

Chapter 13: Bioinformatics and Biological Reality

Ingvar Johansson⁶³

Many bioinformaticians seem to shy away from believing that there is a mind-independent biological reality at all, or believing that we can have knowledge about such a reality. The aim of this chapter is to try to counteract this tendency, and it consists of two main parts. In the first part, I clarify three different positions in the philosophy of science with which it would be fruitful for bioinformaticians to become familiar. When they are spelled out in some detail, it becomes evident that these positions are mutually exclusive, but when seen only vaguely, the false impression may arise that one can sometimes rely on one position and sometimes on another. I label them *Myrdal's Biasism*, *Popper's Epistemological Realism*, and *Vaihinger's Fictionalism*, respectively, and I defend Popper's position.⁶⁴ In the second part of this chapter, I infuse new blood into the common semantic distinction between the *use* and *mention* of terms and concepts.⁶⁵ Both the red and the white blood corpuscles in this new fluid come from the philosophy of intentionality.

1. *Myrdal's Biasism*

Now and again we think of, and even perceive, the world in a way that is closer to how we would like it to be than how it really is. In such

⁶³ Slightly revised from the version originally published in *Journal of Biomedical Informatics* 39 (3:2006), 274-287; used with permission from Elsevier.

⁶⁴ The content of this paper has gradually come to fruition over the course of many conferences and workshops concerned with philosophy and informatics. The conference 'Ontology and Biomedical Informatics' in Rome (May 2005) finally triggered me to make these thoughts as clear as possible. Both biasism (but not Myrdal's) and Vaihinger's fictionalism were, quite independently of me, put on the agenda in Rome by Alexa McCray's talk 'Conceptualizing the World: Lessons from History'.

⁶⁵ I will deny my own preferences and use 'term' and 'concept' instead of 'word' and 'meaning', respectively, in order to conform to the usage of bioinformaticians. To a non-Platonist philosopher such as myself, the term 'concept' suggests too many allusions to entities that exist in some extratemporal realm of their own, independently of human beings. 'Meaning', on the other hand, has no such associations. Meanings exist directly only in people.

situations, we are biased. But how often does this occur? And what are the consequences of such bias for scientific research? One position in the philosophy of science can be captured by the following thesis and proposal:

- Thesis: Every conceptualization and theory is biased.
- Proposal: Admit that you are biased and make the causes of this bias (valuations, social positions and backgrounds, etc.) explicit, both to yourself and to your readers.

This position, nowadays widespread, was put forward already in the fifties by the economist Gunnar Myrdal (who shared the Nobel Prize in economics with Friedrich Hayek in 1974), but only as a thesis about conceptualizations in the social sciences (Myrdal, 1956, 1968, 1973 (Chapter 7)). Myrdal's views quickly reached the general philosophical audience thanks to Ernest Nagel's discussion – and criticism – of them in his classic *The Structure of Science*, 1961.

At the time Myrdal was writing, it was commonly assumed that scientists in their research activities ought to be, and mostly were, neutral with respect to valuations (Myrdal's term) and values that are not purely scientific. In criticism of this assumption, Myrdal claimed (a) that it is impossible for social scientists to free themselves from all such valuations, and (b) that valuations necessarily distort. According to Myrdal, since the value-neutral social scientist is a myth, social science is always more or less biased and more haunted by conflicts than the natural sciences are, and the only thing that scientists can do to become more objective is to find out and clarify, both for themselves and their readers, what kinds of valuations they bring to their research.

Looking at the historical development of the natural sciences, one can now add further that, even though there is much scientific consensus among natural scientists at most points in time, there is nonetheless a great divide between natural scientists belonging to different epochs (contrast Europe, for example, at the times of Newton and Einstein). Such differences, it can then be argued, are due not to the discovery of new facts but to the different cultural values of the centuries and scientists in question. In this way, many people have moved from Myrdal's own biasism, which is restricted to the social sciences, to the generalized version, which applies to all sciences that are not purely formal. Logic and

mathematics are mostly regarded as being outside the scope of biasism, but I have seen no claim that bioinformatics should be regarded as such.

As I will show, biasism (in whatever version) contains at least three serious philosophical flaws, each of which is sufficient reason to reject it. First, it makes no sense to speak of something being to the right if there is nothing that can be said to be to the left; similarly, it makes no sense to speak of bias if it cannot be contrasted with truth. In Myrdal's writings, biasism does not (and cannot without losing its sense) take the concept of truth wholly away. What it does do is to claim that we cannot *know* truths and that, therefore, we should speak of research results as being true-for-certain-valuations instead of being just true.

Biasism does not say that scientists are *sometimes* biased and so put forward distorted research results. The claim of biasism is that this is *always* the case; either only in the social sciences (the restricted thesis) or in all the non-formal sciences (the general thesis). Let me compare biasism in science with issues of legal jurisdiction. Judicial procedure seeks to find non-biased judges and jury members. If biasism were applied to such procedures, it would amount to the claim that there are no non-challengeable persons at all. Because of this generality, the thesis of biasism has to be applied to itself. It then implies the following disjunction: either biasism is false or it is true; but in the latter case it says of itself that it is biased and therefore false. That is, it is necessarily false. Therefore, of course, it should not be adhered to.

However, the self-referential paradox of biasism can be taken away. The defenders of biasism have merely to claim that their thesis lies outside the harmful influences of valuations and that they, therefore, are in a position to state one of the few known truths. Their thesis would then be: All theories are biased, except the theory of biasism. But now, another problem pops up. They have to explain why their thesis – a thesis which belongs to the sociology of knowledge is, in contradistinction to all other scientific and philosophical hypotheses, not influenced by valuations. If their thesis really is true, then it seems to be a mystery why this cannot also be the case with scientific assertions of other sorts. As far as I know, no one has solved this problem; I think it is unsolvable.

There are at least two reasons why many otherwise good researchers do not notice the paradoxical character of biasism. First, it seems to be natural for people who make assertions such as 'Humans are always fools' and 'Humans are always liars' to place themselves outside the scope of what they say; if not altogether, at least at the moment of making the assertion.

Those who have asserted ‘Humans are always biased’ might have followed this habit without noticing it. Second, in the case in hand, it is easy to make a so-called fallacy of composition. That is, from the fact that something is possible in *each* case, one falsely thinks that one can draw the conclusion that this something is possible in *all* cases taken collectively. Obviously, from the fact that in a marathon race *each* starting runner may win, one cannot validly draw the conclusion that *all* runners may win. Similarly, but less obviously, from the fact that *each* scientific hypothesis may be biased, one cannot by means of mere logic draw the conclusion that *all* scientific hypotheses may be biased.

A second argument against biasism is the following. The biasist proposal says that scientists should make the causes of their biases explicit; but according to the biasist thesis, even such a presentation of one’s bias must itself be biased and therefore false. Why? Because to state what has caused one’s bias is as much a hypothesis as are other empirical assertions. According to the thesis, it must be impossible to know what the true causes of one’s bias are. If biasism is true, researchers do not only automatically get a distorted view of what they study, they also get a distorted view of what has caused their distorted research results. There are, so to speak, distortions all the way down. Therefore, there is no reason to follow the proposal.

However, as in relation to the first flaw, the defenders of biasism may attempt to bypass this self-referential oddity by qualifying their position. Confronted by this second curious feature of their position, they may claim that their proposal makes good sense because researchers are *less* biased when they try to find truths about the causes of their bias, than they are when they try to find other scientific truths. For instance, it may rhetorically be asked: Isn’t it easier for an economist to find out what his sex, ethnicity, social background, and social valuations are than to find out how, in some respect, the market works? I have two counter-remarks to this rejoinder. First, the relevant problem is much harder than merely discovering facts about one’s social position and background. The real problem is to discover what causes distortions in one’s own research. In such an undertaking, one has also to take into account the fact that sometimes people with the same social position and background have different opinions. But second, and more decisive, is the fact that this qualification breaks the biasist frame. If there are degrees in the way researchers are biased, there are degrees of distortion in research results; and if there are degrees of distortion, there are degrees of being true or

false, i.e., degrees of being ‘truthlike’ (to anticipate the section on Popper’s realism below). But if there are such differences in degree of distortion between ordinary hypotheses and hypotheses about factors that cause bias, then there seems to be no reason why one should not be able also to detect such differences in truthlikeness between ordinary hypotheses.

The critique that I have presented so far takes seriously the fact that biasism puts forward an all-embracing thesis, which, in effect, replaces the notion of being true with the notion of being true-for-certain-valuations, which, in turn, can ground notions such as true-for-us and true-for-them. Such a replacement leads, as I have pointed out, to inconsistencies. This criticism, it has to be noted, by no means implies that we are never justified in talking about bias in science. In local cases, and having recourse to the notion of truth, we seem now and then to be justified in asserting that some scientists have been biased. But such local accusations of bias must be kept distinct from biasism, which contains a universal thesis.

For several decades now, biasism comes naturally to many people. One causal factor behind this fact might be the following. Nowadays, a large number of people in Western societies earn their living performing research or research-like activities in which the final research report takes the form of a consensus statement written by a group. This is true of public commissions of inquiry, be they initiated by the state or some regional or local authority; it is true of research departments in big firms; and of the managements of many research institutes. In such groups, after the research is performed, there comes a phase in which the final results are negotiated. This process can easily convey the false impression that there are no truths at all, only negotiations about truths and, therefore, only truths-for-certain-valuations, truths-for-us, truths-for-them. As far as I can see, many bioinformaticians have a similar kind of experience when they try to do justice to the advice and opinions of experts in various domains of knowledge.

Finally, let us for a while imagine that biasism has no self-referential problems. Nonetheless, another curiosity appears. All research needs a regulative idea, something that tells the researchers what to look for. Traditionally, the overarching regulative idea has been truth. This does not mean that truth has to be at the center in every phase and corner of research. For example, physicists may now and then be merely playing with possible solutions to some equations without, for the moment, bothering about *truth* at all. Similarly, some biologists may be merely playing with simulations of various biological processes; and researchers

in the humanities may be playing with certain possible interpretations of texts. This means only that there is, even in these playful situations, an indirect connection with the discovery of real natural laws or physical facts, the discovery of real biological processes, and the discovery of true interpretations, respectively. Today's science relies on a division of labor where not every part has to have a direct link to reality. This being noted, the third flaw of biasism can be stated as follows:

- Biasism wants science to get rid of the regulative idea of truth, but it has no adequate alternative to offer.

According to the biasist proposal, researchers should admit that they are biased and make the causes of their bias explicit. But what is the purpose of this proposal? Since a rational person should not seek truth if he or she firmly believes that one cannot even come closer to it because there is bias all the way down, the proposal in effect implies that researchers should exchange the regulative idea of truth for the idea of truth-relative-to-the-researcher's-valuations. The latter could be specified either as the idea that researchers should try only to promote their own long-term interests, or that they should, in the course of their research, only try to have as much fun as possible. Although such a substitution has no logical flaws, it amounts to a substitution of science with something else. It implies, contrary to Myrdal's intention – which was to promote objectivity – that researchers should be allowed to consciously deviate from data and to consciously ignore data that they suspect are problematic *because of their own valuations*. In secret, individual researchers may well have such goals, but these goals cannot possibly be made the public goal of science. Who, for instance, would fund researchers who say that they will use their research money only in order to promote their own egoistic interests or only in order to have fun?

Note that the remarks just made are not at all concerned with the question of how *research problems* within one's discipline *are chosen*. Such choices can, of course, be easily related to valuations and interests. This reminds us that several factors may account for the popularity of biasism, and that one is precisely a conflation of the choice of a regulative idea for one's research (truth, long-term subjective interest, short-term fun) with the choice of research problem. Another factor may be a neglect of the fact that all philosophically minded researchers have to face what might be called an existentialist choice: shall I primarily try to find the true solution

to my problem, or shall I primarily try to find a result that promotes my interests, or shall I primarily try to merely have fun? We can call this choice existentialist, since it is inevitably a personal choice that every researcher must make for herself. It can be made consciously, half consciously, or subconsciously, but it has to be made.

Yet another factor behind the popularity of the biasist proposal, perhaps the most influential, is the fact that there really is something to the proposal when it is viewed from the perspective of the *readers* of research reports. However, this requires that the proposal be put within a traditional truth-seeking framework – a fact which is not noted by the adherents of biasism. In what way can someone reading a research report be helped by coming to know the valuations of the researchers in question? On non-biasist premises, the answer is simple. As soon as there is a division of labor in the knowledge enterprise of a community, the sources of knowledge traditionally discussed in epistemology, namely reason and observation, are complemented by trust (in those who are providing information) (Coady, 1992; Kalman, 1999). In order for laymen to accept knowledge or information from researchers, and in order for researchers to accept knowledge or information from other researchers, the former have to trust the latter; in information science, knowledge engineers normally trust the domain experts. Therefore, the readers may be helped in this trust issue if each researcher states: ‘trust me or not; I have done my best to find the truth with ordinary methodologies, but if you suspect that I have distorted facts in order to further my interests, then I can tell you that my sex, ethnicity, social backgrounds, and social valuations are as follows: ...’ An example, for simplicity’s sake not taken from biology, will make the point more lucid.

Imagine the following situation. Two different investigations have been made about the income distribution for a certain kind of job. According to report A, the average income is 15% higher for men than that for women, but according to report B it is 25% higher. The researcher behind report A states that he is a male income statistician who thinks that men ought to have higher salaries than women, and that, in particular, a 15% difference is too little, whereas the person behind report B states that she is a female statistician who thinks that men and women ought to have the same salaries for the same kind of job, and that a 25% difference is far too much. Whose report should be trusted? Both reports cannot be true, although both can be false. In my opinion, if it is impossible to perform further independent investigations of this matter of one’s own, it is somewhat

rational to trust the person whose values (not sex) one shares. This being so, there is a kernel of truth in the biasist proposal that researchers should make their valuations, social positions and backgrounds, etc. visible, but this kernel has here been placed within a context where traditional truth seeking is taken for granted. That is, the researchers in the example have both asked themselves: ‘what is the truth, what are the facts?’, and their readers ask themselves: have the researchers really found the true income distribution? A researcher who suspects that he (or she) really unconsciously distorts facts ought to perform his (or her) investigation twice. He should first make it, so to speak, spontaneously, and he should then work through it once more with the conscious intention of trying to detect hitherto unconscious distortions.

2. Popper’s Epistemological Realism

Outside the philosophy of science, Karl Popper is most well known for his defense of democracy in *The Open Society and Its Enemies* (Popper, 1945). Within the philosophy of science he is best known for his falsifiability criterion and his advocacy of *fallibilism*. The former consists in the thesis that scientific hypotheses, but not metaphysical assertions, are falsifiable, and that, therefore, scientists (but not metaphysicians) are marked by the fact that they can state in advance what could make them regard their hypotheses as false. Fallibilism is the view that no presumed knowledge, not even scientific knowledge, is absolutely certain. In order to put his falsifiability criterion to real work, Popper connects it with some other general methodological rules. Here, however, I will present only his general epistemological realism.⁶⁶ Although I wholeheartedly accept this realism, I believe that his falsifiability criterion and its concomitant rules have to be rejected (Johansson, 1975). Thus, Popper’s general realism can be dissociated from his methodological rules, from his view that there is a gap between science and metaphysics, and that there is a criterion for detecting this gap. In particular, I will highlight a notion that is crucial to Popper’s realism. He verbalizes it in three different ways: truthlikeness, verisimilitude, and approximation to truth (Popper, 1972). I find this notion extremely important, but unduly neglected outside circles of Popper

⁶⁶ This realism is best spelled out in Popper, 1963, in particular in Chapters 1-4 and 10, and in Popper, 1972, in particular in Chapters 2, 5, and 7-9. His falsifiability criterion and most of his methodological rules are put forward in his 1959.

experts. The core of Popper's epistemological realism can be captured by the following thesis and proposal:⁶⁷

- Thesis: Every conceptualization and theory almost certainly contains some mismatch between theory and reality.
(Compare Myrdal: Every conceptualization and theory is biased.)
- Proposal: Seek truth but expect to find *truthlikeness*.
(Compare Myrdal: Make your valuations, social positions and backgrounds, etc. visible.)

Popper's epistemological realism combines fallibilism with the traditional idea that truth seeking has to be the regulative idea of science. The key to this mix is the notion of truthlikeness (verisimilitude, approximation to truth). The intuition behind this notion is easily captured. Consider the three assertions: (1) The sun is shining from a completely blue sky, (2) It is somewhat cloudy, and (3) It is raining; or at the assertions (1) There are four blood groups plus the Rh factor, (2) There are four blood groups, and (3) All blood has the same chemical composition. In either case, if the first assertion is true, then the second assertion has a higher degree of truthlikeness and approximates truth better than the third one. This is not to say that the second is *more likely to be wholly true* than that the third. The sentences 'X is *probably* true' and 'X has *probably* a high degree of truthlikeness' express relations between an assertion X and its evidence, whereas the sentences 'X is true' and 'X has a high degree of truthlikeness' express relations between the assertion X and facts (truthmakers) in the world. The former sentences express evidential relations, the latter express semantic-ontological relations;⁶⁸ the idea of truthlikeness belongs to a correspondence theory of truth.⁶⁹

⁶⁷ Of course, any epistemological realism presupposes a philosophical-ontological realism. With respect to the spatiotemporal world, Popper has a kind of level ontology (with which I wholly agree), according to which neither biological reality nor mental reality can be ontologically reduced to lower levels. Also, he thinks that thought contents have a kind of objective existence in what he calls the Third World (as contrasted with material reality, which makes up what he calls the First World and mental reality which forms the Second World).

⁶⁸ The possible conflation between being truthlike and being probably true comes more easily in some other languages. In German, for instance, the corresponding terms are 'wahrheitsähnlich' (truthlike), 'wahrscheinlich' (probable), 'Wahrheit' (truth), and 'wahr' (true).

At the end of better and better approximations to truth, there is of course truth. To introduce degrees of truthlikeness as a complement to the simple opposition between true and false is a bit – but only a bit – like switching from talking only about tall and short people to talking about the numerical or relative lengths of the same people. The difference is this. Length corresponds both to real comparative and numerical concepts, but there are no such concepts for verisimilitudes. All lengths can be linearly ordered (and thus give rise to a real comparative concept), and a general numerical distance measure can be constructed for them (which gives us a quantitative concept). Popper thought that such concepts and measures of degrees of truthlikeness could be constructed, but like many others I think that the ensuing discussion shows that this is impossible (Keuth, 2000). That is, we have only a qualitative or semi-comparative concept of truthlikeness. Some philosophers think that such a concept of truthlikeness can be of no use (*ibid.*, 198-9), but this is too rash a conclusion.

To demonstrate that even a semi-comparative concept of truthlikeness can be useful and important, I will use an analogy. We have no real comparative concept for geometrical shapes, to say nothing of a quantitative concept and measure. Nonetheless, we continue to use our qualitative concept of shape; we talk about shapes, point to shapes, and speak informally about similarities with respect to shape. Sometimes we make crude estimates of similarity with respect to shapes and are able on this basis to order a small number of shapes linearly (shape A is more like B than C, and A is more like shape C than D, etc.); we might be said to have a semi-comparative concept. In my opinion, such estimates and orderings of a small number of cases are also sufficient to ground talk of degrees of truthlikeness.

In the same way that a meter scale cannot be used before it has been calibrated in relation to something external to it, a standard meter, so the concept of truthlikeness of theories cannot be used until one has judged, for each domain in which one is working, some theory to be the most truthlike. In this judgment, the evidential relation stages a comeback. As I have said, truthlikeness informally measures the degree of correspondence with facts, not the degree of correspondence with evidence. Nonetheless,

⁶⁹ The correspondence theory of truth says that the truth of an assertion (truthbearer) consists in a relation to reality or in a correspondence with facts (truthmakers). Note that there can be no degrees of falsitylikeness; there are no non-existent facts with which an assertion can be compared. But, of course, one may use ‘being falsitylike’ as a metaphor for having a low degree of truthlikeness.

degrees of evidence must come into play when judging what shall be, so to speak, the standard meter for verisimilitude. Note that such judgments are commonplace decisions even for biasists and social constructivists. They are made every time some course book in some discipline is chosen to tell students some facts.

The notion of truthlikeness is epistemologically very important. Today's history of science tells us that it is no longer possible to believe that science progresses by adding one bit of truth to another in the way brick houses are built by laying bricks on top of each other. Whole theory edifices have often had to be revised, and new conceptualizations introduced; this sort of development will probably continue for a long time, perhaps forever. If, in this predicament, one has recourse only to the polar opposition between true and false, and is asked whether one believes that there are any true theories, be it in the history of science, in today's science, or in the science of tomorrow, then one has to answer: there are not. If, however, one has recourse to the notion of truthlikeness, then one can answer as follows.

There are, so far, no empirical theories known to be true in some absolute sense, but, on the other hand, there are not many absolutely false theories either. Most known theories in the history of science contain some degree of truthlikeness, even if only a very low degree. Today, however, some theories have what is probably a very high degree of truthlikeness. Why? Many modern inventions – and modern standardized therapies – are based on scientific theories, and it seems absurd to think that all such inventions in technology and medicine are based on theories with very low degrees of truthlikeness, to say nothing of the thought that these theories are mere fictions (see next section). Think, for instance, of travel to the moon, images from Pluto, computers, the internet, the GPS system, physiologic contraception, artificial insemination, and organ transplantations.

Let me now add a quotation from Popper in order to show how he himself summarizes his views on truthlikeness (1972, 335):

I have in these last sections merely sketched a programme of combining Tarski's theory of truth with his Calculus of Systems so as to obtain a concept of *verisimilitude* which allows us to speak, without fear of talking nonsense, of *theories which are better or worse approximations to truth*. I do not, of course, suggest that there can be a criterion for the applicability of this notion, any more than there is one for the notion of truth. But some of us (for example Einstein himself) sometimes wish to say such things as that we have reason to conjecture

that Einstein's theory of gravity is *not true*, but that it is a *better approximation to truth* than Newton's. To be able to say such things with a good conscience seems to me a major desideratum of the methodology of the natural sciences.

Just as in ethics there are people who only think in terms of white or black and who always want to avoid nuance and complication, so in science there are people who simply like to think only in terms of true or false. Not many decades ago scientists thought of their research only in terms of being certainly true; today, having familiarized themselves with the history of science, many – and especially in domains like informatics – think of it only in terms of being certainly false or certainly fictional (see next section). In neither of these positions – being certain that one has truth on one's side, or laying no claims to truth at all – must researchers fear criticism; but on fallibilist premises they must once again learn to do so.

Applying the notion of truthlikeness to the history and future of science allows us to think of scientific achievements in the way engineers think of technological achievements. If a machine functions badly, engineers should try to improve it or invent a new and better machine; if a scientific theory has many theoretical problems and empirical anomalies, scientists should try to modify it or create a new and more truthlike theory. As in engineering it is no sin to invent imperfect devices, so in science it is no sin to create theories that turn out not to be true. Rather, the sin in both cases is in not trying to improve on problematic machines and theories. Also, and for everybody, it is of course better to use existing technological devices than to wait for tomorrow's, and it is better to trust existing truthlike theories than to wait for the science of tomorrow.

Most rules have exceptions. Perhaps bioinformaticians, unlike scientists in other disciplines, need not bother about the history of science or think through the conflict between Popperian fallibilism and biasism? Isn't it enough for bioinformatics simply to systematize what the present-day experts in biology tell them? No, it is not. Biological knowledge grows rapidly, and even a young discipline like bioinformatics will no doubt soon have to revise some of its achievements in light of new biological knowledge. In the Gene Ontology, this is taking place before our eyes. For example, the constructors of GO list molecular functions as terms which have *obsoleted*.

3. *Vaihinger's Fictionalism*

In the 1920s and 1930s, Hans Vaihinger's book *The Philosophy of As-If* (1924) enjoyed much success. Viewed from one side this book speaks to the general positivist trend of those times; viewed from another side however it also fits well with the social constructivist trend of more recent decades. The essence of Vaihinger's position is:

- Thesis: Absolute truth, if such there is, is not attainable.
(Compare Popper: There is absolute truth, but it is probably not attainable.)
- Proposal: Regard your theories as referring to fictions; don't concern yourself with truth and falsehood.
(Compare Popper: Regard your empirical theories as referring to the world; try to find out if they are false.)

Vaihinger holds that there is only one kind of real entity, the contents of our sensations (this is the positivist side of his thinking). Things and persons in the ordinary sense, matter and energy as spoken of in physics, and things in themselves as postulated by some philosophers, are all merely fictions. Nonetheless, there are reasons to live *as if* many of such entities are real; the expression 'live as-if X exists' at the heart of Vaihinger's philosophy should be understood as follows:

- If there were Xs and we knew it, then we would have to expect some specific things to happen, and, also, we would have to act in some specific ways. In fact, however, we know that there are no Xs. Nonetheless we ought to create expectations and act *as if* there were Xs.

In some parts of his book, Vaihinger makes clear distinctions between (i) hypotheses (which are directed towards reality and demand verification), (ii) semi-fictions (which abstract away some known features of an entity, as for example the irrationality of humans is abstracted away in the concept of *homo economicus*), and (iii) pure fictions (which are based on no abstraction of this sort); but in the end he turns everything (except the contents of sensations) into pure fictions and says (1924, 108):

we are able ultimately to demonstrate that what we generally call truth, namely a conceptual world coinciding with the external world, *is merely the most expedient*

error. [...] So-called agreement with reality must finally be abandoned as a criterion.

He stresses the importance and necessity of postulating fictions in all areas of life, practical, scientific, as well as ethical. Like contemporary Anglo-American social constructivists, he implicitly takes for granted that we can communicate with each other about such fictions, i.e., he implicitly regards communication as real.⁷⁰ Since the contents of sensations play a subordinate role in his philosophy, it is no accident that his ideas can be summarized in such a way that they become, as here, lumped together with those of present-day social constructivists.

It is interesting to note how similar Vaihinger's and Popper's theses are and, despite this, how dissimilar their proposals are. In my opinion, the small difference between their theses is of no importance at all. Even if Vaihinger had subscribed to the view that there is some low probability that absolute truth is attainable, I am sure that he would have put forward the same fictionalist proposal. Conversely for Popper, even if he had thought that absolute truth is in principle unattainable, he would still have put forward the same falsificationist proposal. What, then, makes Vaihinger and Popper differ so radically in their proposals? My answer is, in short: Vaihinger's lack of the notion of truthlikeness.

False and fictional assertions are in one respect different and in another similar. They are different in that it is possible to tell a lie with a false assertion but not with a fictional one. When we are lying we are presenting as true an assertion that is false, but fictional assertions are beyond the ordinary true-false dimension. The two are similar in that neither refers to anything in reality that corresponds exactly to the assertion in question. A false empirical assertion lacks a truthmaker as a matter of fact; a fictional assertion cannot possibly have one. Therefore, it is easy to confuse the view that all theories are false with the view that all theories are about fictions. Nonetheless, it is astonishing how easily Vaihinger goes from falsehood to fictions. Why does he not believe that there can be degrees of fictionality? The less that has been abstracted away in a semi-fiction, the closer an assertion about it is to a hypotheses, and the more that has been abstracted away, the closer an assertion about it is to a purely fictional assertion. Assertions about semi-fictions might be said to be semi-true, and

⁷⁰ It should be noted that some French post-structuralists, e.g., Derrida, even regard the idea of communication as a fictional idea, and they communicate this thesis in many books.

since being semi-true takes degrees, we have hereby simply created another name for truthlikeness.

If one assertion is more truthlike than another, then it is by definition also less false. However, this falsity content (to take an expression from Popper) can easily be turned into fictionality content, whereupon a more truthlike assertion can also be said to be a less fictional assertion. When we are reading about, say, Sherlock Holmes, we have no difficulty in placing this fictional character in a real setting, London between 1881 and 1904. Not everything is fictional in many works of fiction, and we often have no difficulty in apprehending mixtures of real and fictional reference. Something similar is true when one reads about the history of science. For example, when I read about the false hypothesis that there is a planet Vulcan between Mercury and the Sun, which might explain some seeming falsifications of Newtonian mechanics, then I had no problem in taking Vulcan to be a fictional entity postulated as existing in the real solar system in about the same way as I take Holmes to be a fictional character in a real London. When I read about the false hypothesis that there is a chemical substance, phlogiston, which is exuded by things that are burning (where in truth, as we now know, oxygen interacts with the things in question), then I have no problem in taking phlogiston to be a putative fictional substance in the world of real burnings. When I read about Galen's view that (what we call) the arterial system contains *pneuma* or *spiritus*, then I have no problem in taking this *pneuma* to be fictional, but the arterial system to be real. Those who write about the history of science often make the reader look upon statements which were once false assertions as being assertions about fictions. In retrospect, we should look upon superseded theories as having mixed reality and fiction in something like the way reality and fiction can be mixed in novels. This is to give fictions their due place in science, but such local uses of fictions must be kept distinct from fictionalism, which contains a universal thesis.

I will end this section on fictionalism with the kind of remark with which I started my criticism of biasism. Apart from all other curiosities, fictionalism is self-referentially inconsistent. Fictions are created, but if everything apart from the contents of our sensations is a fiction, then there is nothing except such content that can create the fictions. However, contents of sensation do not have such a capacity. Unfortunately, Vaihinger and most fictionalists do not see the need for this kind of self-reflection.

4. *Use and Mention: In the Light of an Optical Metaphor*

When we look at things such as stones, trees, and walls, we cannot see what is on the other side. But things like water and glass are such that we can look through them to the other side. In the case of spectacles, microscopes, and telescopes, this feature is extremely useful. By *looking through* their lenses, we are able better to *look at* something else. This being-aware-of-x-through-y phenomenon is not restricted to the visual sense. It can be found in the tactile realm as well. One can grip a tool and feel the tool against the palm, but when one is very good at using such a tool, this feeling disappears. Instead, one is primarily aware of that in the world which the tool affects. For instance, when one is painting, say, a wall with a brush, one is only (if at all) indirectly aware of one's grip on the brush, and is primarily aware only of what one is painting. One is *feeling* the surface of the wall *through* the brush. What glasses are for people with bad sight, the white cane is for blind people.

Speech acts, acts of listening, writing, and reading acts – in short, language acts – are, like glasses and the white cane, tools for improving our everyday life. They can be used to convey and receive information, to give and take orders, to make and apprehend emotional outbursts, and to do very many other things. Even though language acts do not have the same robust material character as tools have, they nonetheless display the same feature of being able to be both looked at and looked through. In the former case, one is directly aware of the linguistic entities as linguistic entities, but in the latter case one is aware of them at most indirectly. When, for example, one is conveying or receiving information in a language in which one is able to make and understand language acts spontaneously, one is not looking at the terms, concepts, statements, and propositions in question; nor is one looking at grammar and dialects. Rather, one looks through these linguistic entities in order to see the information (facts, reality, or objects) in question. We are looking at linguistic entities, in contrast, when for example we are reading dictionaries or examining terminologies. If I say: 'look, the cat has fallen asleep', I want my use of the term 'cat' to be transparent and to help the person I am addressing to get information about a state of affairs in the world. But if I say 'In WordNet, the noun 'cat' has 8 senses', then I want someone to look at the term 'cat'.

My distinction between looking *at* and looking *through* is similar to the traditional distinction between the *use* and *mention* of linguistic entities,⁷¹ and it applies both to factual talk and to reading novels. In fictional discourse, terms are *used* as much as they are in talk about real things, but they are used in a very special way. Fictional discourse is *about* fictional characters; it is not about terms and concepts. In fact, we are typically using the same terms and concepts in both fictional and factual discourse.

When one is not using lenses, one can look at them and investigate them as material objects in the world. One can, for instance, try to find out what their physical properties and internal structures are like. In the world of practice, one makes such investigations of tools only when they are not functioning properly and are in need of repairing. Something similar holds true of terms and concepts. Only when our language acts are not functioning well – think for instance of learning a new language – do we normally bother to look *at* terms and concepts in dictionaries.

Furthermore, we are able to switch quickly between looking through and looking at things. Car drivers should look through, not at, the windshield, but when driving they should also have the ability to take a very fast look *at* their windshield in order to see whether, for instance, it has been damaged by a stone. Something similar is true of people using a foreign-language dictionary. They should be able to take a look at a certain foreign term and then immediately start to look through it by using it. Let me summarize:

1. In the same way that we can both look at and look through many material things, we can both look at and look through many linguistic entities.
2. In the same way that we can quickly switch between looking at and looking through glass, we can also quickly switch between looking at and looking through linguistic entities.

And let me then continue the analogy by adding still another similarity:

3. In the same way that consciously invented material devices for being-aware-of-*x*-through-*y*, such as microscopes and telescopes, have provided new information about the world, so consciously

⁷¹ I do not regard the distinction between use and mention as the same distinction as that between object language and meta-language. The use-mention distinction does not split ordinary language into distinct levels.

invented linguistic devices for being-aware-of- x -through- y , such as scientific concepts, have provided new information about the world.

By means of the *invention* of new concepts, we can sometimes *discover* hitherto completely unnoticed facts. Often, we (rightly) regard discoveries and inventions as wholly distinct affairs. Some things, such as stones, can only be discovered, not invented; others, such as bicycles, seem only to be inventions. One person might invent and build a new kind of house, and other persons may later discover it; but the first person cannot both invent and discover it. These differences between inventing and discovering notwithstanding, devices for being-aware-of- x -through- y present an intimate connection between invention and discovery. By means of new inventions of the being-aware-of- x -through- y type, we can discover x . There are many x 's that we can discover only in this way.

The third point above should be understood partly in terms of the notion of truthlikeness: if an existing conceptual system is faced with a conflicting conceptual system which has a higher degree of truthlikeness, the latter should supersede the former. But, conversely, the notion of truthlikeness should also be understood by means of the distinction between looking at and looking through. I introduced the idea of truthlikeness with the three assertions 'The sun is shining from a completely blue sky', 'It is somewhat cloudy', 'It is raining', and I said that, *given that the first assertion is true*, the second one seems intuitively to be more truthlike than the third. A standard objection to such a thesis is that this sort of comparison can show us nothing relevant for a correspondence theory of truth, since what we are comparing are merely linguistic entities (assertions). However, this objection overlooks the distinction between looking at and looking through. Looking at the assertions allows us to see only similarity relations between the assertions themselves; but when we have learned to switch from looking at to looking through such assertions – at the reality beyond – then we can coherently claim that the second is more truthlike than the third.

In the same way that our choice of lens may determine what we are able to see, so too, our choice of concepts determines what we can grasp. However, this is no objection to the thesis of epistemological realism to the effect that we have knowledge about the world: it does not render truth a social construction. When, through a concept, we look at and grasp something in the world, this concept often (i) *selects* an aspect of the world, (ii) *selects* a granularity level (for instance, microscopic or

macroscopic), and (iii) *creates* boundaries where there are no pre-given natural boundaries. The concept nonetheless (iv) *does not create* this aspect, this granularity level, or what is bounded. Think of the concept *heart*. It selects a biological aspect of the human body; it selects a macroscopic granularity level; and it creates a boundary line between the heart and its surroundings which does not everywhere track physical discontinuities as for example where the heart meets the aorta and the veins (Smith, 2001). But, nonetheless, our invention of the concept *heart* does not create our hearts, and there were hearts many millions of years before there were concepts.

5. *The Fallacy of Mixing Use and Mention*

All ontologies in information science contain terms. The builders of such ontologies look mainly *at* the terms in question, whereas the users of ontologies look mainly *through* them. Like the users, the experts in the various specialized domains of knowledge generally look through the terms. However, an ontology such as WordNet presents a special case, for (if it is to be called an ontology at all) it is an ontology of terms and meaning; it is like a dictionary, not like a taxonomical textbook. In its treatment of the term ‘cat’, WordNet begins as follows:

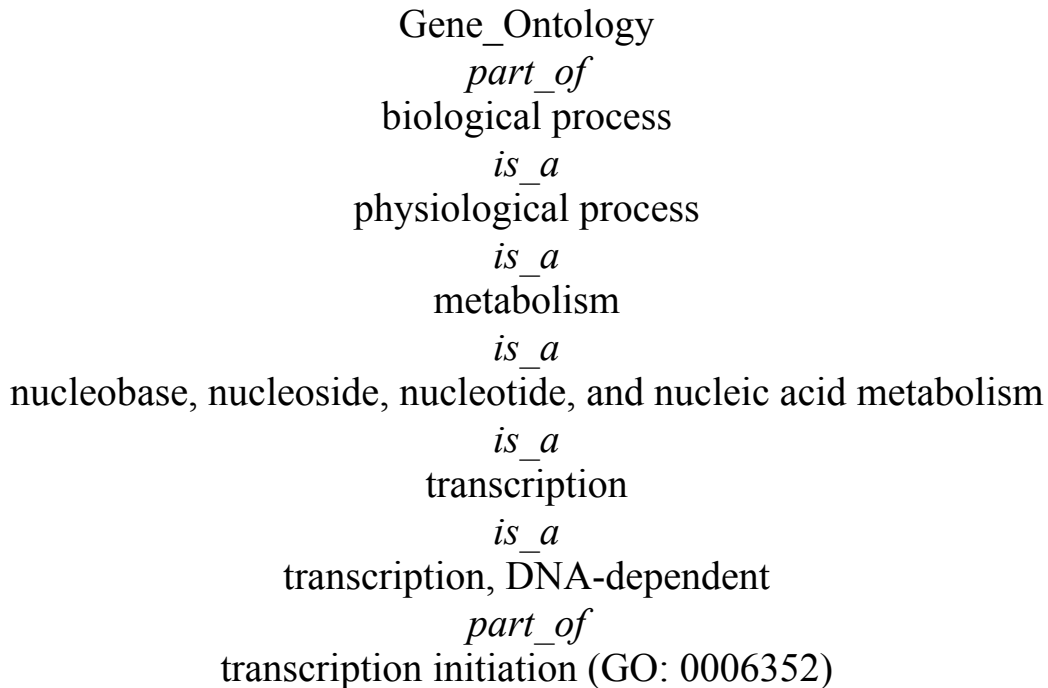
The **noun** ‘cat’ has 8 senses in WordNet.

1. **cat**, true cat – (feline mammal usually having thick soft fur and being unable to roar; domestic cats; wildcats)
2. **guy**, **cat**, **hombre**, **bozo** – (an informal term for a youth or man; ‘a nice guy’; ‘the guy’s only doing it for some doll’) (*WordNet*)

It is doubly clear that the term ‘cat’ is mentioned and not used in WordNet. Both the scare quotes around the term ‘cat’ and the fact that it is preceded by the term ‘noun’ make it clear that WordNet contains no talk of real cats; both scare quotes and context are able to disambiguate between use and mention. Here, therefore, matters are clear. In many biomedical ontologies, however, use and mention are systematically confused.

The Gene Ontology Consortium asserts that ‘[t]he Goal of the Consortium is to produce a structured, precisely defined, common, controlled vocabulary *for describing* [italics added] the roles of genes and gene products in any organism’ (Gene Ontology Consortium, 2000). That is, it is not an ontology for looking *at* terms but for looking through terms.

GO consists of three different ontologies, one for cellular components, one for molecular functions, and one for biological processes. One graph in the latter ontology (as it looks when this was originally written in June 2005 – the problem has since been corrected) can be reproduced as in the figure below; it is to be read from bottom to top. The original graph contains arrows representing the subsumption (*is_a*) and part-whole relations (*part_of*):



When a user of the GO reads this he is, I am sure, looking through the terms. That is, he reads it as signifying something like: ‘each transcription initiation is part of a DNA-dependent transcription, which is a kind of transcription, which is a special kind (nucleobase, etc.) of metabolism, which, like all metabolisms, is a physiological and biological process’. So far so good, but I have stopped at ‘biological process’. What about the last step? Reading it in the same way would yield: ‘Each biological process is part of the Gene Ontology’. But this is obviously false. It should instead be read: ‘the *term* biological process’ is part of the Gene Ontology’s hierarchy of *terms*. Thus use and mention of ‘biological process’ are here mixed. When one reads the ontology from the bottom up and arrives at ‘biological process’, this term should be regarded as *used*, but when one continues reading upwards, it should be regarded as *mentioned*.

Since, as I said earlier, most people are able to switch between looking through and looking at terms, the human users of the GO may perhaps

without noticing do so also here, and no harm is done. However, automated information-extracting systems are not able to make such switches. Obviously, GO would be a better construction without this mixture of use and mention. As the graph stands, it allows a fallacious inference to the effect that, if something is a biological process then it is part of a certain human artifact called the Gene Ontology. This might be called the fallacy of mixing use and mention. The solution is easy: let the graph end in ‘biological process’.

The same kind of fallacy appears as well (at least in June 2005 – this problem has still not been fixed) in the CRISP (Computer Retrieval of Information on Scientific Projects Thesaurus).⁷² There, one finds subsumption relations which can be represented as in the hierarchy below (to be read from the bottom upwards):

```

immunology
  is_a
  antigen
    is_a
    allergen
      is_a
      airborne allergen
        is_a
        pollen

```

Here, ‘antigen’ should be *used* in relation to ‘allergen’ (‘Each allergen is an antigen’), but *mentioned* in relation to ‘immunology’ (‘The *term* antigen is an immunological *term*’). ‘Allergen’ is a term among other terms in the field of immunology, whereas allergens themselves are among the entities that immunology studies.

The Health Level 7 Reference Information Model (HL7 RIM), also, conflates use and mention, with the unfortunate result that the users of the RIM are told by its authors that the RIM *cuts them off from the world*:

Act as statements or speech-acts are the only representation of real world facts or processes in the HL7 RIM. *The truth about the real world is constructed through a combination (and arbitration) of such attributed statements only, and there is no class in the RIM whose objects represent ‘objective state of affairs’ or ‘real processes’ independent from attributed statements. As such, there is no distinction*

⁷² <http://crisp.cit.nih.gov/Thesaurus/index.htm>. Accessed February 4, 2008.

between an activity and its documentation [italics added]. Every Act includes both to varying degrees. For example, a factual statement made about recent (but past) activities, authored (and signed) by the performer of such activities, is commonly known as a procedure report or original documentations (e.g., surgical procedure report, clinic note etc.) (HL7 RIM).⁷³

6. Use and Mention: In the Light of a Good Philosophy of Intentionality

As I pointed out in section 2.1, both terms (words) and concepts (meanings), as they are most commonly used, are invisible, since we most often look *through* them at the entities to which they refer. The need for a distinction between term and concept arises as soon as we discover a synonymy, be it between two terms in the same language or in different languages. For we then have to specify what makes the terms differ and what makes them similar in meaning (i.e., synonymous). The terms differ because they are constituted by different syntactic unities such as letters or words conceived of as purely graphical or acoustic patterns, and they are synonymous (as we say) because they express the same concept. A term is a fusion of a syntactic unity and a concept.⁷⁴ One looks *through* the concept, not through the syntactic unity, i.e., concepts are to terms what lenses are to glasses, microscopes, and telescopes.

The optical metaphor of looking through concepts is sustained by a certain approach in the philosophy of intentionality. The term ‘intentionality’ was introduced into contemporary philosophy by Franz Brentano in the nineteenth century. It refers to phenomena such as perceiving, thinking, reading, and desiring. Intentional phenomena have in common the feature that they contain a *directedness* towards something. Mostly, it is a directedness that originates in a person who is in a so-called intentional state, or who performs an intentional act, towards something else. There are, however, different opinions on how to analyze intentional

⁷³ The quoted statement and others in the documentation of the HL7 RIM are criticized in Vizenor, 2004.

⁷⁴ Those who are amenable to Ferdinand de Saussure’s linguistics can read the last sentence as follows: A sign is a fusion of a signifier and what is signified. Let me add that Saussure consciously abstracted away from his studies all questions concerning looking-through at referents. Some of his present-day followers, however, seem to take the position (criticized in this paper) that there simply are no referents.

phenomena. In my opinion, Edmund Husserl (1970) and John Searle (1983) have come the closest to the truth.⁷⁵ Let me quote Searle (18-19):

it is at least misleading, if not simply a mistake, to say that a belief, for example, is a two-term *relation* between a believer and a proposition. An analogous mistake would be to say that a statement is a two-term relation between a speaker and a proposition. One should say rather that a proposition is not the *object* of a statement or belief but rather its *content*. The content of the statement or belief that de Gaulle was French is the proposition that the de Gaulle was French. But that proposition is not what the statement or belief is about or is directed at. No, the statement or belief is about de Gaulle ...

Intentional phenomena are marked by a tripartition between (intentional) act, (intentional) content, and (intentional) object. Assume that you are reading a physician's report about your heart, which tells you that your heart has some specific features. At a particular moment, there is then your reading *act* and what you are reading about, the intentional *object*, i.e., your heart and its properties. But since your heart exists outside of your reading act, there must be something within the act itself in virtue of which you are directed towards your heart and its properties. This something is called the *content*; in assertions, it consists of propositions.

According to many non-Husserlian and non-Searlean analyses of intentionality, you are in your act of reading directed only towards the proposition, but then there is outside your awareness also a relation of representation between the proposition (the content) and the object (your heart). According to Husserl and Searle, on the other hand, you are, while, reading directed towards your heart (object) by means of the proposition (content). The first kind of analysis leaves no room for any sensible talk of looking through concepts and propositions, but Husserl's and Searle's analyses do. Though Husserl's and Searle's theoretical frameworks differ in other respects, both of them make it reasonable to believe that the metaphorical distinction between looking at and looking through concepts can be embedded within a truly theoretical framework.

⁷⁵ In this respect see Searle 1983, p. 18-9, 57-61, 97, and Husserl 1970, Investigation V, §11 and the appendix to §21 ('Critique of the 'image-theory' and of the doctrine of the 'immanent' objects of acts'). Despite later changes of opinion, Husserl retains his belief in the feature of intentionality that I will highlight.

7. *The Last Word and the Last Word but One*

In the first part of this paper I advocated Popper's realism, in particular his notion of truthlikeness. In the second part I advocated a Husserl-Searlean analysis of intentionality, in particular the view that in assertions one is directed towards the world by looking through terms and concepts. Now, in order to forestall the possible criticism that I cannot explain and make sense of the position from which I am talking, I want to bring in another prominent thinker, Thomas Nagel. I regard myself as speaking from the kind of naturalist, rationalist position that Nagel has tried to work out in his *The View from Nowhere* (1986) and *The Last Word* (1997). Below are two quotations. The first is from the introduction to the latter book, and the second is its concluding paragraph.

The relativistic qualifier – 'for me' or 'for us' – has become almost a reflex, and with some vaguely philosophical support, it is often generalized into an interpretation of most deep disagreements of belief or method as due to different frames of reference, forms of thought or practice, or forms of life, between which there is no objective way of judging but only a contest for power. (The idea that everything is 'constructed' belongs to the same family.) Since all justifications come to an end with what the people who accept them find acceptable and not in need of further justification, no conclusion, it is thought, can claim validity beyond the community whose acceptance validates it.

The idea of reason, by contrast, refers to nonlocal and nonrelative methods of justification – methods that distinguish universally legitimate from illegitimate inferences and that aim at reaching the truth in a nonrelative sense. Those methods may fail, but that is their aim, and rational justification, even if they come to an end somewhere, cannot end with the qualifier 'for me' if they are to make that claim (1997, 4-5).

Once we enter the world for our temporary stay in it, there is no alternative but to try to decide what to believe and how to live, and the only way to do that is by trying to decide what is the case and what is right. Even if we distance ourselves from some of our thoughts and impulses, and regard them from the outside, the process of trying to place ourselves in the world leads eventually to thoughts that we cannot think of as merely 'ours'. If we think at all, we must think of ourselves, individually and collectively, as submitting to the order of reasons rather than creating it (*ibid.*, 143).

Reason, Nagel says, has to have the last word. However, this statement needs to be qualified. As the logician Per Lindström notes with regard to Nagel's book: 'reason has the last word – or perhaps only the last but one, since reality, reason tells us, has always the absolutely last word'

(2001, 3-6). Not only for biologists, but even for bioinformaticians, biological reality has the last word; notwithstanding the fact that bioinformaticians need not consult it too often. Mostly, they can trust the domain experts who provide them with their information.

