the Relation Ontology (RO) (Smith, Ceusters et al. 2005). This provides a basis for the formal definition of the relations used by biomedical ontology developers in a way that promotes interoperability of the ontologies which use them and thereby allows new types of automated reasoning both within and across ontologies.[19]

In some domains, universal parthood assertions of the above-mentioned sort are unproblematic. This holds for example of relations between molecules and their parts in chemistry. It also holds for certain anatomical relations, such as: Neuron *has_constitutional_part* Plasma Membrane.[20] In biology in general and in medicine in particular, however, such universal assertions are problematic because there are variants (for example some humans have a middle lobe of left lung) and pathologies (for example tumors, or missing teeth). Many assertions of relations in the FMA hold, therefore, as a matter of canonical ontology.

This means that an FMA statement such as:

> (1) Skin of Thumb *has_regional_part* Nail of Thumb

---

[19] The current version of the Relation Ontology can be found at http://www.obofoundry.org/ontology/ro.html. An expanded set of upper-level relations, developed to deal with the problem documented in (Grewe 2016), is provided in part 2 of (ISO/IEC 21838 2021).

[20] The two main types of part in FMA are: constitutional parts, which are genetically determined, as in: Hand *has_constitutional_part* Skeleton of Hand; and regional parts, where the part entities are the results of fiat delineation using arbitrary landmarks, as in: Hand *has_regional_part* Digit. (Mejino et al. 2003)

is not an empirical assertion. Thus it is not falsified by the existence of human thumbs from which the nail has been removed. Rather, it is a statement that expresses how Nail of Thumb and Skin of Thumb are supposed to relate to each other in virtue of the workings of the underlying structural genes of the human organism.

4.3   Aristotelian definitions

A further crucial contribution of the FMA to the subsequent development of biomedical ontologies is in the field of definitions. The goal of a dictionary definition is to provide an explanation of the meaning of an expression that is useful to humans. In the ideal case, the dictionary provides an explanation that is built out of terms that are more familiar and simpler in meaning than the term to be defined. Often, however, dictionary definitions will amount to mere paraphrases, and they may be circular, either directly, or indirectly (as when term A is defined using term B, but term B is defined using term A). Often, too, multiple, mutually inconsistent definitions are provided for a single term.

To reach the goal of providing a tool to support logical reasoning, FMA requires a set of logically consistent definitions, with at most one definition for each term, and structured in such a way that each definition provides a statement of individually necessary and jointly sufficient conditions for the correct application of the term defined (Köhler et al. 2006, Seppälä et al. 2017). To address these needs, Rosse and his collaborators introduced the idea of what, drawing on the background discussed above in 1.2, they called 'Aristotelian definitions'.

A definition of the form

(2) S =def. a G which Ds,

where S stands for *species*, G for *genus*, and D for *differerentia(e)*, tells us that, if we know that something is a G which Ds, then we know that it is an S; and if we know that something is an S, then we also know that it is a G which Ds. Here G is the immediate parent of S in the backbone taxonomy of the salient ontology, and D is what sets those Gs which are Ss apart from the rest of the Gs. An example from the FMA ontology is:

Anatomical structure [S] =def. Material anatomical entity [G] which ⌜is generated by coordinated expression of the organism's own genes that guide its morphogenesis; has inherent 3D shape; is such that

its parts are connected and spatially related to one another in patterns determined by coordinated gene expression⌉[D] (Rosse and Mejino 2007)

where ⌈⌉ mark out the collection of sufficient conditions that forms the salient specific difference. Together, G and D specify the essential characteristics of any S. And a 'group of entities that share the same set of essential characteristics constitutes a class of the ontology' (Michael et al. 2001),

5   The Gene Ontology (GO)

5.1   Background

In 1977 Frederick Sanger and his collaborators sequenced the first full genome, that of a virus called phiX174. Since that point, the biological and biomedical sciences have been subjected to a process of upheaval as a result of the need to take account of the gigantic amounts of molecular assay data that have been generated in the wake of the successful completion of the human and the various fly, mouse, fish, yeast and other model organism genome projects. Practically all aspects of what we might call 'old biology' were destined to be transformed as biologists and clinical scientists worked out how to take account of these new data in dealing scientifically not only with the many new kinds of entity being disclosed at the molecular (and finer) levels through the advance of science, but also with all the already recognized phenomena at coarser levels of granularity (cell, tissue, organ, organism, population) upon which biology and medicine had hitherto been based. (Kumar et al. 2004)

But how to make the gigantic quantities of new data discoverable and usable by biologists in a situation where the primary source data lived in many independently developed biological databases? How to transform these many efforts into a single cooperative force? Among the very earliest repositories for the new data, created already in the 1970s, were the first protein structure database (Protein Data Bank, https://www.rcsb.org/pages/pdb50/) and the first mammalian genetics database (created at the Mouse Genome Informatics (MGI) resource of the Jackson Lab, http://www.informatics.jax.org/). These were followed in 1981 by the first depository for nucleotide sequences, established in 1981 at the European Molecular Biology Laboratory (EMBL) in Heidelberg (https://www.embl.org/about/history/). Each of these contributed to the strategy of using molecular assay data deriving from model organisms to advance our understanding of human health and disease, the idea being that clinical scientists could harvest the results of experiments carried out on model organisms in order to draw conclusions relevant to humans by exploiting cross-species sequence correlations. It is this idea which provided initial impetus for the

development of the GO.

By the turn of the millennium the number of biological databases was reaching a level where it had become unmanageable. Attempts to create a federated system failed, not least because it was so hard to get the many groups involved to agree on how the data should be structured and labelled. The fear, too, was that such a federated system would create what Suzanna Lewis refers to as 'a technological behemoth that would be unable to respond to new requirements when they inevitably occurred.'

> The most fundamental questions for the biologists served by the model organism databases revolved around the genes. … One essential aspect of this, which everyone agreed was necessary, was *systematically recording the molecular functions and biological roles of every gene.* (Lewis 2004, emphasis added)

5.2 The origins of the GO

In the 1990s, Michael Ashburner began assembling classifications of molecular functions and biological processes, originally to serve the requirements of FlyBase, the database for *Drosophila* genetics and molecular biology. At around this time different model organism communities began to see that they could solve a significant portion of their data integration issues if a functional classification system were created that was cross-species in nature. The goal was to get the developers of databases focused on sequence (nucleic acid or protein) together with the developers of other specialty biological databases built for different ranges of organisms to agree on how this should be done in a way that would work for all organism communities.

Lewis describes against this background how the GO came into being in 1998:

> In July of that year, Michael Ashburner presented a proposal at the Montreal International conference on Intelligent Systems for Molecular Biology (ISMB) bio-ontologies workshop to use a simple hierarchical controlled vocabulary; his proposal was dismissed by other participants as naïve. But later, in the hotel bar, representatives of FlyBase (me [Ashburner]), SGD [the *Saccharomyces* (yeast) Genome Database] (Steve Chervitz), and MGI (Judith Blake) embraced the proposal and agreed jointly to apply the same vocabulary to describe the molecular functions and biological roles for every gene in our respective databases. Thus we founded the Gene Ontology Consortium. (Lewis 2004; compare Ashburner 1998; Stevens 2013)

Note that the vision was not to create a *database* covering all functions of all genes in all organisms. Rather – and here lay the brilliant insight of Ashburner, Lewis and their collaborators – it was to create a controlled vocabulary for representing *types of molecular functions*, and to use this vocabulary to annotate (or 'tag') occurrences of references to corresponding genes or gene products in literature or in data in such

a way as to make the latter discoverable by third parties from different branches of biology.

The GO became, in effect, an engine for searching in literature and data what was still mostly hidden to outsider communities because it was inadequately or inconsistently described. It was based on annotations created by human beings (PhD biologists), pioneers in the new discipline of biocuration. The GO itself was to a large degree populated through the work of such biocurators. The annotations themselves would then be compiled, in conjunction with the UniProt protein sequence repository (UniProt 2008), to form the GO Annotation database (GOA) (Camon et al. 2004), Then came more sophisticated software tools such as the AmiGO browser (http://amigo.geneontology.org/), which allowed a significant fraction of the world's biological literature and data to be subjected to filtered search, allowing an investigator, for example studying the process of muscle development in *Bos taurus* (cow), to find immediately all proteins documented as involved in this process, all the articles in which this involvement is documented, and the source and nature of the evidence which each of these articles provides.[21]

The result was called 'Gene Ontology', not because it was an upshot of the work on ontologies growing out of the KR and other computer-associated disciplines in the preceding years, but merely because 'ontology' was, in 1998, the word-du-jour. The KR ontologies were in many cases, as we saw, products of a view to the effect that for every different project a new ontology is needed. The more ontologies, after all, the better. But then the results, for all their bangs and whistles, proved (not surprisingly) useless as soon as their authors moved on to the next project. The GO, in contrast, resulted from the insight that a simple controlled vocabulary could *unite* the many sequence-data driven projects springing forth on all sides. It started out, not as a sophisticated computer artifact, but rather as just a simple directed acyclical graph that could serve as the basis for an indefinitely extendable project of annotation. The nodes of the graph are terms[22] (again: nouns and noun phrases, albeit now associated with alphanumeric identifiers, definitions, URIs, and so forth), and its edges are relations (initially just *is_a* and *part_of*).

The fact that the GO was developed and maintained primarily by experts in molecular biology led initially to a certain animosity between the GO community and the community of those who had been developing ontologies on the basis of their computer expertise. However, with the eventual adoption by the GO

---

[21] http://amigo.geneontology.org/amigo/search/annotation?q=muscle%20development

[22] Stefan Schulz (personal communication) points out that 'label' is in some ways preferable to 'term' A text string such as 'Primary malignant neoplasm of lung (disorder)', for example, would never be used by any human author of scientific text. In the end, however, he favors over 'term' the expression 'representational unit', whose advantages are outlined in (Smith, Kusnierczyk and Ceusters 2006).

community of OWL as their ontology development language, and with the ever increasing numbers of powerful software tools and algorithms and research methodologies made possible by the existence of the GO and its sister ontologies, this animosity has now largely disappeared.

Initially, too, there were sceptics on the biology side, above all Sydney Brenner, winner of the 2002 Nobel Prize for his discoveries concerning programmed cell death. In the same year Brenner published a paper entitled "Life sentences: Ontology recapitulates philology", charging the GO Consortium with the desire to transform genomics into what he called 'genamics'.

To do serious theoretical work, Brenner held,

> the network we should be interested in is not the network of names but the network of the objects themselves. The language of these objects is not the Oxford Dictionary of Molecular Biology … but the language of molecular biology itself. [There the] objects have their own names: they are chemical names written in the language of DNA sequences and the arrangements of amino acids on protein surfaces. (Brenner 2001)

What Brenner failed to see was that, even should we all of us become fluent in the language of chemical names, we would still need to connect what we can say in this language with what we need to say in all the languages of old biology, including, not least, the languages of clinical medicine.

Since its inception, indeed, the GO has gone from strength to strength. It is today by far the world's most successful scientific ontology, whether measured along the dimensions of number and variety of associated software applications; quantities of data and literature annotated using its terms; number, size and degree of utilization of major databases incorporating term; numbers of experiments performed with its aid, and so forth.

There are multiple drivers of this success. One of the main ones is that the GO and the GO annotations are hand built by human curators, who use the scientific literature as a basis for their work. The result is an extract of biological knowledge captured using GO (and sister ontology) terms and relations, which has proved itself to be of tremendous utility.

There have, to be sure, been a number of proposals to leave population of the GO to machine learning. The problem with this approach is that it is not possible to create an algorithm that can extract knowledge from scientific literature automatically (Landgrebe and Smith 2021). Algorithms can be used for the sort of approximative text translation that is made available by google translate, but they cannot achieve results

with the sort of accuracy that is required for the scientific purposes of the GO (Landgrebe and Smith 2021).

One very fruitful application of GO is to what is called the enrichment analysis of gene (product) datasets. In intervention studies (for example genetic or pharmacological interventions) or time-series analyses, the GO can be used to obtain an overview of the cellular locations, functions and biological processes in which the gene products are involved in order to develop hypotheses about dependent variables or outcomes analysed in such experiments. The GO can also be used to classify and assess the status of independent variables in order to identify confounding effects (hidden co-variables). Powerful software applications have been developed for these purposes, including the GO-Figure! visualization tool developed by (Reijnders and Waterhouse 2021).[23]

Another reason for the GO's success is that it makes certain sorts of investigations possible that would just not be possible without it. The point is not just that genomic data is annotated with the same shared ontology, nor that this enables such data to be exchanged and integrated. Still more important is that the resulting huge and ever growing unified knowledge base about the functions of genes makes it possible to interpret large-scale measurements of gene expressions (or other -omics measurements) in relation to an unending series of biological phenomena. Examples of studies, selected at random from those published just in recent weeks, use the GO to identify pathways implicated in suicide behavior, breast cancer survival, autism spectrum disorders, involvement of calcium signaling in Schizophrenia, association signals of dental caries, disease modeling in *C. elegans*, and many more.

---

[23] At the same time, care must be taken to avoid misuse of the GO annotation data, for example by failing to take account of the ontological structure of the GO itself, or by ignoring the evidence codes which provide information as to the methods by which the data expressed in annotations were obtained (Rhee et al. 2008).
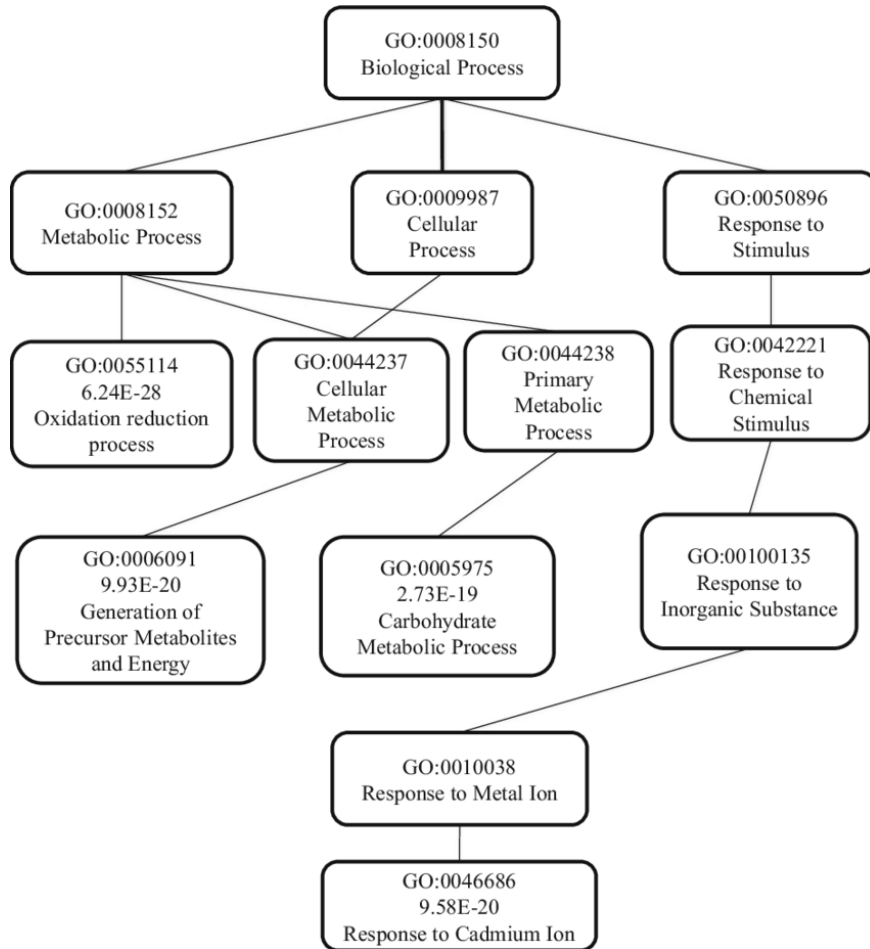
Figure 6: Fragment from the GO Biological Process Ontology, from (Li 2004)

5.3 The GO table of categories

The three questions you want the answers to when you discover a new gene product or complex are:

What does it do at the molecular level of granularity?

To what downstream biological processes does it contribute?

Where is it located in the cell?

The GO is accordingly divided into three sub-ontologies, whose respective root nodes – referred to in the original GO paper (Ashburner et al. 2000) as 'The three categories of GO' – are defined as follows:

Molecular function =def. Biochemical activity (including specific binding to ligands or structures) of a gene product. This definition also applies to the capability that a gene product (or gene product complex) carries as a potential. Examples: 'enzyme', 'transporter', 'ligand'.

Biological process =def. Biological objective to which the gene or gene product contributes. A process is accomplished via one or more ordered assemblies of molecular functions. Processes often involve a chemical or physical transformation, in the sense that something goes into a process and something different comes out of it. Examples: 'cell growth and maintenance', 'signal transduction'.

Cellular component =def. Place in the cell where a gene product is active. Examples: 'ribosome', 'nuclear membrane', 'Golgi apparatus'.

5.4 Function in the GO

The GO has retained its original modular architecture and its general structure and methodology over its more than 20 year history. But it has been subject throughout this entire period to considerable revisions at lower levels. This is primarily a matter of the deprecation of terms deemed obsolete, revision of definitions, or addition of new terms and even of new families of terms, for example covering hitherto underrepresented domains, such as immunology (Diehl et al. 2007).

The passage of time has seen also revisions to the original definitions of the three GO categories, and it is especially in connection with the GO's definition of 'function' that controversy has arisen. The current definition of *molecular function* reads as follows:

(3)   **Molecular function** =def. molecular process that can be carried out by the action of a single macromolecular machine, usually via direct physical interactions with other molecular entities. Function in this sense denotes an action, or activity, that a gene product (or a complex) performs. (http://geneontology.org/, as of August 7, 2021)

At the start, there was built into the GO world view the assumption that *is_a* relations can never span the boundaries between the three GO sub-ontologies. The *functions* at the (chemical) level of granularity of molecules in the GO thus stand in some way opposed both to the *processes* occurring at higher (biological) levels of granularity and to the *locations* in the cell.

Quite rightly, I believe, many of those who first encounter the GO are therefore confused by the fact that its molecular function ontology is populated primarily with terms designating types of activities – terms such as 'ion channel regulator *activity'*, 'regulation of lysozyme *activity'*, 'ceramide floppase *activity'*, 'regulation of phosphatidate phosphatase *activity'*, and so on.

The new definition of molecular function (3) tells us, indeed, that *molecular function is_a molecular process*. In normal usage, however, and also in the usage of many philosophers from Aristotle onwards (Lennox 2021), functions are not a special type of process. Rather, they are certain sorts of historically grounded potentials or capabilities in things that *can be realized* in processes when suitable circumstances obtain. As the term 'function' is normally understood, a function can fail entirely to be realized; or it can be *mis*realized. A well-oiled machine, for example, will indeed perform its function in normal circumstances; but when things go wrong then it can behave (act) in all sorts of non-functional – or, as we might also say, non-canonical – ways.

In a sense, the functions of the macromolecular machines inside an organism are being continuously realized, just as the function of the organism's heart is being continuously realized for so long as the organism is alive. But macromolecular machines change their activity patterns. For example, the sleeping brain is biochemically very active, but the pattern of activity differs from the wake pattern, and it differs again if one suppresses the normal rest pattern by taking sleeping pills or alcohol or both. Strictly speaking, of course, these latter cases are irrelevant to the GO. The GO, too, is a canonical ontology, in that it does not deal, for example, with biological processes which are induced experimentally. Unrealized functions at the molecular level are therefore out of scope. However, the cyclic nature of many activities in organisms means that the referent of 'activity' will even in many canonical cases differ from one phase to the next. The referent of 'function', in contrast, will always be the same.

Defining precisely the meaning of the term 'function' is a non-trivial matter, and philosophers and others have proposed various alternative definitions. In the *Gene Ontology Handbook,* Paul Thomas provides an account of the GO's usage of 'function' according to which it is the standard etiological or 'selected effect' definition of function that is intended by the GO. We believe that his arguments for this interpretation –

an interpretation which we also defend (Spear et al. 2016) – are sound. Unfortunately, as we shall see, he gets tangled in knots as a result of the odd terminological choices bequeathed to him by the GO.

In a simplified version of the selected effect account (based on Millikan (1989), which Thomas also cites), a function is defined as follows:

> (1) A **has function** F =def. A originated as a reproduction (for instance as offspring, or as copy) of some prior item or items that, due in part to possession of the properties reproduced, have actually performed F in the past, and A exists because (causally historically because) of this or these performances.

It is, very briefly, the function of my heart to pump blood *because* my ancestors' hearts' pumping blood through their bodies kept them alive and because I exist because of this. We note that, according to this definition, it would still be the function of my heart to pump blood even if (for example because I am connected to a heart-lung machine) it is currently unable to do so. It would still be the function of my screen to display pixels even if (for example because my machine is switched off) it is currently unable to do so.

As Thomas correctly points out, it is an advantage of the selected effect approach that it explicitly incorporates evolutionary considerations by requiring that the function of any biological entity ultimately derives from its history of natural selection. The approach thereby provides a method for determining which, among the myriad potential alternative effects the actions of a particular entity might have, are properly to be considered as the exercise of its function. One effect of my heart pumping, for example, is to produce sound; but this is not a part of the function of my heart, because this effect was not selected for.

There are several ways, now, in which the terminology of the GO might have been built in such a way as to do justice to the selected effect account of function. The simplest would be to rename the GO 'molecular function' ontology, and to call it instead the 'molecular activity ontology' or – following a practice which, I understand, is already being favored in certain GO circles – the 'molecular functioning ontology', thereby adding the explanation that by molecular 'activity'/'functioning' is meant 'the exercise of a function of a macromolecular machine', and providing as part of its glossary a suitable definition of 'function'.

Thomas, however, chooses to defend the position according to which the current GO labelling is in order as it stands, because, as he sees it, a function – in the selected effect sense – *is* an activity.

Yet in (1), and in all statements of the selected effect account of function with which we are familiar, the idea of *performing a function* is crucial to understanding what a function is. A function is not a process, that is to say, it is not a performance, or realization, or manifestation. Rather, a function is that which is performed, or realized, or manifested *in a process* – an idea which, at least in its reference to 'potential' or 'capability' in its original definition of 'molecular function' (see 5.3 above), the GO already gets right. Potentials or capabilities, too, are the sorts of things which get realized (or actualized) in processes. Functions are in fact a special subclass of capabilities: they are those capabilities which are selected for.

Thomas summarizes his account of molecular function in two places, as follows:

> (2) In the GO, a **molecular function** is a process that can be carried out by the action of a single macromolecular machine, via direct physical interactions with other molecular entities. Function in this sense denotes an action, or activity, that a gene product performs.

> (3) A function as conceived by molecular biologists (in what could be called the 'molecular biology paradigm') refers to specific, coordinated activities that have the appearance of having been designed for a purpose. That apparent purpose is their function.

As to (2), a process is not something that can be 'carried out by an action'; for as we have seen, the process *is* the action. And as to (3): 'function is an activity' and 'function is an apparent purpose' – you can choose one of these; but you can't choose both.

5.5   Extending the GO

Interestingly, in elucidating his account of function, Thomas draws on Jacques Monod's idea in *Chance and Necessity* (1971) of *teleonomy*. This Monod defines as 'the characteristic of objects endowed with a purpose or project, which at the same time they exhibit through their structure and carry out through their performances' (p. 9). This applies to artifacts such as screwdrivers, which are *designed* to have a certain purpose. For living systems, however, we cannot talk of design. Rather, as Thomas writes, 'what appears to be a future-goal-oriented action by a living organism is, in fact, only a blind repetition of a genetic program that evolved in the past.' More completely, however – for there are two series of blind repetitions here – he should write:

1. the *program* is copied over and over again through (blind) biological processes of copying the relevant entity (for instance the relevant macromolecular machine)

2. the *execution* of each copy of the program is repeated in the successive realizations of that entity's function.

Teleonomy, for Monod, is present at all levels of a biological system, from proteins (which he calls 'the essential molecular agents of teleonomic performance') to 'systems providing large scale coordination of the organism's performances … [such as] the endocrine and nervous systems' (op. cit., p. 62). At all levels, indeed, we have objects, and systems and parts of objects, performing (activities) which realize apparent purposes (functions), such as pumping blood, regulating chemical levels in the blood, removing damaged cells from the blood, and so forth. And in each of these cases we have to deal, not only with functions of the body involving groups of cells interacting via molecules or ions, but also with functions of parts of the body at higher levels of granularity than molecules, which we can therefore call *biological* functions. And interestingly, although Thomas's (2019) paper deals almost exclusively with functions at the molecular level, its title is "The Gene Ontology and the Meaning of Biological Function".

Table 1 depicts, on this basis, an extended version of the GO architecture, in which an explicit division is drawn between levels of granularity along the vertical axis and kinds of entity along the horizontal, where the shaded cells correspond (roughly) to the coverage domain of the original GO.

		Ontological Categories		
		Function	**Process**	**Object**
Level of Granularity	**Cell, organ, organism**	to pump blood	pumping blood	heart
	Molecule	to catalyze a biochemical reaction	acceleration of a chemical reaction	ribozome

Table 1: Extended GO ontology to accommodate biological functions. Adapted from "STOP!" (see footnote 26)

6   The Open Biological and Biomedical Ontologies (OBO) Foundry

6.1   The birth of OBO

As we learn from the subtitle of the landmark paper (Ashburner et al. 2000), the GO was originally conceived as a 'tool for the unification of biology'. In other words, the GO was built to foster the melding together of biological data sets and techniques across disciplines, across levels of granularity in the organism, across species, and across geographically dispersed communities of originators and compilers of data and of researchers using these data. An example of success in this regard is the way in which the GO enables communication across all the disciplines collaborating for example in a field such as aging research, which involves the study of model systems of human aging in organisms as diverse as yeast, reptiles and whales.

Already in 2001, the trail laid by the GO opened the way for the creation of a series of controlled, cross-species vocabularies for neighbouring areas of biology by proposing the creation of a public ontology repository, originally (we imagine for a very short time) referred to as 'GOBO', for 'Global Open Biomedical Ontologies' (Ashburner and Lewis 2003), and subsequently dubbed the 'OBO Library'.

The rules for building ontologies for this library can be found in a tutorial presented by Ashburner and Lewis at the Intelligent Systems for Molecular Biology (ISMB) conference in 2005 on "Principles of Biomedical Ontology Construction" (http://bit.ly/2GUkpoh). Most important are that the ontology must be shared without limit, and thus that it must be in the public domain and easily findable;[24] that it is used in application to actual instances of important scientific data; and that it is maintained in such a way that, where such application leads to identification of errors and gaps, the latter will be promptly rectified.

In the case of the GO, this strategy produced a positive snowball effect, making the GO increasingly attractive to successive cohorts of new users, who themselves identify new errors and gaps, giving rise to a regimen of continual improvement of a sort that was unknown to ontologies before the GO.

6.2   The birth of the OBO Foundry

The OBO *Foundry* was first conceived at a meeting held in Leipzig in 2004 on the topic of *The Formal Architecture of the GO.* Other groups from the KR and OWL communities had attempted earlier to interest

---

[24] The OBO community here anticipates the modern FAIR approach (Wilkinson et al. 2016)

the GO community in the benefits of a more ambitious approach to ontology development, especially as concerns the treatment of logic and definitions. Where these earlier efforts had failed,[25] some of the new arguments presented at this meeting drawing on the perspective of ontological realism met with greater success.[26]

No less important was the introduction of new rules to promote coordinated ontology development, the idea being that ontologies would be admitted as members of the Foundry initiative only if their developers had committed in advance to certain principles, for example relating to working within set boundaries (for example of *proteins* or *cell types*) and agreeing to collaborate on those terms which relate to entities in areas where boundaries overlap.

The details of Foundry organization were then worked out in a series of meetings, some of them under the auspices of the then newly established National Center for Biomedical Ontology (NCBO).[27] Building the Foundry was viewed as amounting to distinguishing within the original OBO Library as a whole an inner compartment comprising, at any given stage, those ontologies certified to have satisfied both the Library principles and also a series of additional principles specific to the Foundry, designed to advance the quality and interoperability of its included ontologies (Smith, Ashburner, et al. 2007).

The core goal of the OBO Foundry – where 'OBO' is now understand as meaning 'open biological and biomedical ontologies' – like that of the OBO Library, is to create a situation where ontologies would support efficient knowledge accumulation in the life sciences by providing recommended sets of terms for annotating data in each life science domain – thus one set of terms for proteins (Chen et al. 2020), one set of terms for small molecules (or chemical entities of biological interest: Hastings 2012), one for plants (Cooper et al. 2018), and so forth. The terms in each ontology would be accompanied not only by natural language definitions designed to ensure that the terms are correctly used by those (humans) involved in creating annotations of biological literature and data, but also by *formal definitions* designed to promote computer-aided reasoning with the resulting annotated data.

For the *Foundry* ontologies a layer of governance was introduced – in some ways analogous to the editorial board of a scientific journal – and a process of review was established which would certify conformance

---

[25] One member of the OWL community remarked to me at the time that 'a meeting on the formal architecture of the GO? Well … that would have to be a very short meeting, then'.

[26] These included arguments in a presentation by the author entitled "STOP!" (for: Smart Terminologies through Ontological Principles – http://ontology.buffalo.edu/04/STOP_GO_5_04.ppt), illustrating how the realist perspective can help in the identification of errors in the GO. See also (Smith, Köhler and Kumar 2004).

[27] http://ncorwiki.buffalo.edu/index.php/NCBO_Sponsored_Dissemination_and_Training_Events_2005-2015

to the Foundry principles. The current set of principles includes the requirement to use a standard ontology language (currently OWL) and use of Basic Formal Ontology (see below) as shared top level. In fact BFO makes itself manifest already in the terminology used in the top two rows of Table 2, which depicts the initial structure proposed for the OBO Foundry, a structure which in effect extends Table 1 to include also ontologies external to the GO.

RELATION TO TIME GRANULARITY	CONTINUANT				OCCURRENT
	INDEPENDENT		DEPENDENT		
ORGAN AND ORGANISM	Organism (NCBI Taxonomy)	Anatomical Entity (FMA, CARO)	Organ Function	Phenotypic Quality (PaTO)	**Biological Process** (GO)
CELL AND CELLULAR COMPONENT	Cell (CL)	**Cellular Component** (FMA, GO)	Cellular Function (GO)		
MOLECULE	Molecule (ChEBI, SO, RNAO, PRO)		**Molecular Function** (GO)		Molecular Process (GO)

Table 2: The projected structure of the OBO Foundry from around 2005 (shaded regions correspond to the 3 original GO ontologies).

In a parallel development, there arose at about the same time what would become a much larger biomedical ontology repository extending the original OBO Library idea, namely the NCBO BioPortal (https://bioportal.bioontology.org/), which provided the advantage of providing access to ontology-structured versions of SNOMED CT, HL7, MeSH and other major resources from the world of medical terminology.

The BioPortal adopted a very liberal strategy of acceptance of ontologies, which was in a sense at the opposite extreme from the strategy of the OBO Foundry.[28] This, however, created for the BioPortal a

---

[28] Later there arose the Ontobee portal (http://www.ontobee.org/) (Ong 2017), a biomedical ontology repository that is optimized for the purposes of the OBO Foundry using the technology of a linked data server. Ontobee is a linked ontology data server, and supporting ontology term dereferencing, linkage, query, analysis and integration.

problem of redundancy and lack of consistency between the (now of the order of) 500 ontologies listed, a problem which was further exacerbated as new ontologies were developed incorporating reuse of terms and definitions from multiple already established ontologies but supplying them with new term identifiers and new URIs. This conflicts with the OBO Foundry goal of creating a set of mutually consistent and non-redundant ontologies for the life sciences that would promote for each term a unique recommended natural language definition, formal definition, and URI.[29]

6.3   Basic Formal Ontology (BFO)

BFO was adopted as required top level for ontologies in the Foundry in order to make available a common set of categories (highest-level universals) that would serve as the shared starting point for the definitions of lower-level universals included in the coverage domains of the separate biomedical ontologies in the Foundry.

BFO itself was developed as a very small representational artifact with the narrowly focused task of providing an upper ontology which could be used to support the integration of domain ontologies developed for purposes of scientific research. As Tables 1 and 2 make clear, the structure of the set of ontology modules of the OBO Foundry is derived, in effect, by taking the cross-product of BFO's top-level categories with the multiple granular levels (of molecule, cell, organ, organism, population) relevant to biology. The FMA was the first extensively populated ontology to take advantage of the theoretical foundations of a top-level ontology and thereby extend the latter into the biomedical domain (Rosse, Kumar et al. 2005, Rosse and Mejino 2007).

Adopting BFO allowed

(i)   the explicit formulation of aspects of the development methodology and architectural structure of the OBO Foundry (and other) ontologies in ways that have helped steer their subsequent development (Arp et al. 2015),

(ii)   providing a readily applicable technique for formulating definitions of terms in these ontologies (Seppälä 2017),

(iii)   formalizing relations (Smith, Ceusters, et al. 2005),

---

[29] The most recent versions of the BioPortal go some way to solving this problem by generating search results in such a way that the original source definition of a term is returned first in the list of all the ontologies which use that term with that definition.

(iv)   supporting the strategy of cross-product definitions (Mungall 2011), whereby definitions in one OBO Foundry ontology will draw on terms defined already in other such ontologies, for example as GO terms whose definitions represent molecules draw on ChEBI definitions of those molecules, as described for example in (Hill et al. 2013),

(v)   formally encoding the OBO Foundry principles as operational rules and applying the resultant checks across the full OBO suite of ontologies, thereby demonstrating how a sizable federated community can be organized and evaluated on objective criteria that help improve overall quality, interoperability and sustainability (Jackson 2021).

6.4   The evolution of BFO

There have been four releases of BFO thus far.[30] Version 1 was released in 2001, and the influence of Aristotle's table of categories on this first version can be seen in the similarity of terminology and structure as between the upper rows of Figure 2 and those of Figure 7. Another influence was the top-level ontology DOLCE (Masolo et al. 2004). BFO shared with DOLCE from the very start a quadripartite architecture based on two orthogonal divisions of entities into disjoint categories of 1. continuant vs. occurrent and 2. independent vs. dependent. Material objects (Aristotle's substances) are independent continuants; qualities are dependent continuants; and processes are occurrents. Entities in all of these categories exist on the level of both universals and instances (as in Table 1).[31]

The release of BFO 1.1 in 2007 was prompted by the need to enable coverage of information artifacts, nucleotide sequences, and similar (copyable) entities, a need which arose with the birth of two new ontologies, the Ontology for Biomedical Investigations (OBI) in 2006 (Bandrowski et al. 2016, Vita et al. 2021)[32], and the Information Artifact Ontology (IAO; see Ceusters and Smith 2015), which provided terms used to represent entities such as publications, footnotes, protocols, databases, and so forth.

The release of BFO 2 in 2015 reflects the transition from an OWL DL to an OWL 2 formalization, as well as the addition of term IDs and of temporalized relations. A preliminary version was released for review at the 2012 meeting of the International Conference on Biomedical Ontology.

[30]Successive versions are available through http://ontology.buffalo.edu/bfo.

[31] A further influence was lessons learned from work on a framework that would link the quantitative data studied in the new field of Geographic Information Science with qualitative data pertaining to the hills and valleys, rivers and lakes that form the subject-matter of what we might call old geography (Mark and Smith 2004).

[32] OBI is now an OBO Foundry ontology. It was created as a generalization of the Functional Genomics Investigation Ontology (FuGO) (Whetzel 2006).

By the year 2020, BFO has come to serve as something of a stable attractor to ontology developers (http://basic-formal-ontology.org/users.html, Haller 2020), thereby giving rise to powerful network effects analogous to those brought by the QWERTY keyboard and the TCP/IP internet protocol, whereby each successive new user BFO raises the value of the artifacts created on its basis by earlier users, in another positive feedback loop.

As a consequence of these developments, the Joint Technical Committee on Information Technology (JTC 1) of the International Standards Organization (ISO) and the International Electrotechnical Commission (IEC) have approved in 2021 the ISO/IEC 21838: Top-Level Ontologies (TLO) standard. Part 1 of this standard sets forth the requirements for being a top-level ontology. Part 2 documents BFO in a way that demonstrates satisfaction of these requirements.

The release of BFO-2020 includes a more careful treatment of definitions. All non-primitive terms have been provided with English language definitions (which means: statements of individually necessary and jointly sufficient conditions). All primitive terms have been provided with elucidations, which means: statements of necessary conditions together with specifications of examples of use. Additional improvements concern the logical formalization of BFO. Along with an OWL version of BFO-2020, the ISO standard provides also an axiomatization in common logic (BFO-2020-CL) and a translation thereof into FOL. A proof of consistency of BFO-2020-CL is provided, together with a proof that BFO-2020-OWL is derivable therefrom. English-language definitions and elucidations provided in the standard are formulated in such a way as to be as close as possible to BFO-2020-CL while at the same time serving as an access route to the content of BFO-2020 for human users.
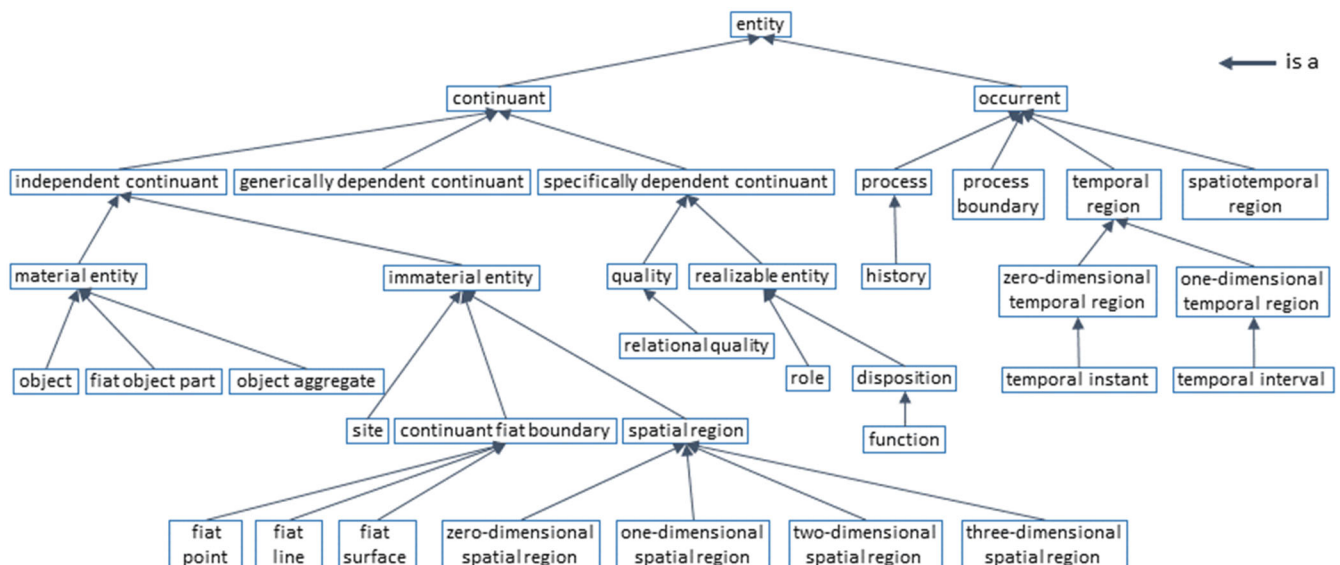
7   Conclusion

Like the Gene Ontology, and like the Planteome ontologies, which are seen by their developers as 'integrative tools for plant science' (Walls et al. 2012), the OBO Foundry as a whole is a *tool for the unification of biology*. Indeed, all the OBO Foundry ontologies continue in their way the project of the Vienna circle to achieve the unification of science. They do this, however, not from the starting point of logic and philosophy, but rather from the starting point of biology and ontology. And they do this more successfully, because their project of unification is deeply interwoven, through multiple different sorts of multidirectional interactions, with ongoing developments in biology and, increasingly, in the clinical sciences.

References

Arp R, Smith B, Spear A. (2015) *Building ontologies with Basic Formal Ontology*, MIT Press, Cambridge, MA.

> To the Copy-Editor: Please note that names of ontologies are like names of countries (United States of America); they must always be capitalized, as marked in *yellow* below. Also SNOMED CT is an acronym, and must be capitalized thus.

Ashburner M (1998) On the representation of 'gene function' in databases. *Discussion paper for ISMB, Montreal, 1998.* Version 1.2 June 19 1998.
http://biomirror.aarnet.edu.au/biomirror/geneontology/docs/gene_ontology_discussion.html

Ashburner M and Lewis SE (2003) On ontologies for biologists: The Gene Ontology – untangling the web. In: Bock GR, Goode JA, eds. *'In Silico' simulation of biological processes*. John Wiley & Sons, New York.

Ashburner M and Lewis SE (2005) Principles of biomedical ontology construction, Tutorial, Intelligent Systems for Molecular Biology (ISMB), Detroit, MI, http://bit.ly/2GUkpoh/.

Ashburner M, Ball CA, Blake JA, et al (2000) Gene Ontology: tool for the unification of biology. *Nature Genetics*. 25(1), 25-29.B

Baader F, Horrocks I, Sattler U (2008) Description logics. *Foundations of Artificial Intelligence*. 3:135-

79.

Bandrowski A, Brinkman R, Brochhausen M, Brush MH, Bug B, et al. (2016) The Ontology for Biomedical Investigations. *PLoS One*. 11(4):e0154556.

Barnes J, ed. (2006) *Porphyry introduction*. Oxford University Press, Oxford.

Bauer S, Grossmann S, Vingron M, Robinson PN (2008) Ontologizer 2.0—a multifunctional tool for GO term enrichment analysis and data exploration. *Bioinformatics*. 24(14):1650-1.

BIPM (2019) International Standard System of Units, 9th edition, Sèvres, France.

Bodenreider O, Smith B, Burgun A (2004) The ontology-epistemology divide: a case study in medical terminology. *Formal Ontology in Information Systems. Proceedings of the Third International Conference (FOIS 2004),* pp. 185-195.

Bodenreider O, Smith B, Kumar A, Burgun A (2007) Investigating subsumption in DL-based terminologies: A case study in SNOMED CT. In *Artificial Intelligence in Medicine*, 39, 183-195.

Bona JP, Ceusters W (2018) Mismatches between major subhierarchies and semantic tags in SNOMED CT. *Journal of Biomedical Informatics*, 81:1-15.

Botti Benevides A, Bourguet JR, Guizzardi G, Penaloza R, Almeida JP (2019) Representing a reference foundational ontology of events in SROIQ. *Applied Ontology*. 14(3):293-334.

Brenner S (2002) Life sentences: ontology recapitulates philology, *Genome Biology* 3(4): 1006.1-2. doi: 10.1186/gb-2002-3-4-comment1006

Brunczwik A (1748) *Tractatus in Aristotelis logicam*, https://classic.europeana.eu/portal/en/record/2048128/39246, last accessed July 26, 2021.

Camon E, Magrane M, Barrell D (2004) The Gene Ontology Annotation (GOA) Database: sharing knowledge in Uniprot with Gene Ontology. *Nucleic Acids Research*. 1; 32 (Database issue): D262–D266.

Carey S (2009). *The origin of concepts*. Oxford University Press, Oxford.

Carnap R (1928) *Der logische Aufbau der Welt*, Berlin: Weltkreis. Translated RA George as *The Logical Structure of the World*, University of California Press, Berkeley, CA, 1967.

Casati R, Varzi AC (1994) *Holes and other superficialities*. MIT Press, Cambridge, MA.

Ceusters W (2011) SNOMED CT's RF2: Is the future bright? *Studies in health technology and informatics*, 169, 829–833.

Ceusters W (2021) The place of Referent Tracking in biomedical informatics. Terminology, Ontology and their Implementations', Springer Book 87020228, Springer Nature Switzerland AG (this volume), pp. XX-XX.

Ceusters W, Elkin P, Smith B (2007) Negative findings in electronic health records and biomedical ontologies: a realist approach. *International journal of medical informatics*. 76:S326-33.

Ceusters W, Mullin S (2019) Expanding evolutionary terminology auditing with historic formal and linguistic intensions: a case study in SNOMED CT. *Studies in Health Technology and Information* 264:65-9.

Ceusters W, Smith B (2003) Ontology and medical terminology: why descriptions logics are not enough. *Proceedings of the Conference Towards an Electronic Patient Record* (TEPR 2003), San Antonio, 10-14

May 2003 (electronic publication).

Ceusters W, Smith B (2015) Aboutness: towards foundations for the Information Artifact Ontology. *Proceedings of the Sixth International Conference on Biomedical Ontology* (ICBO). CEUR 1515. pp. 1-5.

Ceusters W, Smith B, Kumar A, Dhaen C (2004) Mistakes in medical ontologies: where do they come from and how can they be detected? in Pisanelli DM (ed.) *Ontologies in Medicine. Proceedings of the Workshop on Medical Ontologies, Rome October 2003*, Amsterdam: IOS Press, Studies in Health Technology and Informatics, vol 102, 145–64.

Ceusters W, Smith B, Kumar A, Dhaen C (2004a) Ontology-based error detection in SNOMED-CT®. MEDINFO, IOS Press, Amsterdam, pp. 482-486.

Chen C, et al. (2020) Protein ontology on the semantic web for knowledge discovery. Scientific Data 7, 337. https://doi.org/10.1038/s41597-020-00679-9

Cheyer A, Gruber TR (2010) Siri: A virtual personal assistant for iphone, an ontology-driven application for the masses. Presentation at Open, International, Virtual Community of Practice on Ontology, Ontological Engineering and Semantic Technology.

Cimino JJ (1998) Desiderata for controlled medical vocabularies in the twenty-first century. Methods of Information in Medicine. 37(4-5):394-403.

Cooper L, et al. (2018) The Planteome database: an integrated resource for reference ontologies, plant genomics and phenomics. *Nucleic acids research*. 46(D1):D1168-80.

Dahms HJ (2018) Mission accomplished? unified science and logical empiricism at the 1935 Paris Congress and afterwards. *Philosophia Scientiæ. Travaux d'histoire et de philosophie des sciences*. 22-23, 289-305.

Diehl 2007 AD, Lee JA, Scheuermann RH, Blake JA (2007) Ontology development for biological systems: immunology. *Bioinformatics*. 1;23(7):913-5.

Egdahl A (1907) Linnaeus' *Genera Morborum*, and some of his other medical works. *Medical Library and Historical Journal*. 5 (3): 185–193.

Feyerabend P (1978) In defence of Aristotle: Comments on the condition of content increase, in G Radnitzky and G Andersson (eds), *Progress and Rationality*, Reidel, Dordrecht, pp. 143-180.

Grene M (1963) *A portrait of Aristotle,* Faber and Faber, London.

Grewe N, Jansen L and Smith B (2016) Permanent generic relatedness and silent change, *Formal Ontology in Information Systems. Proceedings of the Ninth International Conference* (FOIS 2016) *Ontology Competition*, (CEUR 1660), 1-5.

Gruber TR (1993) A translation approach to portable ontology Specifications, *Knowledge Acquisition*, 5, 199-220.

Guarino N (1995) Formal ontology, conceptual analysis and knowledge representation. *International journal of human-computer studies*. 43(5-6):625-40.

Guarino N and Poli R, eds. (1995) Proceedings of the International Workshop on Formal Ontology in Conceptual Analysis and Knowledge Representation, special issue of *International journal of human-computer studies* 43(5-6):623-965.

Guarino N, Oberle D, Staab S (2009) What is an ontology? In: *Handbook on ontologies.* Springer, Berlin, Heidelberg, pp. 1-17.

Guarino N (2017) BFO and DOLCE: So far, so close…. *Cosmos + Taxis*. 4(4); 10-18.

Gurcan MN, Tomaszewski J, Overton JA, Doyle S, Ruttenberg A, Smith B (2017) Developing the Quantitative Histopathology Image Ontology (QHIO): a case study using the hot spot detection problem. *Journal of Biomedical Informatics*. 66:129-35.

Haendel MA, Neuhaus F, Osumi-Sutherland D, Mabee PM, Mejino JL, Mungall CJ, Smith B (2008) CARO – the Common Anatomy Reference Ontology. In A Burger et al. (eds) *Anatomy ontologies for bioinformatics: Principles and practice*. Springer, London, pp. 327-349).

Haller A, Polleres A (2020) Are we better off with just one ontology on the Web? *Semantic Web*. 11(1):87-99.

Hastings J, et al. (2012) The ChEBI reference database and ontology for biologically relevant chemistry: enhancements for 2013. *Nucleic acids research*. 41(D1):D456-63.

Haugeland J (1985) *Artificial Intelligence, the very Idea*, MIT Press, Cambridge, MA.

Hayes PJ (1978) Naive physics I: Ontology for liquids, Working Papers of the Dalle Molle Institute, No. 35, Geneva, 66pp.

Hayes PJ (1979) The Naive Physics Manifesto, in D Michie (ed) *Expert Systems in the MicroElectronic Age*, University of Edinburgh Press, 1979.

Hayes PJ (1985) The Second Naive Physics Manifesto, in Hobbs and Moore (eds.), 1-36.

Hayes PJ (1985s) 1985a. Naïve physics I: ontology for liquids, in Hobbs and Moore (eds.), 71-108.

Hayes PJ (2013) Early use of the word 'ontology' in AI (via John Sowa), http://ontolog.cim3.net/forum/ontolog-forum/2013-11/msg00016.html, last accessed July 4, 2021.

Hill DP, Adams N, Bada M, Batchelor C, Berardini TZ, Dietze H, Drabkin HJ, Ennis M, Foulger RE, Harris MA, Hastings J (2013) Dovetailing biology and chemistry: integrating the Gene Ontology with the ChEBI chemical ontology. *BMC genomics*. 14(1):1-0.

Hobbs JR and Moore RC, eds. (1985) *Formal theories of the commonsense world* (Ablex Series in Artificial Intelligence), Intellect Books.

ISO/IEC 21838 (2021) Information Technology—Top-Level Ontology (TLO), Part 1: Requirements, Part 2: Basic Formal Ontology, International Standards Organization, Geneva.

ISO/IEC 24707 (2018) Information Technology—Common Logic (CL): A Framework for a Family of Logic-Based Languages, International Standards Organization, Geneva.

Jackson RC, Matentzoglu N, Overton JA, Vita R, Balhoff JP, Buttigieg PL, Carbon S, Courtot M, Diehl AD, Dooley D, Duncan W, et al. (2021) OBO Foundry in 2021: operationalizing open data principles to evaluate ontologies. *bioRxiv*. https://doi.org/10.1101/2021.06.01.446587

Jansen L (2007) Aristotle's categories, *Topoi*, 26, 153-158.

Johansson I (2021) Quantities as metrical coordinative definitions and as counts: On some definitional structures in the new SI brochure. *Journal for General Philosophy of Science*, 1-23.

Köhler J, Munn K, Rüegg A, Skusa A, Smith B (2006) Quality control for terms and definitions in

ontologies and taxonomies. *BMC Bioinformatics*, 7(1), 1-12.

Kumar A, Smith B, Novotny DD (2004) Biomedical informatics and granularity. *Comparative and functional genomics*. 5(6-7):501-8.

Landgrebe J and Smith B (2021) Making AI meaningful again. *Synthese*. 198(3):2061-81.

Landgrebe J and Smith B (in preparation) BFO Physics Ontology – Mechanics Module, draft manuscript.

Leitgeb H and Carus A (2021) Rudolf Carnap, Supplement A: *Aufbau*, *The Stanford Encyclopedia of Philosophy* (Summer 2021 Edition), (ed.) EN Zalta https://plato.stanford.edu/archives/sum2021/entries/carnap/aufbau.html/.

Lennox JG (1984) Marjorie Grene, Aristotle's philosophy of science and Aristotle's biology, *Proceedings of the Biennial Meeting of the Philosophy of Science Association*, 2: 365-377.

Lennox JG (2021) Aristotle's biology, *The Stanford Encyclopedia of Philosophy* (Fall 2021 Edition), EN Zalta (ed.).

Leroi MA (2014) *The lagoon: how Aristotle invented science*. Penguin Books, London.

Lewis SE (2004) Gene Ontology: looking backwards and forwards. *Genome Biology* 6:103.

Li X et al (2014). Pop's Pipes: poplar gene expression data analysis pipelines. *Tree genetics & genomes*. 10(4):1093-101.

Linnaeus C (1758) *Systema naturæ per regna tria naturæ, secundum classes, ordines, genera, species, cum characteribus, differentiis, synonymis, locis*, Laurentii Salvii, Stockholm.

Linnaeus, C (1759) *Genera morborum*, Steinert, Upsalla.

McCarthy J (1980) Circumscription – A form of non-monotonic reasoning, *Artificial Intelligence*, 5: 13, 27-39.

McCarthy J (2000) Concepts of Logical AI, *Logic-based Artificial Intelligence*, Springer, pp. 37-56.

Marcus G and Davis E (2019) *Rebooting AI: Building Artificial Intelligence we can trust*. Vintage, New York.

Mark DM and Smith B (2004) A science of topography: Bridging the qualitative-quantitative divide. In: Geographic Information Science and Mountain Geomorphology. Springer-Praxis, Chichester, England, pp. 75-100 (2004)

Masolo C, Borgo S, Gangemi A, Guarino N, Oltramari A (2004) WonderWeb Deliverable D18: Ontology Library (http://wonderweb.semanticweb.org/deliverables/documents/D18.pdf).

Mejino Jr JV, Agoncillo AV, Rickard KL, Rosse C (2003) Representing complexity in part-whole relationships within the foundational model of anatomy. In *AMIA Annual Symposium Proceedings*, p. 450-454.

Michael J, Mejino Jr JV, and Rosse C (2001). The role of definitions in biomedical concept representation. *AMIA Annual Symposium Proceedings*, pp. 463-467.

Millikan RG (1989) In defense of proper functions. *Philosophy of Science* 56:288–302.

Monod J (1971) *Chance and necessity*. Alfred Knopf, New York

Moore GH (1988) The emergence of first-order logic. In P Kitcher and W Asprey (eds) History and

philosophy of modern mathematics 11:95-135. University of Minnesota Press, Minneapolis.

Mungall CJ, Bada M, Berardini TZ, Deegan J, Ireland A, Harris MA, Hill DP, Lomax J (2011) Cross-product extensions of the Gene Ontology. *Journal of biomedical informatics*. 44(1):80-6.

Mungall CJ, Torniai C, Gkoutos GV, Lewis SE, Haendel MA (2012) Uberon, an integrative multi-species anatomy ontology. *Genome biology*. 13(1):1-20.

Munsche H and Whitaker HA (2012) Eighteenth century classification of mental illness: Linnaeus, de Sauvages, Vogel, and Cullen. *Cognitive and behavioral neurology*, 25(4), 224-239.

Neuhaus F and Smith B (2008) Modelling principles and methodologies: relations in anatomical ontologies. In A Burger et al. (eds) *Anatomy ontologies for bioinformatics: Principles and practice*. Springer, New York, pp. 289-306.

Neurath O, Carnap R, Morris C (1938-1968) *Foundations of the unity of science*: *toward an International Encyclopedia of Unified Science*, 2 volumes, University of Chicago Press, Chicago.

Noy N, McGuinness DL (2001) Ontology development 101. Knowledge Systems Laboratory, Stanford University.

Ong E, Xiang Z, Zhao B, Liu Y, Lin Y, Zheng J, Mungall C, Courtot M, Ruttenberg A, He Y (2017) Ontobee: a linked ontology data server to support ontology term dereferencing, linkage, query and integration. *Nucleic acids research*. 45(D1): D347-52.

Quine WVO (1948). On what there is. *Review of Metaphysics* 2 (5), 21-38.

Reijnders MJMF and Waterhouse RM (2021) Summary visualizations of Gene Ontology terms with GO-Figure! *Frontiers in Bioinformatics*, https://doi.org/10.3389/fbinf.2021.638255

Rhee SY, Wood V, Dolinski K, Draghici S. (2008) Use and misuse of the Gene Ontology annotations. *Nature Reviews Genetics*. 9(7):509-15.

Rosse C (2001) Terminologia Anatomica; considered from the perspective of next-generation knowledge sources, *Clinical Anatomy* 14(2):120-33.

Rosse C, Mejino Jr JV (2003) A reference ontology for bioinformatics: The Foundational Model of Anatomy, *Journal of Biomedical Informatics*, 36:478–500.

Rosse C and Mejino Jr JV (2007) The Foundational Model of Anatomy Ontology. In: A Burger et al. (eds), *Anatomy ontologies for bioinformatics: Principles and practice*, Springer, New York, pp. 59-117.

Rosse C, Kumar A, et al (2005) A strategy for improving and integrating biomedical ontologies. In *AMIA Annual Symposium Proceedings*, pp. 639-643.

Rudnicki R (2019) *An Overview of the Common Core Ontologies*, CUBRC, Buffalo.

Sallam HN (2010) Aristotle, godfather of evidence-based medicine, *Facts, Views and Visions*, 2(1); 11-19.

Schulz S, Balkanyi L, Cornet R, Bodenreider O (2013) From concept representations to ontologies: a paradigm shift in health informatics? *Healthcare informatics research*. 19(4):235-42.

Schulz S, Martínez-Costa C (2015) Harmonizing SNOMED CT with BioTopLite: an exercise in principled ontology alignment. In *MEDINFO 2015: eHealth-enabled Health* 2, IOS Press, Amsterdam, pp. 832-836.

Schulz S, Steffel J, Polster P, Palchuk M, Daumke P (2019) Aligning an Administrative Procedure Coding System with SNOMED CT. In JOWO Joint Ontologies Workshops, 2019 (CEUR 2519).

Schulz S, Stenzhorn H, Boeker M, Smith B (2009) Strengths and limitations of formal ontologies in the biomedical domain. *Electronic Journal of Communication, Information and Innovation in Health* (Special Issue on Ontologies, Semantic Web and Health), 3 (1), 31-45. doi:10.3395/reciis.v3i1.241en

Seppälä S, Ruttenberg A, Smith B (2017) Guidelines for writing definitions in ontologies. *Ciência da informação* 46 (1), 73-88.

Shaw M, Detwiler LT, Brinkley JF, Suciu D (2008) Generating application ontologies from reference ontologies. Proceedings of AMIA Annual Symposium, pp. 672-6.

Smith B (2005) Against fantology, in JC Marek & ME Reicher (eds), *Experience and Analysis*. Vienna: HPT&ÖBV. pp. 153-170.

Smith B (2006) From concepts to clinical reality: An essay on the benchmarking of biomedical terminologies. *Journal of Biomedical Informatics*, 39 (3):288-298.

Smith B (2008) Ontology (science), in C Eschenbach and M Grüninger (eds), *Formal Ontology in Information Systems. Proceedings of the Fifth International Conference (FOIS 2008)*, Amsterdam: IOS Press, pp. 21-35.

Smith B, Ashburner M, Rosse C, Bard J, Bug W, Ceusters W, Goldberg LJ, Eilbeck K, Ireland A, Mungall CJ, Leontis N et al. (2007) The OBO Foundry: coordinated evolution of ontologies to support biomedical data integration. *Nature biotechnology*. 25(11):1251-5.

Smith B, Ceusters W (2010) Ontological realism as a methodology for coordinated evolution of scientific ontologies, *Applied Ontology*, 5, 139–188.

Smith B, Ceusters W, Klagges B, Köhler J, Kumar A, Lomax J, Mungall C, Neuhaus F, Rector AL, Rosse C (2005) Relations in biomedical ontologies. *Genome biology*. 6(5):1-5.

Smith B, Köhler J, Kumar A (2004) On the application of formal principles to life science data: A case study in the Gene Ontology. In Erhard Rahm (ed) *Data Integration in the Life Sciences, First International Workshop, DILS 2004*, Leipzig, Germany, March 25-26, 2004, (Lecture Notes in Computer Science 2994), Springer, 79–94.

Smith B, Kusnierczyk W, Ceusters W (2006) Towards a reference terminology for ontology research and development in the biomedical domain. *Proceedings of KR-MED* (CEUR, vol. 222), pp. 57-65.

Smith B and Welty C (2001). Ontology: Towards a new synthesis. In *Formal Ontology in Information Systems*), ACM Press, pp. 3-9.

Spear AD, Ceusters W, Smith B (2016) Functions in Basic Formal Ontology. Applied Ontology. 11(2):103-28.

Stevens H (2013) *Life out of Sequence*. University of Chicago Press, Chicago.

Thomas PD (2017) The Gene Ontology and the meaning of biological function. In *The Gene Ontology handbook*. Humana Press, New York, pp. 15-24.

Thomas PD, Hill DP, et al. (2019) Gene Ontology Causal Activity Modeling (GO-CAM) moves beyond GO annotations to structured descriptions of biological functions and systems. *Nature genetics*. 51(10):1429-33.

UniProt Consortium (2008) The <mark>Universal Protein Resource (UniProt</mark>). Nucleic Acids Research 36 (Database issue).

Vita R, Zheng J, Jackson R, Dooley D, Overton JA, Miller MA, Berrios DC, Scheuermann RH, He Y, McGinty HK, Brochhausen M (2021). Standardization of assay representation in the <mark>Ontology for Biomedical Investigations.</mark> *Database*. https://doi.org/10.1093/database/baab040

Walls RL, et al. (2012) Ontologies as integrative tools for plant science. *American journal of botany*. 99(8):1263-75. https://doi.org/10.3732/ajb.1200222

Whetzel PL et al. (2006) Development of <mark>FuGO:</mark> an ontology for functional genomics investigations. OMICS. 2006 Summer;10(2):199-204. doi: 10.1089/omi.2006.10.199

Wilkinson MD, et al. (2016) The <mark>FAIR</mark> Guiding Principles for scientific data management and stewardship. Scientific data. 3(1):1-9.