4

A Manifesto for Android Epistemology:

an essay on the limits of human cognition

PASS

I read and commented on an earlier version not so different from this one.

Introduction

It has long been the case that human sensory apparati are outperformed by various kinds of mechanical devices. Perhaps there was a time when these devices could not be compared with human sensory apparati, because they could not function independently of human sensory apparati. The early telescopes and microscopes did not "sense" anything; but merely presented images to persons' sight which would not be presented without these devices. However, nowadays most instruments of scientific observation and measurement produce data which may be easily transformed to any informational medium. The devices which encode these data in forms accessible to persons tend to be physically separable from the devices which produce the data; even when this is not actually true, the portion of the device devoted to inscription tends to be based on wholly different physical principles than the portion which produces data. Insofar as data production is generally independent of inscription techniques, the devices which produce data may be said, in a very real way, to sense. When we claim that machines "sense" we claim that they can make many of the same distinctions that we do, in the same circumstances. Machines can, in addition, make many distinctions which we are unable to make.

The time is also at hand, or very soon at hand, when machines can outperform humans cognitively. In order to claim that a machine can outperform a person cognitively (or otherwise) we must specify what it is that is trying to be achieved, by either party. We can do this very easily -- the goal relative to which I should like to compare the cognitive performance of machines and persons is the prediction and control of the outputs of data-producing systems.

If we are to be metaphysical realists[1] we might consider the physical world to be a data-producing system, otherwise we might be more restrained about what we consider to be data-producing systems. It makes little difference, really. At very least, we can hardly help but consider the many more of less familiar devices in modern laboratories to be data-producing systems. Again, whether these devices are assumed to show facts about a world external to them makes little difference. An ideology of science has it that scientists are trying to predict and control the outputs of modern laboratory machines (for whatever reason). We have no reason not to take this ideology at face value; and if it is true then scientists will surely welcome the machines which can better predict and control the outputs of laboratory devices than they themselves can. If it turns out that the purpose of science is really other than as the ideology claims, I am not concerned. My only purpose is doing as philosophers of science normally do: telling scientists how better to carry out their chosen goals.

Human Cognitive Limits

The general motivation behind my belief that computers can cognitively outperform persons is that persons have fairly sharply limited cognitive abilities. Some of these are pointed out by David Faust, in his exemplary book, The Limits of Scientific Reasoning[2]. We do not need to cite the empirical studies which Faust discusses, for we can find some pretty clear limitations from our armchairs. It seems clear that we cannot conceptualize the meaning of things nested to depths of more than seven (and this is a distant upper bound). Some examples are nested intentions, meaningfully dependent clauses, or quantifiers in first-order logic. Grice[3] and others have maintained that our communication involves nested intentions of depth three or four, but no more[4].

In the case of logic, our inability to comprehend quantifiers of deep nesting is significant. It has been said that mathematics only goes seven quantifiers deep, though one rarely encounters a theorem with more than five or six nested quantifiers. Beyond this we are unable to give any particular significance to a sentence[5]. There are, however, infinitely many theorems with more than seven nested quantifiers. In general, these all have correlaries with fewer than seven nested quantifiers, many of which are of independent mathematical interest. With the formalization of logic by Frege, Tarski[6] and others we can easily recognize the soundness of a proof involving such a deeply nested sentence; but it is unlikely that we will often (or ever) think of such proofs.

Conceptualizing deeply nested expressions is certainly not the only area where human cognition is limited, but it does allow a simple contrast with the "cognitive" capacities of machines. I do not have any clear idea of why a machine would manipulated nested intentions, or dependent clauses, but there are many efforts to get machines to produce formal proofs in first-order logic. It is an elementary fact of first-order logic that no proof necessarily entails using a step with more deeply nested quantifiers than the theorem being proven (at least in familiar proof systems). However, it might be that a heuristic method of arriving at a proof involves using formulas with more quantifiers than the theorem in question. In general, there is no effective[7] way of transforming a proof using deeply nested quantifiers into one using less deeply nested ones, even where it is known that some such proof exists. The point of this is that persons could follow such a heuristic method at best very poorly. The only time that persons can do math or logic particularly well is when they can give a pretty good semantic conceptualization to the formulae with which they are dealing. It is, of course, in principle possible to have humans perform the very same syntactic transformations which machines do; but they must do this millions or billions of times more slowly and with thousands of times the likelihood of error.

Much the same situation exists in science as in logic. That is, some types of formulae are too complex for persons to conceptualize. For example, imagine some data-producing system which has twenty dependent indicators[8] and twenty independent indicators; and the behavior of these indicators could be well described by a system of twenty equations. A person cannot conceptually grasp the "meaning" of the whole system of twenty equations, except in special cases. The only cases where a person can conceptualize the whole system of equations are ones where each equation is of the form:

$D_i = F_i(I_1, \ldots, I_n)$;

for D_i a dependent variable indexed by i, F_i a function indexed by i, and for n the number of independent variables and $m < n+1$, I_m is an independent variable.

For comparison, the general case of a system of equations is of the form:

$F_{yi}(D_1, \ldots, D_k) = F_{xi}(I_1, \ldots, I_n)$;

where for n the number of independent variables and $m < n+1$, I_m is an independent variable; for k the number of dependent variables and $m < k+1$, D_m is a dependent variable; for F's functions indexed by i, and by 'x' or 'y' (i.e. independent or dependent).

If one does not believe what I have said about human inability to conceptualize systems of equations for forty (twenty dependent, twenty independent) variables, then the example can be expanded to one hundred variables, or one thousand. In certain cases there are known rules for manipulating the systems of equations described above (e.g. when the functions are polynomial), and of course humans can follow these rules. There are two things to notice here. The first is that machines can generally perform these syntactic manipulations of equations much faster and much better than humans. The more important thing to notice is that it is incredibly unlikely that a person would ever develop such a system of equations, no matter how well it describes the data or how badly simpler systems of equations do. Again, a person could in principle run through the same syntactic manipulations of the data by which a machine could arrive at the "correct" system of equations. However this amount of calculation would typically take a person a whole lifetime (or several) to perform.

Complexity and "The World"

It may be that scientists simply never encounter the degree of complexity which I have described, in their normal work. I could, of course, artificially make up a series of data which demanded a complex system of equations to correctly describe. However, there is no reason to think that any scientists would be concerned with this aspect of my data-productive activity. If, in fact, the data-producing systems of concern to scientists are describable in terms of the types of formulas which humans can conceptually grasp, then two possible explanations present themselves.

The first is that scientists have chosen to only concern themselves with systems which they might have some chance of predicting and controlling. The second is that the systems which scientists happen to be concerned with for reasons other than these systems' complexity also happen to be simple ones. Either of these two explanations could serve if the data-producing systems of concern to scientists were simply describable. However, I think it unlikely that the systems which scientists actually do concern themselves with are simply describable ones. I do not have any metaphysical and characteristically philosophical arguments for the complexity of particular systems, so my opinion does not come out of reflective certainty. However, I can point to several familiar systems where possibly complex relations of data are specifically not studied, and perhaps where complex relations exist.

One clear case where scientists study only simple relations on possibly complex data-producing systems is in most statistical social science -- in particular, sociolinguistics. A familiar type of study in sociolinguistics is one in which the relative frequency of two items occurring (in some sense) in alternation with one another is measured. An attempt is then made to correlate, for each individual or situation studied, the relative frequency of these two items with some other linguistic or extra-linguistic factors. A pioneering study of this sort is Labov's[9] result that the frequency with which clerks in department stores in NYC use /r/ versus /-/ in words in which they occur in alternation depends on the "status" of the store. Studies which are methodologically very similar occur throughout the social sciences.

It is presumed in these studies that the data-producing systems of interest are linguistic communities. However, what is pointedly not studied is *absolute* frequency of items. In addition, these studies limit themselves to a small number of extra-linguistic indicators. In this area we should suspect that a large number of indicators are significant. It may well be that many extra-linguistic factors which could be shown to correlate with frequencies of linguistic items in complex studies cannot be revealed by simple studies. This might happen, for example, if a positive value on one indicator (assuming it is Boolean) overrides the effect of a second factor. If we only included the second factor in our study we would get enough "noise" from the first factor as to make it unlikely we would detect the effect of the second factor.[10]

We encounter similar problems with the linguistic items we consider. It may be that we do not have a simple system of patterned variation between two items, but rather (for example) an alternation between two (or more) linguistics patterns, each of which contain item variations. If the choice of pattern overrides or changes the system of item variation then we would have to include the several patterns as well as the items which vary in every pattern to avoid "noise" from factors not included. An example of a complexity greater that item variation is shown by Erica

Garcia[11], as regards some aspects of the pronoun system of Spanish. This complexity is of a very limited sort, even so.[12]

The methodology I should like to see used will often reveal the sorts of complexity I mention in the above paragraphs, as well as many other sorts which are difficult to describe or conceptualize. The remarks I make will be specifically directed to a way of conducting sociolinguistic studies. However, they may very easily be extended to any area of social science with numerically codifiable indicators. The methodology I will suggest involves a wanton disregard for the conceptualizations we have about out subject matter. This is important in that many of our conceptualizations grow out of ideological reification of our cognitive bumbling.

My method of conducting a sociolinguistic study is as follows. Collect recorded speech and measure the absolute frequency of as many items as possible: phonetic, semantic, syntactic, et al. In addition, get as much information as possible about extra-linguistic factors surrounding the speech. This body of data will be huge, unseemly, and unanalyzable by persons. A machine, however, can conduct pairwise, three-wise, etc. correlation checks on all variables; and proceed in heuristic fashion to conduct correlation checks on various linear and non-linear combinations of variables. Importantly, we humans do conduct checks only of those correlations intuitively sensible to us. It is otherwise with a machine. Any patterns found by a machine in this manner, as always, require cross-validation.

There are several things I would expect to find by the sort of study described. Most generally, I would expect to find correlations among linear and non-linear combinations of data. These are correlations unlikely to be discovered or conceptualizable by persons. When persons analyze data they generally only look for correlations among pairs of variables, or pairs of additive combinations of variables. This probably often underestimates the complexity of the data-producing systems of interest. Another pattern which I would expect to find are clusters of strongly correlated indicators. These are what we idiomatically call "stereotypes". I expect that linguistic communities, and human communities generally, are divided into "stereotype groups".

If there actually do turn out to be stereotype clusters of strongly correlated indicators, they may or may not match our ordinary stereotypes. These clusters, if they exist, give the potential to define many social terms (e.g. dialect, class, gender) in much more rigorous and data-dependent ways than has hence been done. It also allows for the possibility of discovering hitherto unknown social terms.

We may examine a physical science, rather than a social one, and find much the same ignored potential for complexity. The following quote is suggestive for what shall follow. David Faust quotes the physicist Martin Deutch:[13]

> A typical high-energy physics laboratory abounds in impressive sights and sounds. . . .
>
> I am sure that a scientist of 150 years ago, told to proceed with experimentation in this laboratory, would have no difficulty in making extensive observations on the change in the appearance of the lights as various knobs are turned, and on the variation in the pitch of the generating noise. But, of course, none of these changes directly accessible to sensory impression are really relevant to the experiment actually being carried out. . . . The only visible feature intimately connected with the actual experiment may be the position of a knob controlling the current through a magnet and a mechanical recording device indicating the rate of arrival of electrical signals from a small counter. . . . Both the change made -- the position of the control knob -- and the resulting effect -- the reading of the

> recorder -- seem almost negligible in the totality of material involved in the experiment. . . .
> . . .To make matters worse, there are probably several other knobs which could be turned with much more drastic effects of our recorder. In addition, there are literally dozens of similar recording devices and hundreds of other similar knobs connected with the experiment, some of which may show correlations much more marked than those under investigation. . . .
> . . .We see, then, a situation in which. . . a large number of related phenomena remain uninvestigated. . . .
> . . . .in a typical experiment in modern physics the apparatus involves complex influences frequently of much greater order of magnitude than the phenomenon investigated. Each of these influences can occupy a lifetime of experimental investigation before it is fully understood.

There are several interesting things we can distill from this quote.

The first is the incidental role that the motor and sensory makeup of humans makes to the physics lab. Modern knobs, as is not specifically mentioned in the quote but is well enough known, serve to do just one thing: regulate electrical current (and capacitance, etc.). This function may be carried out in very many other ways than by knobs of roughly the size of human fingers; in particular, the function may be carried out by means available to electronic computing machines. On the flip side of this, the quote points out that the recording device in question records the arrival of electrical signals. The apparatus sending signals presumably does not vary its behavior according to what sort of device receives them. This merely points to what I said at the start of this paper, that modern devices produce data in a manner entirely independent of the humans who analyze the data.

The second point of interest to distill from our quote is that "a typical high-energy physics lab" has hundreds of knobs and dozens of recording devices; that is, it has hundreds of independent variables determining dozens of dependent variables. In that the data produced by typical recording devices may be divided into numerous variables, the lab may in fact have more then dozens of dependent variables. Closely related to this complexity we should notice that (as described in the quote) experiments as performed only try to correlate one or a few of the independent variables with one or a few of the dependent variables. This is to look for a very simple explanation of a system which might act according to very complex regularities. If we stop worrying about the "real world" which (according to a realist) underlies the behavior of the equipment, then we see that all human physicists do is to look for pairwise correlations of indicators. Even these pairwise correlations are only studied very selectively. The most that we could ever get out of strict pairwise correlation of dependent with independent variables would be a system of equations, each of the form:

$D_1 = F_{11}(I_1) + F_{12}(I_2) + \ldots;$

For each F a function indexed for dependent variable by i and for independent variable; and for each I is an indexed independent variable. The optimal description of the data-production of the lab could, of course, be much more complex.

The method of finding the possibly complex regularities of a physics lab which I should like to follow is similar to the method I proposed in social sciences. Basically: ignore the preconceptions which scientists have and check every possible regularity. This is overstated, as there are infinitely many possible regularities; it gives a feel, though, of my proposed method. The best place to start in machine analysis of the data-production of the lab is with reanalysis of available data. We do this not because the recorded data are the best possible, but merely because it is likely to be cheaper to start with existing data.

This analysis will almost certainly pick out the regularities established by the scientists who originally did the experiments. It will also probably find many correlations which human scientists did not, especially ones of higher order of complexity than humans can conceptualize. Most probably,some of the strongest correlations discovered will be the ones which scientists discovered; not because of the inherent nature of the system, but because of the choice of experiments done.

The analysis suggested will also provide an heuristic guide to conducting new experiments. Less firmly established correlations can be checked, well established correlations can be utilized to minimize "noise" on a possible correlation, and research strategies can be developed. Even if scientists retain control over which variable correlations are of the most interest for external reasons, a machine can both propose experiments which are most likely to demonstrate or refute such correlations and find correlations of higher order complexity, of these variables. Almost none of this is available to humans working with paper and pencil, if only because of the sheer number of calculations involved.

Scientific Explanation

The question might arise, if I am more or less correct about the complexity of data-producing systems, of what good are explanations of data-regularities which are not conceptualizable by humans. I suppose this is a fair question, but it seems to presume that there is an opposition here with data regularities which humans can conceptualize. I am not quite sure what good the explanations which humans are able to conceptualize are. I am not so cynical as to doubt that there are any reason to do science; but neither do I think that there is a moral imperative to "discover the truth". The only eminently good reasons to do science (that I can think of) are concerned with building technologies which will generally help the quality of life of most people. It is not a consequence that "pure" science should be discouraged, any more than the lack of inherent worth should discourage people from playing games, from listening to music, or from doing any of the other activities which we do as expressions of human needs -- but neither is "pure" science any more worthy, in general, than other means of fulfilling human needs.

If we look at some of the cases where doing science has fairly direct technological ramifications, it seems clear that machines are generally better at the cognitive portion of the task. For example, in aerodynamics it may be that wind velocity, wind material porousness, wing lift, microturbulance, macroturbulance, size of wing, shape of wing, temperature of air, pressure of air, crystalline structure of wing material, and other factors are related in a more or less complex manner. If this degree of complexity exceeds the degree which humans can conceptualize then it is unlikely that humans will discover the optimal explanations of the data regularity. Machines can better find these complex regularities than can people; so if we want better airplanes then we would be better off letting machines do the cognitive work in aerodynamics. If finding more complex regularities does not lead to technological ramifications, then no harm is done by letting humans do the work of science.

There does seem to be something more in the practice of science than a search for technological ramifications. It is probably the case that no one seriously believes that particle accelerators have potential for producing useful gadgets. The main purpose of such scientific enterprises as particle accelerators is not producing technological ramifications (at least not if thought of in terms of technologies which everyday people use). The purpose is closer to "finding out how the world is"; which demystified amounts to finding data regularities. Some have claimed that this purpose is represented in basic human instincts. Failing this, it is probably true that the purpose is well represented in modern Capitalist and Socialist societies. If the purpose of "pure" science is finding data regularities, then we must use machines to find the more complex ones. If the purpose of "pure" science is merely to allow humans to conceptualize the data regularities of various systems, then machines must play a much less essential role.

Complexity and Positive Law

The limits of human cognition have practical and political consequences, and these should be regulated by positive law to a much larger extent than they are currently. Some areas of data-complexity where machines are necessary for optimal analysis have important practical consequences. These are not generally areas of "pure" science. It has little effect on most people's lives how adequately the behavior of cyclotrons is understood. Medical prediction and control, criminal sentencing and probation, toxicology and pollution studies, epidemiology studies, affirmative action, economic trends, and other factors do have a great deal of effect on our lives. In these areas, as elsewhere, human cognition is far less able to predict and control data regularities than are machines. However, there is little positive law in the USA concerning our normative relation (our epistemic stance towards, if you like) toward machine analysis of data. There should be a great deal, and probably will be in the near future (though whether it will be formed on a rational basis is something else altogether). I should like to make just a very few remarks on how I think such positive law should be constituted, at least at the level of generality.

Before I make these comments I wish to point at some possible misuses of machine discovered data regularities. What I can say does not apply uniquely to the regularities which machines discover, but it may be more crucial therein insofar as the machines which I have described analyze data without regard to what "facts of the world" the data stand for. An example where a correlation can de dangerously misused regards IQ tests.

It is well known that non-whites in the USA get considerably lower scores on American language IQ tests than do whites.[14] This kind of correlation can be made without any concern for the referents of 'IQ' or 'race'. A machine which analyzed existing records regarding IQ scores would find this correlation and possibly also much more complex ones of this sort (perhaps race, class, gender, parents' occupation, and score stand in a more or less complex relation). There are a number of mutually incompatible reactions we can have to this correlation. If we believe that there "really is" a factor in humans named 'IQ' (and that there really is a factor named 'race') then we will either conclude that non-whites are less intelligent, or that the existing tests do not really indicate this factor. If the latter we would probably think that better tests could be designed. On the other hand, if we believe that there is no such factor as IQ, then we should probably think the above correlation speaks very negatively of the activity of giving IQ tests.

My own belief, for example, that there is no such factor as IQ comes out of my conviction that there is no such factor as race, in anything like the sense which is normally attributed to it. Perhaps what 'IQ' is really a measure of is some function of race(?), class, gender, and parents' occupation. The point of all this is that mere correlation underdetermines consequences. We have an additional tendency to think that some terms, but not others, refer; and to think that causation operates in a particular direction. Hume first pointed out that the only experiential evidence we have of causation is correlation.[15] The rest must come out of our inclination to reify certain correlations. Machines cannot tell us which correlations to reify.

There are some areas where it is clear what use we should put correlations to. These areas are not much affected by what we want to consider causes and what we want to consider mere correlation. In these areas there stands a need for greater positive law regulation. Some of these areas were mentioned above: medical prediction and control, criminal sentencing and probation, etc. The first of these is dealt with by David Faust[16], in a fair degree of detail. The particular studies may be found in his book, I shall merely repeat the jist of them.

Doctors and other clinicians are often not able to make as good judgments about the presence or severity of a condition as are mechanical procedures working with the same data. In fact, even when doctors are given the predictions made by mechanical means, they choose to differ from them at the wrong times in their own prognosis or diagnosis. The mechanical procedures are not merely principally possible, but are becoming widely available as instantiated in computer based expert systems.

This is where positive law ought to enter. If an expert system has a demonstrated record of better prediction than human doctors, then it ought to constitute clear medical negligence to disregard the prediction of such a machine. This is an area where clear positive law statutes would do well to set the course of judicial decisions; the state of precedent clearly underdetermines a clear legal direction in deciding cases based on medical disregard for expert systems.

I shall make a few remarks about mechanical data analysis as regard criminal sentencing and probation, but shall leave out any remarks about the other areas of legal concern which I have mentioned. I believe the general pattern of my thinking will be more than clear with this one more example.

Correlation should be extended to routine aspects of sentencing and probation. Sentencing should be based on likelihood of recidivism (and what sort is likely). A large number of social factors could enter in to determining this likelihood: sociological, psychological, economic, criminological, and otherwise. There is an abundant source of past data from which to draw correlations, and these correlations are likely to be unconceptualizable by humans. Here is a simple and known correlation[17]: non-psychotic unpremeditated murderers of close family members or friends are no more likely to commit repeat offenses than are the general public (unless other overriding factors exist). There is no rational reason for sentencing such offenders at all. Sentencing and probation should be determined mandatorily based on such statistical and mechanical correlations. I maintain that this is the only plausible way of extending equal protection, and not basing such matters on the cognitive and social prejudices of juries, judges, parole boards, and law-makers.

The purpose of the last part of this paper has been to point to areas where human cognitive limitations have significant social impact. It is no longer the case that we must shrug our shoulders at these areas. The alternative to leaving cognitive decisions to humans now largely exists. I have argued that in these areas the use of superior mechanical means of decision/prediction should not be left to the discretion of individuals, but should be socially legislated. In an overall way this would lead to a more just and more humanitarian society -- and one better able to accomplish social goals.

Footnotes

1. Hilary Putnam uses this term in the third volume of his philosophical papers (Putnam, Hilary Realism and Reason. Cambridge University Press 1983). The distinction he makes is between 'metaphysical realism' and his own later position which he calls 'internal realism'. Metaphysical realists believe that certain truths exist in the world itself, without any regard to our position as observers, experimenters, etc. Putnam does not propose an outright anti-realist position, but he does propose that our position as observers, et al. partially determines the truth of sentences. The relevance of this to our discussion is the following. For a metaphysical realist, the world produces data without any regard for our position as data-collectors (or for the position of the machines which collect data on our behalfs). For an Putnam's internal realist (or also, for example, for George Lakoff's 'experiential realist' -- Lakoff, George Women, Fire, and Dangerous Things. University of Chicago Press 1987; or for Ian Hacking's effect driven science -- Ian Hacking, Representing and Intervening. ???) data are produced through our interaction with the world, not by the world in itself. This latter perspective indirectly emphasizes the flexibility which scientists have in their definition of the subject matter.

2. David Faust, The Limits of Scientific Reasoning, University of Minnesota, Minneapolis 1984.

3. Grice, H. "Meaning," Philosophical Review. 1957, no.66.

4. So, for example, in communicating a piece of information I wish you to believe this information; and also wish you to believe that I believe it; wish you to believe that I wish you to believe it; wish you to believe that I wish you to believe I believe it; and so on. This sort of nesting goes on indefinitely, and if these go amiss at any depth the communication has failed its full purpose. For example, if you do no believe that I wish you to believe I believe the piece of information then you are presumably suspicious of the honesty with which my original statement is made. However, we cannot actually conceptualize these nestings beyond a certain point. Grice deals with this by saying that we merely assume that "all is well beyond a certain point;" that is, we assume that none of the appropriate beliefs fail beyond the point where we consciously conceptualize the nestings -- and hence that there is no duplicity in the communicative act.

5. I use the term 'sentence' in the formal sense given it by modern mathematical logic. More broadly, our understanding of any well-formed formulae is dependent on the depth of quantifier nesting. I merely use the more common term 'sentence' for stylistic reasons.

6. *[formalization of logic -- frege, tarski, godel]

7. 'Effective' here means that no more steps are involved in transforming the original proof than a polynomial on its size. If we use some heuristic involving deeply nested formulae in a given proof, then clearly we have also shown that there exists some proof using less deeply nested formulae (that is, perhaps the final line of the proof has n-deeply nested quantifiers, but our proof which proceeded along some heuristic has lines n+1 deeply nested quantifiers, or greater). We may find such a proof using less deeply nested formulae by simply exausting all proofs using formulae which do not exceed nesting of a certain depth (this involves being able to limit the maximum length of each line to some finite length). The only point here is that this process does not, in any way, constitute a transformation away from the original proof, although it is a *finite* process.

8. I use three terms somewhat interchangeably in this paper. These are, 'indicator', 'variable', and 'factor'. Insofar as a distinction should be made among these, it is as follows. 'Factors' is the most metaphysically realist of the terms. A factor is an actual underlying fact about a particular part of a system in question. An indicator is some arrangement of the data collection which "indicates" the presence or degree of the underlying factor. An indicator may show a numeric or boolean value, whereas a factor does not have a value in this sense. Where a factor may merely be present, the corresponding indicator might show the value **YES**. A variable is the value of an indicator, or rather a schematic position where a value may be substituted. We need not assume that indicators really "indicate" something outside themselves, or that there are really factors in systems. These are merely convenient ways of talking. If we liked we could stop using 'factors' at all, and use 'readings' or the like in place of 'indicators'

9. William Labov, The Social Stratification of English in New York City. Center for Applied Linguistics, Washington D.C. 1966.
In a 1962 study Labov examined the social distribution of the sound /r/ (as opposed to its New York allephone /-/) in New York City department stores. Three stores of clearly different "status" were studied -- as well as various other independent variables: floor within the store; the sex, age, occupation, race of the informants. In each situation (between 68 and 125 interviews in each store), the response "fourth floor" was anonymously elicited in first unstressed, then stressed register. The presence of /r/ in each position was recorded by the surveyer, along with the other data mentioned above. Several correlations were found amongst the data; including a strong correlation between the "status" of the store and the presence of /r/.

10. For example, we could imagine that the subjects of Labov's study behaved as follows: whenever a supervisor is present, in any store, /r/ becomes entirely prevalent. However, in the absence of a supervisor, the "status" of the store determines the relative frequency. Had Labov, were this the situation, studied only the effect of "status" he might well have missed any significant result. To have found any significant result in this case, Labov would have had to include the independent variable 'presence of supervisor'. This sort of situation I describe is easily imaginable as the use of /r/ is related, by Labov, to degree "linguistic security" on the part of the speakers. Such a factor is quite likely to be affected by authority relations. Similarly, all those many factors which are actually not studied in actual studies may have just this kind of influence

11. Garcia, Erica C. & Ricardo L. Otheguy, "Being Polite in Ecuador," in Lingua: International Review of General Linguistics 61(2-3), Elsevier Science Publishers 1983.

12. In Ecuadorian Spanish, for example, two patterns of variation exist between the clitic pronouns *le* and *lo*(*la*). Garcia and Otheguy show that in other Latin American countries a semantic distinction is made between the pronouns on the basis of the felt activeness of their object. However, in Ecuador, due to the contact with the language Quechua, not all speakers control the gender distinction between *lo* and *la*. Those who do not control the distinction tend to avoid *lo*(*la*) altogether; while those speakers who do control the gender system tend to exagerate their use of *lo*, and especially *la* (the form completely unavailable to those lacking the *lo*/*la* distinction). The system involving the full gender distinction possesses higher status in Ecuador than does the genderless system -- one consequence being that where in most Latin American countries the more active *le* form is used as a sign a respect, in Ecuador the *lo*(*la*) form is used in these contexts because of its association with the higher status full-gendered system. A study which attempted to link respectful address with use of the *le* form would not be successful if it did not

include region as an independent variable (that is, if it used a mixture of Ecuadorans and others as informants -- who obey contradictory patterns in their use of *le* versus *lo*(*la*)).

13. Faust, op.cit. p.16 -- taken from Deutsch, M. 1959 "Evidence and Inference in Nuclear Research," in Evidence and Inference. edited by D. Lerner. New York: Free Press of Glencoe.

14. *[citation on IQ test stuff(whites v. others)]

15. see Hume, David, An Inquiry concerning Human Understanding. Bobbs-Merrill Educational Publishing: Indianapolis 1980 (1748).

16. Faust, op.cit., see chapter three in particular.

17. *[repeat murder offenders data]