Lecture Notes in
The Philosophy of Computer Science

Matti Tedre

Department of Computer Science and Statistics, University of Joensuu, Finland

This .pdf file was created March 20, 2007.
Available at cs.joensuu.fi/~mmeri/teaching/2007/philcs/

4 Scientific Reasoning

- What do scientists mean when they say they know something?
- What is scientific explanation?
- An introduction to the basic schools in the philosophy of science.
- Is science objective?

Epistemology is a wide and diverse field, but generally it is characterized as the study of knowledge. Typically, epistemologists are concerned with distinguishing knowledge from opinions and good reasoning from poor reasoning (Cruz, 2003). One of the basic questions in epistemology is "What do we know?". This question is not even intuitively easy, and a great number of different approaches to the theory of knowledge has been proposed, all of them leading to even more questions: "What does knowing here mean?", "How do we know which beliefs of ours are true?" and "Are there *any* propositions that are absolutely true?". These questions are examples of *epistemological questions*. Regarding the term *knowledge*, epistemology, in its conventional forms, is concerned more about *propositions* like "*x* is true", rather than procedural knowledge, skills, arts, technology, or other knowledge about "how to do *y*".

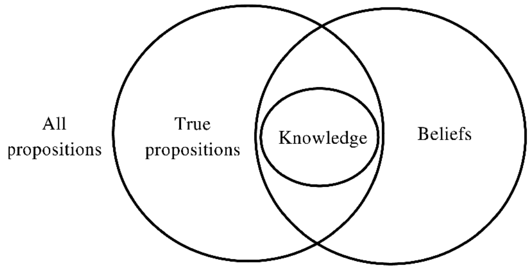

Figure 7: A classical view of knowledge

The classical tripartite of what is knowledge is depicted in Figure 7. The tripartite is that (1) there are propositions that are *true*, that (2) there are things that people *believe*, and that (3) *knowledge* is something that people believe and that is true. Already Plato, in Theaetetus, raised the question of whether true opinions are knowledge, though he did not get far with the question.

Perhaps partly because of its simplicity, partly because of its intuitiveness, it is common to hear the saying that *knowledge is justified true belief* (often called *the JTB analysis of knowledge*; see Steup, 2006).

THE JTB ANALYSIS OF KNOWLEDGE

An individual *S* knows that "*p*" if and only if:

i. *p* is true;

ii. *S* believes that *p*;

iii. *S* is justified in believing that *p*.

Substituting *p* with propositions like "the Earth is round" or "There is snow and ice at the summit of Mt. Everest" gives a good idea of the intuitiveness of the JTB analysis of knowledge. Edmund L.

Gettier's short article "*Is Justified True Belief Knowledge*" (Gettier, 1963), however, effectively refuted the JTB analysis of knowledge. What Gettier's article says is that if one, for wrong reasons, believes something to be true, and if, for other reasons, that belief luckily happens to be true, it cannot be said that one's belief is really *knowledge*. Although Gettier's refutation seems naïve at first sight, sometimes the best way to convey a point is to make it simplistic or to be blunt. For instance, suppose that Jussi, who is a computer programmer, has extensively tested a program and believes that

1. Program *w* works. (*P*)

(Actually, Jussi's program *w* has many bugs, but Jussi is not aware of those bugs.) Suppose that Jussi, for some strange reason, proposes that

2. Program *w* works or it rains in Helsinki ($P \vee Q$)

Because Jussi believes that program *w* works, and because he is familiar with logic, he also believes that proposition 2 is true. Suppose that by sheer luck, it happens to rain in Helsinki, but Jussi does not know about the rain. Now, logically speaking, proposition 2 is indeed true. Condition (*i*) of JTB analysis of knowledge is met: *It is true* that program *w* works or it rains in Helsinki. Condition (*ii*) of JTB analysis of knowledge is met: Jussi *believes* that proposition 2 is true. Also condition (*iii*) of JTB analysis of knowledge is met: Jussi has a *justified reason* to believe that proposition 2 is true (he has extensively tested the program and believes that it works). According to JTB analysis of knowledge, it should be said that Jussi *knows* proposition 2. Now common sense says that Jussi does not really *know* proposition 2, but he is just lucky because it happens to rain in Helsinki. Gettier's text is indeed usually considered to be an effective counterexample to JTB (Steup, 2006). Although "*knowledge is justified true belief*" is a nice slogan, knowledge is clearly more complex a concept than what JTB analysis suggests. Many modern epistemologists take the problem posed by Gettier's article as a sign of vagueness of the term *knowledge* or as a sign that there are still more conditions of true knowledge (Cruz, 2003; see also Steup, 2006).

Generally speaking, when philosophers have criticized the JTB analysis of knowledge, its condition (*i*) (*p* is true) has not been challenged very often—it is usually agreed that what is false cannot be "known". For instance, the statement that "magical dragons regularly eat sailors near the coastline of India" is false. Therefore, although one might *believe* that such things happen, it would be odd to say that anyone could *know* it. Also condition (*ii*) (*S* believes that *p*) is rarely challenged—it is usually agreed that if you say you *know* something to be true, you also need to believe it. (Note that in this context there is a difference between *knowing that* p *is true* and *being familiar with* p.) However, the condition (*iii*) (*S* is justified in believing that *p*) is controversial. Especially in the case of scientific knowledge, justification is a fiercely debated issue. The question of scientific justification is discussed later in this chapter. Before that, however, it is good to discuss briefly what *scientific explanation* is.

4.1 Explanation and Understanding, Very Shortly

I borrow the title of this section from philosopher Georg Henrik von Wright's (1916-2003) book *Explanation and Understanding* (Wright, 1971). *Explanation* is one of the oft-stated purposes of scientific research—the others being *description*, *exploration*, and *prediction* (Wright, 1971:1;

Okasha, 2002:1). The word pair *explanation* and *understanding* has famously been used in reference to the dichotomy between the aim of natural science (to explain) and the aim of history (to understand). Although practically all explanation can be considered to aim at furthering out understanding, the term *understanding* also has a psychological twist and special connections with intentionality and semantics, which the term *explanation* does not have (see Wright, 1971:6). (Note that not all scientists and not all philosophers of science think that science should explain anything. Those philosophers and scientists who advocate *descriptivism* argue that the only things one can require of science are description and prediction (Bunge, 1998b:59-61).)

Philosopher of science Mario Bunge listed six elementary problem forms in (natural) sciences: *which*, *where*, *why*, *whether*, *how*, and *what* (Bunge, 1998:196). Examples of such elementary problems are, for instance, “Which processes have the property *A*?”, “Under which circumstances (where) is *x* true?”, “What causes *p* to happen (why *p* happens)?”, “Is *q* true or false?”, “How does *c* happen?”, and “What properties does *c* have?”. In addition, he noted a number of classes of substantive scientific problems (Bunge, 1998:209). Bunge's list of substantive scientific problems includes a number of *empirical problems*; for instance, *problems of data collection* (such as observing, counting, and measuring) and *problems of making* (such as constructing and calibrating instruments). His list also includes a number of *conceptual problems*; for instance, *describing* (characterizing individuals and classes), *arranging* (classing and ordering sets), *elucidating* (interpreting signs and refining concepts), *deducing* (computing values, proving theorems, checking solutions, etc.), *building* (introducing new concepts, making empirical generalizations, building theories, etc.), and *metalogical problems* (uncovering and removing inconsistencies, proving independence, etc.) (Bunge, 1998:209).

Furthermore, research in education, social sciences, and humanities often concerns detailed and rich descriptions of phenomena, descriptions which are not aimed at universal generalizations but at deeply understanding the particular phenomena in question. That research is often aimed at (1) *exploring phenomena* rather than testing hypotheses; (2) emphasizing *unstructured data* instead of analytic categories; (3) focusing on *cases* in detail instead of large populations; and (4) *explicitly interpreting* the meanings and functions of human actions (Atkinson & Hammersley, 1994).

But science does not deal only with problems of empirical investigation and theory-building such as the ones above. There are also *strategy problems*, which include *methodological problems* and *valuational problems* (cf. Bunge, 1998:210). Methodological problems are about how research should be done. Bunge noted that methodological problems include problems of measurement and other conventions, problems of techniques, problems of experiment design, problems of theory design, and the analysis and criticism of all of those problems. He continued that contrary to some frequent claims, natural sciences are not value-free. Problems of valuation arise both in empirical science where scientists frequently make choices between research factors such as range, accuracy, reliability, versatility, and cost; and in theoretical science where scientists frequently make choices concerning their theories' coverage, depth, support from other fields, and even formal elegance (Bunge, 1998:210). Valuational problems are an integral part of scientific research although they do not show in the body of scientific knowledge (one could justly argue that not making value choices explicit is a weakness of science, especially of scientific reporting).

Mario Bunge warned against passing *pseudoexplanations* as real explanations (Bunge, 1998b:10-11). He wrote that *labeling* is an oft-used form of pseudoexplanation. It is not scientific explanation to slap a name tag on a phenomenon and claim to have explained it. For instance the following ones are labeling-type pseudoexplanations: "Why do people often fight each other? Because they have an aggression instinct.", "Why is this plant so widespread? Because of its encroaching nature.", "Why do some people behave more intelligently than others? Because of their high IQ." (cf. Bunge, 1998b:11). None of the explanations above actually explain anything—they just tag the phenomenon with a label (*aggression instinct*, *encroaching nature*, and *IQ*, respectively).

4.1.1 Causal Explanations

In natural sciences, explanations of phenomena are done with hypotheses, theories, and laws. Consider, for instance, the following call for explanation: "Why is this object (a rose) red?". One explanation might be, "That object is red because when white light arrives at its surface, the surface absorbs certain kinds of wavelengths and scatters other wavelengths" (members of our color research group may propose much more complex answers). That answer relies on the physical properties of the object, on the current lighting conditions, on the laws of physics, and on a convention of which wavelengths are considered to be red. It is an explanation of what *causes* an object to appear red.

Strictly speaking, in David Hume's (1711-1776) terms, causal explanation (*a* caused *b*) must meet three criteria (Hume, 1739:B.I, P.III, S.XIV). Firstly, *a* and *b* must be "contiguous in time"—that is, they must happen temporally close to each other, and *b* must not happen before *a*. Secondly, *a* and *b* must be "contiguous in place"—they must happen spatially close to each other. Thirdly, in Hume's words, there must be a "necessary connexion" between *a* and *b*. Terms *a* and *b* can refer to events, situations, or phenomena. In the explanation above, Hume's first and second condition are clearly met: white light *meets* the object surface, *and then* red light scatters. But Hume's third condition is a tricky question.

Although many causal explanations seem to apply universally, they are definitely not *necessary truths* or *logically necessary propositions*, that is, truths that simply could not be otherwise because of the meaning of their terms, such as "no bachelor is married" and "2+2=4" (cf. Shapiro, 2000:21-22). Already Hume noted that although it seems intuitively true that a billiard ball that hits another billiard ball *causes* the other ball to move, this cannot be *logically* proven true (Hume, 1739:B.I, P.III, S.XIV). According to Hume, one can say that a billiard hit another ball, *and then* the other ball moved. Although one cannot prove that an event caused another event, one can say that the two events "have been always conjoined together, and which in all past instances have been found inseparable" (Hume, 1739:B.I, P.III, S.VI). Of course, that one cannot *logically* prove causality is not a good reason to begin to suspect causality and adopt a habit of jumping off rooftops. Scientists have to have faith that causes and effects that have been constantly conjoined (such as those denoted by terms like gravity and impulse) continue to be conjoined. (Yet there again: What is the meaning of the famous logical fallacy *post hoc, ergo propter hoc*—"This happened after that, so that must cause this"—if one can never establish causality with certainty?)

According to philosopher Carl Hempel's (1905-1997) famous account of what constitutes a *scientific explanation* (Hempel, 1965:331-496), the above-mentioned explanation for the question

"Why is this object (a rose) red?" is certainly a scientific explanation. That is, the explanation above refers to *general laws* (from physics, especially chromatics), it includes *particular facts* (the material properties of that object), and from those general laws and particular facts, one can derive a conclusion that the object is red. Hempel's model has been referred to as *the covering law model*, as *the deductive-nomological model*, and as *the subsumption theory*. Hempel himself used the term *deductive-nomological model*, which refers to its deductive character and its reliance on laws (*nomos* in Greek).

Hempel's Model of Scientific Explanation	
General laws &	(*explanans*)
Particular facts	(*explanans*)
↓	
Phenomenon to be explained	(*explanandum*)

Hempel's model is simple: a scientific explanation relies on one or more general law(s), it includes particular facts that are at place before the phenomenon, and from these laws and facts (called the *explanans*), one can in a logical, unambiguous, and deductive manner derive the phenomenon to be explained (called the *explanandum*) (Hempel, 1965:336). In other words, if the question were, "Why q?", a scientific explanation of q follows logically from a number of general laws and particular facts (which do not include q) to q. For instance, if the question were, "Why is this rose red?", a scientific explanation includes explanans such as (1) the laws of chromatics (laws), (2) the particular physical properties of the object (fact), and (3) the particular lighting conditions (fact). From these explanans one can deduce the explanandum that the object indeed must appear red.

So in Hempel's model, explanations have at least the following characteristics: they answer to *why* questions, they refer to formulas and laws, and they are deduced logically. Bunge noted that there are two more characteristics: Firstly, explanation in terms of Hempel's model locates every phenomenon and explanation in a larger system of interrelated phenomena and laws; and secondly, the actual working process works inversely to deduction (Bunge, 1998b:7-8). Whereas in deductive explanation one has a number of explanans from which one logically proceeds to the explanandum, in the working process, scientists have observed a phenomenon and try to make hypotheses about the explanans that explain that phenomenon (Table 3).

Deductive process (explanation)	***The actual working process (finding explanans)***
Laws and facts	Which laws and facts? (unknown)
↓ *deductive process* ↓	↑ *hypothesizing* ↑
Phenomenon to be explained	Phenomenon to be explained (given)

Table 3: The Process of Scientific Explanation

The actual working process in Hempel's model of scientific explanation can be characterized with the *hypothetico-deductive model*, a description of the scientific method originally conceived by Karl Popper (see Figure 3 on page 79). The hypothetico-deductive model starts with observations of a phenomenon. Scientists can freely propose wild guesses, conjectures, and hunches about the phenomenon, and freely create falsifiable hypotheses as their attempts to explain the phenomenon (see page 120 of these lecture notes for *falsifiability*). All hypotheses have logical consequences, and it is a scientist's duty to test those logical consequences. When the test results support the predictions derived from the hypothesis, the tests are said to *corroborate* the hypothesis, and when tests conflict with the predictions derived from the hypothesis, the tests are said to *falsify* the hypothesis.

Causal explanations are best suited for the scientific cycle that begins with observations of a phenomenon that scientists wish to explain. Causal explanations cannot be used in formal fields which have no empirical connection, fields such as mathematics and logic. For instance, explanations such as "Algorithm a has two nested loops over the input data $\Rightarrow$ For algorithm a: $T(n)=O(n^2)$" are not causal explanations (the statements are not connected in time and place). But also in natural sciences there are many scientific statements that portray relationships between terms, properties, or statements, and which do not have causal connections. Consider, for instance, the statement "Molecule m has two hydrogen atoms and one oxygen atom $\Leftrightarrow$ m is a water molecule". Having two hydrogen atoms and one oxygen atom does not *cause m* to be a water molecule. The statement above just defines the term *water*.

One should ask which explanations in computer science are causal explanations and which are not. Similar, one should ask how well suited Hempel's model of scientific explanation is for computer science. Quite some computer scientists indeed work in areas where Hempel's model can provide guidelines for scientific explanation, and where causal explanations are needed. But engineering-oriented branches, for instance, do not aim at explanation at all; theoretical computer science does not explain causes; and many human-oriented branches seem to require more than causal explanations.

4.1.2 Functional Explanations

In the section above, scientific explanation was discussed from the physicist's viewpoint. But another scientist might give a different answer to the question "Why is this object (a rose) red?". A biologist might say, "Roses are red because red color attracts some pollinating insects and hence increases the flower's potential to reproduce". Or perhaps, "The function of red color is to attract pollinating insects". The biologist's explanation explains why it is *beneficial* for a rose to be red, and that explanation is based on the argument that red color gives the rose better reproductive potential than any other color would give. Functional explanations seem to be best suited for biology, but for instance, those proponents of evolutionary psychology who argue that how people think is a result of evolutionary processes, also employ functional explanations (see, e.g., Pinker, 2006).

Although physicist's and biologist's explanations are of different kind, they both seem to explain the phenomenon well, from their own viewpoints. The physicist's explanation is a causal explanation. That is, it relies on cause and effect: when light meets the surface of a rose, properties of that surface cause certain kinds of wavelengths to be absorbed or scattered. In physics all explanations of events are causal explanations (yet explaining *events* is not all that physicists do—they also make statements about the properties of the world, such as "matter is made of atoms" or "there are neutrons"). The biologist's explanation is a *functional explanation*. That is, it relies on the fact that some properties of living organisms are beneficial to individual organisms, to populations, or to whole species. For instance, being red is beneficial for roses because red color attracts some pollinating insects. In biology there are causal explanations and functional explanations.

Causal explanations are characteristically employed to explain nonliving things whereas functional explanations are characteristically employed to explain living things. The crucial difference is that functional explanations are often of form *what is beneficial for c*, but with nonliving things such explanations are nonsensical. With biological phenomena, if the explanans are (1) insects pollinate roses, (2) red color attracts some pollinating insects, and (3) pollination increases the reproductive potential of a rose (and that is beneficial for roses), then one could argue that "roses are red" follows from the explanans (but is not straightforwardly *deduced* from the explanans). With physical phenomena it would be odd to say that it is beneficial for a surface to scatter red wavelengths and absorb other wavelengths.

4.1.3 Intentional Explanations

Von Wright went on to note that causal and functional explanations do not fit well history, social sciences, or other sciences that attempt to explain people's actions (Wright, 1971:Ch.4). For instance, causal and functional explanations cannot explain why Andrés gave Helena a red rose. Because people's actions have purposes and because those actions have intentions behind them, yet another kind of explanation is needed. Von Wright introduced the following inference schema, which he called *practical inference*:

> A *intends to bring about* p.
>
> A *considers that he* [sic] *cannot bring about* p *unless he does* a.
>
> *Therefore* A *sets himself to do* a.
>
> (Wright, 1971:96)

Explaining why *A* did *a* requires knowing *A*'s intention to bring about *p* and knowing *A*'s belief that doing *a* will bring about *p*. For instance, consider that a group of bystanders see Carolina kick her computer. Those bystanders cannot explain or understand Carolina's actions unless they know that inside Carolina's computer there is a loose cable, and Carolina has found out that a brisk kick on the side of the computer temporarily fixes the problem. Carolina's actions can be explained by her *intentions* (temporarily fixing the computer) and her *beliefs* (kicking the computer temporarily fixes it).

It is another matter altogether to what extent one can ever hope to know or understand other people's intentions and beliefs—they are often inextricable from other beliefs, valuations, and values that people hold. Not all intentions and beliefs can be easily isolated and communicated to other people. Nonetheless, von Wright argued that *functional explanations* and *intentional explanations* are two subprovinces of *teleological explanations* (Wright, 1971:16).

It is important to note that at least two of the aforementioned kinds of explanation—*causal* and *intentional* explanation—have their uses in computer science but have their limitations, too. Some branches of computer science, such as signal processing, require causal explanations (e.g., "Random movement of electrons causes noise in electronic circuits"). Functional explanations ("it is beneficial for the electronic circuits to...") or intentional explanations ("the intention of the electronic circuits is to...") would sound weird. Some other areas in computer science may require understanding the intentions and beliefs of users ("The users clicked icon *i* because they believed that clicking it brings about *p*"). And some areas of computer science, such as branches of theoretical computer science and some engineering-oriented branches, do not explain the world at all.

4.1.4 Statistical Explanations

There is one more particularly common, inductive-type of explanation in modern sciences that must be noted here—*statistical explanations*[22]. Statistical explanations differ from the above-mentioned *explanations-in-a-strong-sense* in that statistical explanation is not aimed at explaining the reasons of why things happen, but at depicting *regularities* that arise from an interplay of a large number of objects or events (Bunge, 1998b:39). Statistical explanations are formulated in reference to statistical patterns or frequencies, they assume randomness, and they require a relatively large sample of independent objects or phenomena.

Descriptive statistics offers summaries or analyses, which are derived from observation data, and used to summarize or describe some selected aspects of that data. The statistical tools in descriptive statistics often include things like means, medians, and deviations. *Inferential* statistics offers characterizations of or predictions about a population of objects or phenomena, characterizations or predictions which are inferred from a smaller sample of that population using descriptive statistics. Mario Bunge described statistical explanation as follows:

> [A] *statistical explanation consists in discovering how collective (statistical) patterns result from the interplay of numerous individuals of a certain kind (molecules, persons, etc.) Such patterns emerge when certain individual variations cancel out and build up* stable averages*, so that the former can be neglected (at the macrolevel) and attention can be focused on the latter.*
>
> (Bunge, 1998b:39, emphasis in original)

Statistical explanations in computer science often describe probabilities, such as

Algorithm *a* outperformed algorithm *b* in a fraction *f* of the cases in sample *C*
Sample *C* is representative of the whole population of possible cases
---
Algorithm *a* outperforms algorithm *b* in a fraction $f \pm c$ of the whole population of possible cases

(where *c* is a small constant depicting the margin of error) and

A fraction *f* of α are β
x is an α

The probability that *x* is a β equals *f*

Note that the former statement does not explain *why* algorithm *a* outperformed algorithm *b*, and the latter statement (cf. Bunge, 1998b:41) does not explain *why* a fraction of α are β. The statements above depict frequencies of phenomena in a given population, but they do not give reasons for those frequencies. If one asked, "Why did algorithm *a* outperform *b*?", it would be odd to answer, "Because algorithm *a* outperforms *b* in 94%±1% of all cases".

Although statistical models may be weak in explanation of phenomena, they can often be used for description of phenomena, and they are particularly good at the fourth aim of science, prediction of future phenomena. Statistical explanations are superior to other forms of explanation especially

22 Woodward, 2003, is a good text about scientific explanation, including statistical explanations.

with prediction of indeterministic or incompletely understood phenomena. For example, causal and functional explanations can be used to explain the mechanisms of why penicillin (usually) cures a streptococcus infection, but one cannot justify a *general law* that stated that in all cases of streptococcus infection, penicillin will lead to recovery. Using inferential statistics, one can still say that receiving penicillin will cure patient with high likelihood (Hempel, 1965:381-382). Statistical explanations can be used to inform choices, policies, or interventions, and statistical analysis of observation data can inform scientists about the correctness of their hypotheses.