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 26, 2007.
Available at cs.joensuu.fi/~mmeri/teaching/2007/philcs/

4.2 Attempts to Verify Facts

In the history of science there have been a number of attempts to formalize the soul of science, and of those attempts, two are often highlighted: *empiricism* and *rationalism.* The empiricists and the rationalists both believe that scientific knowledge is derived from the facts rather than from opinions (Markie, 2004). However, they differ greatly in their views on the nature of those facts and on how can one get to know those facts. In today's computer science modern variations of both empiricist and rationalist traditions are well-represented. That is due to the fact that computer science deals both with experimental, observable things (a posteriori knowledge) and with theoretical, axiomatic systems (a priori knowledge).

The empiricist tradition can be seen in fields such as software engineering, computer vision, and human-computer interaction. The rationalist tradition can be seen in fields such as computational complexity theory and formal methods. But many or most branches of computer science incorporate aspects from both empiricist and rationalist traditions; take, for instance, information retrieval, image processing, and distributed computing.

That some branches of computer science incorporate and mix both empiricist and rationalist ideas can be a progressive trait: theoretical results can be empirically checked and empirical research can be built upon formal analysis. However, mingling research traditions can also be a source of confusion, because there is a fundamental difference between the methods, results, and constraints of each tradition. In this chapter, though, the focus is on the empiricist bent of computer science.

In the history of science there have been a great variety of empiricist approaches to research. Here two of them are discussed—positivism and inductivism—because both of those approaches entail the idea that science can prove some facts to be true. The terms *positivism* and *inductivism* are used in the following pages in the manner they are used in the history and philosophy of science (most commonly in the context of natural sciences, especially physics).

There is so little difference between the early 20th century versions of *logical positivism* and *logical empiricism*, that the terms logical positivism and logical empiricism are here used synonymously. For some reason, the term *logical positivism* seems to be used more often in a negative sense whereas the term *logical empiricism* more often carries a positive connotation. Although logical positivism is not an active school in the philosophy of science today, it had such a fundamental influence on scientific thinking that it is covered here. And although inductivism has changed radically from its naïve form presented here, that naïve form links directly to the problem of induction, and is therefore portrayed here.

4.2.1 Logical Positivism

Logical positivists held that there are final truths—truths that are universal and unchanging. The positivists' trust in scientific facts is well portrayed in Morris Cohen and Ernest Nagel's 1934 book, in which they stated, "we are sure that foolish opinions will be readily eliminated by the shock of facts" (Cohen and Nagel, 1934:402). Logical positivists (hereafter also *positivists*) believed that researchers can find final truths about the world and prove them correct, i.e., *verify* them. Hence the term *positivism*: one might call a proof a positive result and a refutation a negative result. Positivists shared the beliefs that (1) there can be a *united method* for all scientific investigation; (2) scientists in all disciplines should adopt the quantitative, *exact methodology* of natural sciences; and (3) phenomena must be explained in terms of *causal explanations*, not in terms of intentions, goals, or purposes (Wright, 1971:4). There are three assumptions about *facts* that are central to positivism:

> *(a) Facts are directly given to careful, unprejudiced observers via the senses.*
>
> *(b) Facts are prior to and independent of theory.*
>
> *(c) Facts constitute a firm and reliable foundation for scientific knowledge.*
>
> (Chalmers, 1999:4)

In other words, the first assumption states that one can get to know *facts* by observing a phenomenon—listening, seeing, smelling, touching, tasting (*facts* about the world are unbiased, unambiguous, and independent of anyone's feelings towards them). The second assumption states that facts are what they are regardless of any theoretical frameworks (though theories can be derived from facts). The third assumption states that *scientific knowledge* can be derived from facts. Those assumptions seem commonsensical, but under closer scrutiny all of them face unsurmountable difficulties.

First of all, the assumption that "*facts are directly given to careful, unprejudiced observers via the senses*" seems difficult to maintain. One would have to assume that two normal observers would, under the same conditions, always observe the same thing. Chalmers wrote that an example of the fact that senses alone cannot be the basis of observation is clear for anyone having had to learn to see through a microscope. When a beginner looks through a microscope at a slide prepared by an instructor, he or she rarely discerns the appropriate cell structures even though the instructor has no difficulty discerning them when looking at the same slide through the same microscope, Chalmers noted (Chalmers, 1999:7).

Even though it is evident that the human biological properties are a major cause of what one sees, what one thinks he or she sees is also dependent on one's previous experiences, cultural upbringing, and expectations. When one looks at the horizon at the Himalaya, one person may see a mountain range, another person can see forms of terrain, and yet another person may see a number of individual mountains. Although all the people see (*perceive*) the same thing, what they interpret they see (*discern*) is an individual experience, and there is no "correct" interpretation (cf. Smith, 2002c:239-240).

Second, the assumption "*facts are prior to and independent of theory*" is also difficult to maintain. When scientists make science, they do not share the things that they see, hear, taste, smell, or feel—they share *statements* about those things. It is impossible that every scientist would repeat every experiment that science is built on. Scientists base a lot of their science on what they read or hear—on statements of facts. Single statements of facts, such as "Matter is made of atoms", rarely make sense alone, but they require a linguistic and conceptual framework within which they are interpreted. To make sense, the atomic statements are anchored in and related to this framework.

If statements of fact are not determined in a straightforward way by the senses, and if observation statements presuppose some knowledge, it cannot be the case that one first establishes the facts and then derives knowledge from them (Chalmers, 1999:12). In computing fields that problem can be seen in the often cherished distinction between *data*, *information*, *knowledge*, and *wisdom*. It is common to prune those concepts to some sort of a variation of the following: *datum* is a basic attribute of information; *information* is data with a conceptual commitment and interpretation; *knowledge* comes from knowing how to use the information; and *wisdom* adds the understanding of when and where the knowledge is applicable. This simplification faces some fundamentally difficult objections.

Data is not collected randomly, and in the data collection process, only a minor part of (infinite) available data is collected. The choice of which data to collect and which not, is done by using what one already knows about the domain. The data is collected using formal or informal data structures, which are usually conceptual aggregates that are created to record or model some aspects of the phenomenon as well as possible. Because data structures are built according to what one already knows about the domain, data cannot be independent of one's previous knowledge (cf. Brian Cantwell Smith's term *inscription error*, found in Smith, 1998:49-50).

Third, there is a problem with the assumption that "*facts constitute a firm and reliable foundation for scientific knowledge*". *Scientific knowledge*, in this interpretation, constitutes of facts that are combined logically. Now, there is a common misconception that has to be made clear. Logical truths do not carry any information about reality (Bremer, 2003). Logic can only tell that if the premises are true and if the argument is valid, then the conclusion must be true.

For instance, let premise p_1 be "Every human being can use the C++ programming language" and p_2 be "My grandmother is a human being". From these premises, it is a valid deduction d that "My grandmother can use the C++ programming language". In this case, it happens to be that p_1 and d are false, as my grandmother has never been interested in learning the C++ language. But this does not affect the fact that this logical argument is perfectly valid. That is to say, logical derivations cannot tell anything about the truth of its premises. The truthfulness of premises in natural sciences must be ascertained by some other means than logic.

So, according to the rules of logic, if one can be sure that the premises are true, one can be equally sure that everything logically derived from the premises will also be true. So, the fate of the positivist science relies on the degree to which its premises (facts) are true. But essentially, affirming that scientific facts are true can be very difficult. The fact that sun has risen every day for the last 4.5 billion years does not conclusively prove that the sun will rise tomorrow (and one "day" it probably will not). A science built solely on logical deduction works only if one accepts the facts on which that science is built. While reading the rest of this week's lecture notes it might be useful to frequently ask whether accepting scientific facts is an act of faith.

4.2.2 Naïve Inductivism

The inductivists understood the problem of basing growth of knowledge on logical deduction and careful observation of facts, and they took another approach. The inductivists base their scientific facts on experimentation, not on single observational facts. The inductivists believe that instead of logical deduction, scientific facts can be made by generalizing from a large number of observations. Philosopher Alan Chalmers summed up the (naïve) principle of induction as, "*If a large number of **A**'s have been observed under a wide variety of conditions, and if all those **A**'s without exception* *possess the property **B**, then all **A**'s have the property **B***" (Chalmers, 1999:47). He wrote that in order to have a valid inductive argument, the following conditions must be met:

> *1. The number of observations forming the basis of a generalization must be large.*
>
> *2. The observations must be repeated under a wide variety of conditions.*
>
> *3. No accepted observation statement should conflict with the derived law.*
>
> (Chalmers, 1999:46)

The problems of the first condition of the inductive argument are apparent: "What is a *large* sample?", "Is a sample of hundred, thousand, or million large?", "Does the concept of *large* depend on the object of investigation?". Certainly, in a study of a rare disease, sample size has to be different from a study of quarks. If one studies web pages to find "cultural markers", how large a number of observations is enough for a generalization? Should one study one hundred web pages or one thousand—or all of them? Statistics can offer some answers about reliability levels, but only referential ones. "Qualified" facts, no matter how well they meet some preordained probability classes, can hardly be called *facts* in the strict sense, in which early inductivists wanted to take them.

It is also hard to draw the line with condition 2 of the inductivist argument ("*the observations must be repeated under a wide variety of conditions*"). Fulfilling condition 2 requires answers to questions such as "What kinds of variation should there be?" and "How much variation should there be?". It is an insuperable problem with inductivism that the list of variables can be extended indefinitely by endlessly adding further variations to the test setting. For instance, in the case of user interfaces, the variations added can be such as the size of the screen, the brand of the computer used, and the distance to the closest accordionist. Unless negligible (*superfluous*) variations can be eliminated, the conditions under which an inductive inference can be accepted can never be satisfied (Chalmers, 1999:48). The problem, however, is telling significant variables from superfluous variables.

The same problem that faced positivism, the problem of theory-independence, haunts also inductivism. When the researcher excludes some variables as negligible or superfluous, the researcher has to rely on some prior knowledge about the test situation. This prior knowledge requires some prior inductive arguments, which, in turn, require other prior inductive arguments. David Hume noted that the problem of induction is how the principle of induction itself is justified (Hume, 1777:SIV,PI:20-27;PII:28-33). Hume wrote that one can use two possible arguments to justify induction.

First, one could try to justify induction by appealing to logical necessity. But nothing really proves *logically* that the future has anything to do with the past (although *empirically* and *intuitively* that seems clear). Second, one can try to justify induction by appealing to the past successes of induction: induction has worked in the past, so it will probably continue to work in the future. But in this second case induction is justified with—induction! That is:

Induction worked in case c_1
Induction worked in case c_2
...
Induction worked in case c_n

Induction works

It seems quite strange to argue that induction is a good mode of argumentation by using an inductive argument. But although inductive arguments are not conclusive proofs-in-the-strict-sense, they are commonplace in everyday science-making. One might look at, for instance, how Newton came up with his law of universal gravitation, “Every single point mass (object) in the universe attracts every other point mass in the universe”. Certainly he did not observe every single object in the universe. Instead, he made a number of observations that all built up his confidence about his law, and at some point he became convinced enough to announce his law. One can say that Newton found very compelling evidence for his law, but his findings do not rule out the logical possibility that somewhere in the universe there might be objects that defy Newton's law. Mario Bunge has listed a number of commonplace, *heuristic* inductive arguments (Bunge, 1998b:327-329—see Table 4).

Analogy	***Pattern***		
Substantive analogy *Similarity of components*	a is $P_1, P_2, P_3, \ldots$, and P_n b is $P_1, P_2, P_3, \ldots$, and P_{n-1} It is likely that b is also P_n		
Structural analogy *Form similarity*	Systems with a similar form (structure or law) often share a number of other properties a and b have the same structure (they “obey” formally similar laws) It is likely that a and b share further properties		
First-degree induction *From instances to lower-level generalizations*	All A up to the n-th level were found to be B It is likely that all A are B		
Second degree of induction *From lower-level to higher-level generalizations*	Law L holds for every set S_i up to the n-th It is likely that law L holds for every set S_i		
Statistical generalization *Sample-population inference*	S is a random sample of the population of U The observed frequency of P's in the random sample S of U is p The expected frequency of P's in U is near p		
Statistical specification *Population-sample inference*	The observed frequency of P's in the population U is p S is a random sample of U The expected frequency of P's in S is near p		
Weak modus ponens	If p then q p is plausible q is plausible	“If p then q” is plausible p q is plausible	“If p then q” is plausible p is plausible q is plausible
Weak modus tollens	If p then q $\neg q$ is plausible $\neg p$ is plausible	“If p then q” is plausible $\neg q$ $\neg p$ is plausible	“If p then q” is plausible $\neg q$ is plausible $\neg p$ is plausible
Strong reduction	If p then q q It is possible that “p” is true		

Analogy	***Pattern***		
Weak reduction	If p then q q is plausible It is possible that p	"If p then q" is plausible q It is possible that p	"If p then q" is plausible q is plausible It is possible that p

Table 4: Oft-recurring Patterns of Inductive Inference (adapted from Bunge, 1998b:327-329)

The patterns in Table 4 are not *rules*, but heuristic patterns of thinking. Their conclusions never follow with certainty. What follows, are educated guesses—Bunge calls this inference "seductive logic" (Bunge, 1998b:329). The patterns in Table 4 are used when scientists generate hypotheses; when they make conjectures about how the world might work.

Hume's problem of induction has been answered in a number of ways, but one specifically lucid reply is considered here. Philosopher Peter Strawson wrote that induction is so fundamental to our reasoning and thinking that it is strange to require that claims about induction should be treated by the same token with other claims (see, e.g., Salmon, 1974 in Swinburne, 1974:56; Okasha, 2002:28). Strawson compared induction to *the legislation*. It makes perfectly sense to consult the law books to check if some actions are legal. But it makes no sense to consult the law books to check if the law is legal. That one cannot justify the law relying on the law does not, however, mean that the law is exempt from criticism—one can ask, for instance, does a particular legal system serve its purposes and if some other legal system might serve those purposes better. In a similar sense one can utilize inductive reasoning in making science, but debates about inductivism itself should take place on a different level, with a different vocabulary.

In conclusion, neither positivism nor inductivism provide a durable basis for *facts* in the strictest sense of the word. They face difficulties that have not been overcome. Yet very few natural scientists today, given some time to think, would actually argue that they mean their statements about the world to be facts in the strictest, universal, timeless sense. Many scientists see—consciously or intuitively—their statements about the world in the *falsificationist* sense.

4.3 Proving Is Impossible But Falsifying Is Possible

In the previous section two kinds of approaches to science were discussed. The first one was positivism, advocated by a group of famous philosophers and scientists who met weekly in Vienna in the 1920s and 1930s, a group called the *Vienna Circle*. In their opinion, there should be strict criteria for categorizing scientific claims either true, false, or meaningless. The second one was inductivism; inductivists use observation and experimentation in trying to find universals: "Jussi's program produced the correct output with input x_1; his program produced the correct output also with input x_2". An inductivist would continue this until finally, after a large enough number of tests has been done, he or she would come to a universal conclusion: "Jussi's program produces correct output with any input x". However, there might always be a special case where Jussi's program would not work correctly.

4.3.1 Falsificationism

Seeing the problems of *proving* scientific claims true, sir Karl Popper proposed an alternative to the logical positivism, the alternative that is often referred to as *falsificationism*. Popper's 1934 magnum opus *The Logic of Scientific Discovery* (Popper, 1959[1934]) was a groundbreaking work in the philosophy of science. Popper turned the positivists' thinking around and stated that there could not be a way to prove universal truths about the world, but observations could be used in falsifying claims. For instance, if the claim is “Jussi's program produces the correct output on any input *x*”, then no number of observations of successful functioning can *prove* the claim, but one single observation of erroneous functioning can *falsify* the claim, that is, to show that the claim is incorrect.

However, if one sees science from the falsificationist point of view, it can never be said that a theory is true, but it can only be said that it is the best theory currently available—since it has not yet been successfully falsified. Falsificationist science does not constitute of *proven facts* but of theories that nobody has yet been able to falsify. All scientific theories are exposed to ruthless experimental testing, and only those theories that pass all the tests they are subjected to, constitute *scientific knowledge*.

In the falsificationist philosophy, it does not really matter how scientists arrive at their theories. Theories can be speculative, or they can be just guesses or hunches. The only thing that can be required of a scientific theory is that it must be *falsifiable.* That is, there must be some kind of possible observations that can falsify (refute) the theory. If scientists indeed observe such phenomena that falsify the theory, the theory must be abandoned. For instance, the statement “*Algorithm* a *is faster than algorithm* b” is falsifiable, because if one can show a situation where algorithm *b* outperformed algorithm *a*, the statement would be falsified.

Note, however, that it is totally different from the previous, falsifiable statement to argue, for instance, that “*The time complexity of algorithm* b *is higher than that of algorithm* a” (this argument is about axiomatic systems that are not empirically testable). Or, “*Using input* i, *algorithm* a *outperformed* b *by a factor of* 12:1” (this argument is about a previous single test, which cannot be refuted but which can be repeated and questioned). Or, “*Using data set* S (n=1986) *algorithm* b *was, on average, more than 22% slower than algorithm* a” (this argument is a statistical report of previous tests). Comparing algorithms is not a trivial issue, and it can be in a number of different ways.

Tautologically true statements, such as “*Program* p *works correctly or it is flawed*”, are not falsifiable, and thus cannot constitute a part of scientific knowledge (they are fundamentally uninformative). Statements that entail possibility, such as “*Educational tool* t *can help learning*” are usually not falsifiable, because they cannot be disproven—no matter how many cases one finds where educational tool *t* does not help learning, there can always be a case where educational tool *t* can help learning. The argument “*There are computable functions which the Turing Machine cannot compute*” is not falsifiable because of the *definition* of computable functions: computable functions are defined as those functions that can be computed by the Turing Machine. The previous statement is meaningless because if a function cannot be computed with the Turing Machine, it is not considered to be computable.

But there are also a plenty of meaningful, non-tautological statements in computer science that are not falsifiable, and as such, could not be considered to be scientific statements in falsificationist view of science. For instance, statements such as “*Program* p *halts with input* i” are not experimentally falsifiable. That is, if it seems that the program does not halt with input *i*, perhaps one should wait just a little bit more, and a little bit more, and a little bit more, ... ad infinitum. There is always a possibility that if one waited one more minute, the program would finish.

4.3.2 Criticism of Falsificationism

The crucial criticism of falsificationism at large came from the field of natural sciences, where falsificationism allegedly suits best. The criticism was that a statement can never be conclusively falsified, because the possibility cannot be ruled out that some part of the complex test situation, other than the statement under test, is responsible for erroneous prediction. This weakness in falsificationism is pointed out in what is often called the *Duhem-Quine thesis*.

Pierre Duhem (1861-1916), who was a physicist and a philosopher of science, is usually credited with the thesis, but the logician, Willard van Orman Quine (1908-2000), developed Duhem's thesis substantially. Originally Duhem wrote, "*an experiment* [...] *can never condemn an isolated hypothesis but only a whole theoretical group*" (Duhem, 1977[1914]:183-188). That is, an experiment cannot falsify only one aspect of a theory, but the whole theoretical framework. This is because one cannot be sure if the abnormal findings result from a flaw in the theory, the instruments used in the experiment, the theories about how the instruments work, or something in the test setting.

In his famous essay *Two Dogmas of Empiricism* (Quine, 1980:20-46), Quine noted that *fact verification* survives only in the supposition that every single statement can be isolated from all other statements and confirmed or refuted. Quine argued further that all human knowledge has significance in every statement. One cannot make single scientific statements that are disconnected from other human knowledge. Quine's critique of verification applies to falsification too—if tests seem to falsify a theory or a statement, nothing can logically prove that the theory is flawed, but any other part of the test-situation, any other part of science, or even any other part of human knowledge, can lie behind the anomalous results.

> UNDERDETERMINATION
>
> It has occurred often in the history of science that there are two or more competing theories that explain the same phenomenon equally well. If there is not enough evidence to decisively single out the best theory among its rivals, the theory choice is said to be *underdetermined* by evidence. There are variations of the underdetermination thesis. For instance, if there is not enough evidence or knowledge to decisively determine between two contradicting statements, the choice between those two statements can be said to be underdetermined by evidence.

The problem raised by the Duhem-Quine thesis is especially lucid in empirical computer science. For instance, suppose that the statement is, "*Jussi's program produces correct output with any input*". That statement is perfectly falsifiable: if one finds that with input *j* Jussi's program produces incorrect output, the statement is falsified. However, if program testing team found a case of incorrect functioning, they still could not rule out the possibility that the program might be correct but the incorrect functioning is caused by the operating system, the physical machine, cosmic rays, or something else. This is not merely philosophical speculation. Quite many experienced programmers have, for instance, been baffled with a strangely functioning program, and painfully tracked it down to a compiler bug. The strength of formal proofs of program correctness is that they can be used to limit one source of errors in computer systems.

Popper acknowledged the problem caused by the Duhem-Quine thesis, and responded that when it is impossible to decide by experiment between two theories (or decide which part of the whole system is faulty), techniques of measurement have to be improved first (Popper, 1959:56,108). This does not eradicate the problem raised by the thesis, though. That is, it does not eradicate the problem that if one gets abnormal findings, one cannot know, with certainty, which part of the whole test situation plus theoretical framework is faulty.

Underdetermination is indeed an interesting problem in computer science. For instance, if there are different models that faithfully simulate and predict a phenomenon, how does one determine which model to utilize? However, it is even more interesting to note that for the majority of phenomena, there is an infinite number of possible modeling schemes, all incomplete and imperfect. If all the models of a certain phenomenon are defective, but in different ways, it could be said that the phenomenon is *underrepresented* by the models. The interesting question, from the computer scientists' point of view is, “Which one of these incomplete models to choose?”. If a number of models model and predict well different aspects of a phenomenon, the choice is ultimately subjective and beyond formalization: “Which aspects of the phenomenon should be emphasized at the cost of others?”.

Another problem with falsificationism, a historical one, comes from the fact that most scientific theories have actually been falsified in their early stages. For instance, the early forms of sir Isaac Newton's (1643-1727) gravitational theory was falsified by observations of the moon's orbit, in early versions of Niels Bohr's (1885-1962) atomic theory there were inconsistencies with classical electromagnetic theory and observations, and James C. Maxwell's (1831-1879) kinetic theory of gases was falsified by Maxwell himself (see Chalmers, 1999:91-92). Paul Feyerabend commented on falsificationism, arguing that no new and revolutionary theory is ever formulated in a manner that permits one to say under what exact circumstances one must regard it as endangered: *Many revolutionary theories are unfalsifiable* (Feyerabend, 1975). He argued that theories have formal flaws, they may contain contradictions, ad hoc adjustments, and so on. Falsificationist criteria are clear, unambiguous, and precisely formulated. This would, Feyerabend wrote, be an advantage if science itself was clear, unambiguous, and precisely formulated. “*Fortunately, it is not*” (Feyerabend, 1975).

Conceding to this criticism, Popper answered:

> *I have always stressed the need for some dogmatism: the dogmatic scientist has an important role to play. If we give in to criticism too easily, we shall never find out where the real power of our theory lies.*
>
> (Popper, 1970:55)

Thus Popper at once relinquished the falsificationist thesis. If dogmatism has a positive role to play in an account of science that is based on ruthless criticism, a number of questions arise: “Where does one draw the line?”, “When do scientists need to be critical and when do they need to be dogmatic?”, “Is it not so that if scientists are allowed to be both dogmatic and critical, then scientists can hold any beliefs they choose to, and as long as they choose to?”.

Take, for instance, three examples from software engineering: “GOTOS increase code entropy”, “Strong typing reduces run-time errors”, and “Good modularization reduces maintenance costs” (Snelting, 1998). Those statements are good examples of statements that are not very clearly falsifiable (if they are falsifiable at all). Although intuitively plausible, in the falsificationist paradigm those statements ought to be classified pseudo-science and they could not constitute a part of computer science proper.

First, GOTOS may not always increase code entropy. An example of this is the debate that Edsger W. Dijkstra's article “GO TO Statement Considered Harmful” (Dijkstra, 1968) raised. After Frank Rubin's response to Dijkstra (Rubin, 1987), almost twenty different versions of Rubin's example were published in the same journal (Communications of the ACM)—many of them for GOTOS, many against. In addition, when structured programming was introduced, it soon was accepted and taken into use in almost all organizations, although no research was ever performed to demonstrate that the claimed and hyped value of structured programming existed (Glass, 2005). Computer scientists trusted their gut feeling about structured programming.

Second, there are few or no empirical experiments that tested the “strong typing-theory” with two versions of the same language, one strongly and one weakly typed. Such empirical experiments are a sine qua non of falsificationist science. Furthermore, the distinction between strongly and weakly typed languages is vague. It is difficult to say under which conditions the “strong-typing theory” is falsifiable.

Third, the problem of the “modularization argument” is its use of the term *good*. Goodness is exceedingly unambiguous term, and without specifically specifying what is considered to be good modularization, the value of the modularization argument in a falsificationist science is dubious. In addition, the reduction of maintenance costs might not be measurable in any straightforward, objective manner.

From the falsificationist perspective, those three intuitively correct “folk theorems” are bad hypotheses, and they could not constitute a part of computer science. They might best be treated as *engineering heuristics*. The problem of computer science, from the falsificationist perspective, is that in computer science not nearly all hypotheses are measurable and falsifiable. There is a lot of software engineering, computational modeling, computer visualization, scientific computation, artificial intelligence, cognitive science, and human-computer interaction that do not meet the falsificationist criteria, but that have contributed greatly to our knowledge about automatic computation and our scientific practices.

4.4 Science as a Contract

In light of previous chapters, it seems that on the quest for shedding light on the growth of scientific knowledge or on the mechanisms of science at large, the philosophy of science before the 1960s was too confined to either theories, observations, experimentation, or the relationships between them. The accounts of science presented above are highly idealized *normative* accounts of how good science *should be done*. But in reality, scientists do not seem to follow any strict guidelines to the letter. A good *descriptive* account of science should take into consideration how scientists *really* work.

In the early 1960s, Thomas Kuhn; who was a physicist, a historian of science, and a philosopher of science; noted that the history of the progress of natural sciences cannot be explained by the positivist or the falsificationist accounts. Kuhn's research showed that successful scientists have not actually worked as the positivists or falsificationists argued they should have worked. Kuhn's book *The Structure of Scientific Revolutions* (Kuhn, 1996[1962]) was especially a criticism of falsificationism, which was the prevailing account of science at the time. Kuhn boldly confronted Popper:

> *...what scientists never do when confronted by even severe and prolonged anomalies* [is give up the paradigm.] *Though they may start to lose faith and then to consider alternatives, they do not renounce the paradigm that has led them into crisis.*
>
> (Kuhn, 1996:77)

In other words, when scientists find contradictions between their theories and their experimental findings, they do not do the falsificationist thing and abandon their theories. Instead, Kuhn argued, they cast aside the abnormal findings, or try to somehow find a way to accommodate their abnormal findings with their theories. Kuhn hit a nerve in Popper's theory, when he criticized falsificationism:

> *No theory ever solves all the puzzles with which it is confronted at a given time; nor are the solutions already achieved often perfect. On the contrary, it is just the incompleteness and imperfection of the existing data-theory fit that, at any time, define many of the puzzles that characterize normal science. If any and every failure to fit were ground for theory rejection, all theories ought to be rejected at all times.*
>
> (Kuhn, 1996:146)

In other words, Kuhn noted that at any given moment of time, scientific theories cannot fully and without problems explain everything. Imperfection and incompleteness of scientific theories do not, however, mean that scientists should abandon those theories—imperfection and incompleteness work as pointers for scientists towards those parts of the theories they need to refine. Kuhn also criticized the falsificationists' constant struggle to try to falsify any and every theory. From the Kuhnian point of view, scientists *need* theories that are taken for granted, and that are not questioned all the time.

4.4.1 The Structure of Scientific Revolutions

In *The Structure of Scientific Revolutions* Kuhn introduced a whole new vocabulary to the philosophy of science. He also presented a dynamic new image of how science works. Figure 8 presents Kuhn's model of how scientific progress happens.

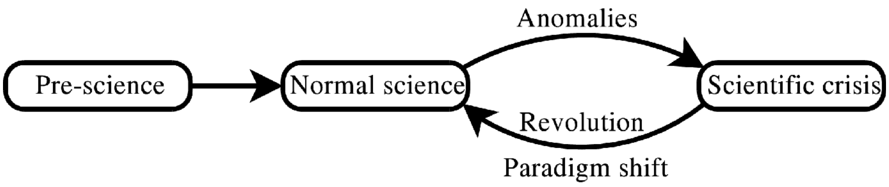

Figure 8: Kuhn's Model of Scientific Progress

In the first stage of Kuhn's model, *pre-science*, the particular scientific field has disagreeing coteries or competing theories of explanations. For instance, in the early 1900s there were a number of competing approaches to automatic computation. At the time, different kinds of analog differential analyzers for automatic calculation of integrals were quite successfully built. In this early exploratory phase, there were many promising candidates for a *scientific paradigm*, and many of the experiences and heuristics that could possibly pertain to the development of a science of automatic computation seemed equally relevant. According to Kuhn, early (pre-science) fact-gathering is nearly a random activity, usually restricted to the wealth of data that lie ready at hand (Kuhn, 1996:13-15). For instance, software engineering has been characterized as being in the pre-paradigm stage (Wernick & Hall, 2004).

> PARADIGM
>
> In his 1962 book *The Structure of Scientific Revolutions* Kuhn introduced the term *paradigm*, but he never made it very clear what the term actually means. Computer scientist and philosopher Margaret Masterman found more than twenty different ways in which Kuhn used the term (Masterman, 1970). In the subsequent versions of his book, Kuhn explicitly made a distinction between two senses of the term *paradigm*. In one sense, it refers to *a disciplinary matrix*, which means an entire constellation of beliefs, values, techniques, and so on shared by the members of a given community of researchers. In the other sense, it refers to *an exemplar*, which is a sort of a set of concrete puzzle-solutions which, when employed as models or examples, can *replace explicit rules* as a basis for the solution of the remaining puzzles of normal science (Kuhn, 1996:175,182-191). Kuhn himself came to prefer terms other than *paradigm* in his later work.

In Kuhn's theory, if one of the pre-science schools seems better than its competitors, it will slowly gain support among scientists. When science stabilizes enough, and the scientists working with the field have developed strong enough a consensus about the theories and tools of that particular science, the commonly supported constellation of theories, beliefs, values, techniques, and so on, can be said to become *normal science*. Those scientists who adopt those theories, beliefs, values, techniques, and so on, practice normal science. A mature scientific paradigm is made up of (1) general theoretical assumptions and laws and (2) the techniques for their application (Chalmers, 1976:108). The paradigm need not (and in fact never does) explain all the facts with which it can be confronted (Kuhn, 1996:16-17).

In the 1930s and 1940s automatic computing was still in a pre-science state: there were competing theories and techniques for automatic computation, and none of those theories and techniques had established superiority over the others. A number of researchers experimented on electronic circuit elements and fully-electronic computation, yet many older members of the scientific establishment defended analog and hybrid computing (Flamm, 1988:48; Campbell-Kelly & Aspray, 2004:70-83). In the first half of the 1940s, relatively young researchers in the Moore School of Electrical Engineering at the University of Pennsylvania gradually came to understand, firstly, the advantages of fully-electronic computation, and, secondly, the stored-program concept. Copies of their drafts were circulated widely, and after a few successful implementations of stored-program computers, the scientific community gradually adopted what might be called the *stored-program paradigm*. Currently the stored-program paradigm has gained enough momentum that it is, despite its limitations, largely taken as an unquestioned foundation for successful automatic computation.

It is, however, difficult to say if the stored-program paradigm suffices as a *scientific paradigm*. The conception of the stored-program paradigm was a definite shift to a technical (e.g., von Neumann-architecture) and theoretical paradigm (e.g. Church-Turing Thesis). The stored-program paradigm, however, entails only a technical model and a theoretical framework. It does not dictate forms of inference, logic of justification, modes of argumentation, practices of research, conventions for settling scientific disputes, or other aspects of a scientific paradigm. Regarding inference, logic, argumentation, or other kinds of conventions and scientific practices, computer scientists hold various views.

No matter how well a paradigm serves the purposes of science, in the course of their work, scientists every now and then come across with phenomena that their current normal science cannot explain coherently; these are, in Kuhn's theory, called *anomalies*. When many enough anomalies accumulate, scientists cannot trust their normal science anymore, and the discipline drifts into a *scientific crisis*.

According to Kuhn, during the crisis state there appears a number of competing approaches that can explain the anomalies that led science to crisis. Different new approaches can usually explain some aspects of the phenomenon well, yet ignore some other aspects of the same phenomenon. The paradigms compete for support, and at some point one of the competing paradigms wins the others, causing a *scientific revolution*. A complete *paradigm shift* happens when the opponents of the revolutionary paradigm are convinced or a new generation of scientists replaces the old one.

4.4.2 Characteristics of a Paradigm

Kuhn took a firm stand on the term "problems". Since the outcomes of research problems in normal science can be anticipated, often so accurately that what remains to be found is uninteresting per se, the *method* of achieving that outcome is often the interesting part, the unknown. "*Bringing a normal research problem to a conclusion is achieving the anticipated in a new way, and it requires the solution of all sorts of complex instrumental, conceptual, and mathematical puzzles*" (Kuhn, 1996:36). Thus, Kuhn called this sort of research activity *puzzle-solving*.

The deliberate selection of the phrase *puzzle-solving* over *problem-solving* underscores the constrained nature of normal science. For instance, Steve Fuller argued that most scientists are narrowly trained specialists who try to work entirely within their paradigm until too many unsolved puzzles accumulate (Fuller, 2003:19). Also, Kuhn emphasized the game-like characteristic of normal science:

> *On the contrary, the really pressing problems, e.g., a cure for cancer or the design of a lasting peace, are often not puzzles at all, largely because they may not have any solution. Consider the jigsaw puzzle whose pieces are selected at random from each of two different puzzle boxes. Since that problem is likely to defy (though it might not) even the most ingenious of men, it cannot serve as a test of skill in solution. In any usual sense it is not a puzzle at all. Though intrinsic value is no criterion for a puzzle, the assured existence of a solution is.*
>
> (Kuhn, 1996:36-37)

Thus, in Kuhn's theory, one of the things that acquiring a scientific paradigm brings along is criteria for choosing problems that (according to the paradigm) can be assumed to have solutions. Those problems that are not reducible to the puzzle form, may be rejected as philosophical speculation, as a concern of another discipline, or sometimes as just too problematic to be worth the effort. In computer science, rejections like those have often been a part of the scientific turf wars as well as a part of the debates over the definition of computer science as a discipline.

A paradigm can even insulate the scientific community from those societally important problems that are not reducible to puzzle form. From the Kuhnian perspective, it seems paradoxical that there is not much correspondence between the difficulty of a particular science's problem field, and the "hard" image of that science—quite the contrary. The more ambiguity there is about the premises, methodology, and the goals of a science, the more "soft" the science is, regardless of the difficulty of the problem field of that science. *Uncertainty*, a characteristic which common sense would attribute to difficult problems, is typically attributed to "soft" sciences. *Clarity*, *predictability*, and an *expected fit with existing knowledge*, which common sense would attribute to simpler problems, are characteristics of a "hard" science. The soft-hard division does not seem to tell much about the difficulty of the problems that a scientific discipline deals with.

A narrow focus on problems is not merely detrimental to science. In Kuhn's theory, normal science owes its success to the ability of scientists to regularly select problems that can be solved with conceptual and instrumental techniques close to those already existing. Researchers conducting normal science *do not aim at radical novelties of fact or theory*, and, when successful, find none (Kuhn, 1996:52). This sort of narrow focus enables the researchers to concentrate resources on well-constrained areas, which may lead to fast and deep development in that particular area. The researchers who practice normal science, engage throughout their careers in what Kuhn called "mopping-up" work (Kuhn, 1996:24).

Kuhn wrote that there are three foci of normal science, which are neither always nor permanently distinct. The first focus is the class of facts a paradigm has shown to be particularly revealing of the nature of things. The goal of the researchers investigating this class is to *expand on the phenomena with more precision and in a larger variety of situations*. For instance, computational complexity is such class of facts, and the understanding of computational complexity is continuously expanded. Nowadays the number of named complexity classes is in the tens or hundreds, depending on how one counts them.

The second focus is on those facts that can be *compared directly with predictions from the theories of the paradigm*. Improving the scientists' agreement on facts within a paradigm or finding new areas in which this agreement can be demonstrated at all, presents a constant challenge to the skill and imagination of the experimentalist and observer. Using the predictions, theories, and techniques (linear programming, heuristics, branch-and-cut) from the first class of facts has led to remarkable results in working out seemingly infeasible problems. For example, the traveling salesman's problem has been solved for the 24 978 cities in Sweden (see optimal tour of Sweden).

The third focus of normal science is on the *fact-gathering activities of science*. It consists of empirical work undertaken to articulate the paradigm theory, resolving some of its residual arguments, and permitting solutions to new problems. Examples of this kind of work are determining physical constants, finding faster algorithms to analyze graph structures, or conducting usability tests. Furthermore, fact-gathering activities may incorporate ways of applying the paradigm to a new area of interest (Kuhn, 1996:25-30). Computer science and computational models have indeed been used in an astonishing number of studies conducted in a variety of fields. There are computational models in physics and chemistry, in meteorology and biology, in economics and neuroscience; even computational models of culture and social phenomena exist.

4.4.3 Revolutions in Science

Normal science succeeds by carefully constricting the prerequisites, goals, and means of scientific research. However, new and unsuspected phenomena have been repeatedly uncovered and radical new theories have again and again been brought forth by scientists (Kuhn, 1996:52). If a scientist fails to fit observations or theories into the dominant paradigm, that is usually seen as a failure of that scientist rather than as a flaw in the paradigm. That is, if a scientist comes up with research results that do not nicely fit the theoretical-technical framework, those *results* are considered to be a failure. In falsificationism, observations that clash with the dominant view of the world are called *falsifications of a paradigm*, but in Kuhn's view, those puzzles that resist solutions are seen as *anomalies* (and single anomalies do not rock science much). So, whereas hard-core falsificationists say that anomalous results should lead to abandoning the theory, Kuhn argued that what really happens, is that anomalous results lead to questioning the *results* but not the theory.

In the Kuhnian theory, when anomalies are encountered, scientists continue to explore them, and this exploration closes only when the paradigm has been adjusted so that the unexpected (i.e. the anomaly) has become the expected. In short, until the theory has been adjusted so that the new fact (anomaly) can be frictionlessly assimilated, the new fact, according to Kuhn, is not quite a scientific fact at all (Kuhn, 1996:52-53).

The more there arise anomalies that do not fit the dominant paradigm, the more insecurity among the researchers grows. Some of the anomalies can be more or less forcefully fit into the dominant paradigm, but gradually the dominant paradigm starts to seem increasingly dubious to researchers. At some point normal science drifts into a crisis state. Although the history of science shows that researchers try to cling to normal science even when it seems flawed, at some point a better alternative is found. Kuhn wrote,

> *The transfer of allegiance from paradigm to paradigm is a conversion experience that cannot be forced. Lifelong resistance, particularly from those whose productive careers have committed them to the old tradition of normal science, is not a violation of scientific standard but an index to the nature of scientific research itself. The source of resistance is the assurance that the older paradigm will ultimately solve all its problems* [...] *That same assurance is what makes normal or puzzle solving science possible.*
>
> (Kuhn, 1996:151-152)

German physicist Max Planck (1858-1947), the father of quantum theory, worded the same thing in a more tragic form:

> *A new scientific truth does not triumph by convincing its opponents and making them see the light, but rather because its opponents eventually die, and a new generation grows up that is familiar with it.*
>
> (Planck, 1949:33-34)

Kuhn used the famous duck-rabbit picture (Kuhn, 1996:114, see Figure 9[23]) to portray the idea of how the same data can be seen in very different ways. Although nothing in the picture changes, people can ultimately learn to see the lines on the paper as a duck, as a rabbit, or both. Similar, when scientists change their theoretical perspective, they can learn to see their research results and research data very differently.

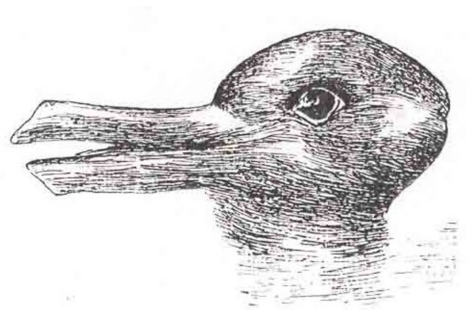

Figure 9: Duck-Rabbit Illusion

4.4.4 Problems With Kuhn's Theory

The Structure of Scientific Revolutions was one of the most influential books on the character of (natural) science in at least the second half of the 20th century, if not the entire 20th century. Interestingly, the book just provides a general account of scientific change in about 200 non-technical, very lightly referenced pages, in the manner of an extended encyclopædia entry as the book was in fact originally conceived (Fuller, 2003:18-19). Kuhn's ambiguity makes it problematic to read his work: Kuhn did not make it clear whether his account is a *descriptive* or a *normative* one. That is, it is not clear if Kuhn only describes how scientists *actually* work, or if Kuhn makes claims of how scientists *should* work.

23 Source: Jastrow, 1899. Copyright expired.

Even Popper conceded that Kuhn's normal science passes as a descriptive account of scientific practice, but Popper certainly did not agree with it as a normative account (Popper, 1970 in Lakatos & Musgrave, 1970). Feyerabend implied that Kuhn may have wanted to leave himself a second line of retreat: Those who dislike the implied derivation of values from facts can always be told that no such derivation is made and that the presentation is purely descriptive (Feyerabend, 1970). Feyerabend alluded to David Hume, who centuries ago wrote a rationalization of why one cannot tie normative claims to descriptive claims (referred to as the *is-ought problem* or *Hume's Guillotine*).

HUME'S GUILLOTINE

The *is-ought problem*, raised by David Hume (1711-1776) is one of the most fundamental questions in philosophy. Hume noted that many authors often make claims about *how things ought to be* based on *how things are* (Hume, 1739: BIII:507-521). But these two kinds of claims are from different realms altogether. Descriptions of *what is* come from observations, and their truthfulness is based on how well they correspond to reality. Prescriptions of *what ought to be* are moral statements, and they are based on something like “right desire”, and there seems to be no obvious way of judging their truthfulness. A complete severing of normative statements (what ought to be) from descriptive statements (what is) is sometimes called *Hume's Guillotine* or *Hume's Law*.

Kuhn replied to Feyerabend's accusations: “*Surely Feyerabend is right in claiming that my work repeatedly makes normative* [and descriptive] *claims*” (Kuhn, 1970). Kuhn continued, stating that the answer is that his claims should be read in “both ways at once”. Kuhn made his view clear: “*If I have a theory of how and why science works, it must necessarily have implications for the way in which scientists should behave if their enterprise is to flourish*” (Kuhn, 1970). Although this sounds quite reasonable at first hand, it is certainly dubious in the light of the is-ought problem.

Following Kuhn's logic, if it were descriptively true that unethical science flourishes, then a normative statement “unethical science is desirable” follows from that. However, Hume's Guillotine states that one cannot derive normative clauses from descriptive statements. One cannot say what is desirable science based on what kinds of science work well. Statements of “what works well” and “what is desirable” are from different realms. And the is-ought problem prevails no matter which word one uses instead of *unethical*, be it *free*, *independent*, or even *progressive.*

Critics of Kuhn identified also other weaknesses in Kuhn's theory; they wrote that Kuhn's normal science combines the worst qualities of the Mafia, a royal dynasty, and a religious order (Fuller, 2003:46; Feyerabend, 1970). Normal science, according to Kuhn, is accountable only to itself. But if scientists are part of a society, if they consume the resources of a society, and if the results of their science are catalysts of change in society, it is hard to argue that they could be accountable only to themselves. Yet, Fuller wrote, Kuhn managed to succeed simply by ignoring the issue, and that Kuhn left his readers with the impression—or perhaps misimpression—that, say, a multi-billion dollar particle accelerator is nothing more than a big scientific playpen (Fuller, 2003:46).

In addition, Feyerabend claimed that wherever one tries to make Kuhn's ideas more definite one finds that they are false. Feyerabend asked, “*Was there ever a period of normal science in the his-*

tory of thought? No—and I challenge anyone to prove the contrary." (Feyerabend, 1975. Feyerabend elucidated this point in Feyerabend, 1993). In a sense, "periods of normal science" may be an illusion of history. The further back in history one looks, the less data about the era's scientific disputes can be found. Is that a proof of greater scientific consensus or an indication of lack of data? The manner of writing histories of science as series of culmination points reinforces this picture of history of science and invention as a predetermined route through series of small revolutions. Also, it depicts history as a series of mini-fables triggered off by revolutions.