ARTIFICIAL ABSTRACT ACTIVITY

Michael Sussna

December 19, 1984

for Artificial Intelligence, Fall 1984
given by Professor John Ritschdorff
at Marist College, Poughkeepsie, N.Y.

ARTIFICIAL ABSTRACT ACTIVITY

What is Artificial Intelligence? Some might say that we shouldn't worry about a strict definition, that it is just a label and we shouldn't attribute too much importance to it or take it too literally, that it covers too many things to fit under one definition, or that even if we could pin it down we shouldn't, so that it is left open-ended and able to change to include new developments. Others might say it would be useful to try to define AI, that a definition would help us to decide whether a software package, for example, really incorporates AI or not, or whether an area of study should come under the heading of AI or not. In fact, as we shall see, the effects of trying to define AI can be much more far-reaching.

Of those who would define AI, some would not try to take the two words "Artificial" and "Intelligence" literally, but rather describe AI functionally or in some other way. For example, we could say that AI consists of two reciprocal activities: studying human intelligence in order to build devices with intelligence, and building devices with intelligence in order to study human intelligence. What about devices that need not be intelligent themselves but which help our intelligence? What about our learning how to be more intelligent because of these studies? Should these two developments be considered part of AI or not? If not, might they not be closer to intelligence than studies of motor skills, dexterity, and even vision? A definition would help us decide, but should we shape our definition to conform with what should be included and what should not, or include or exclude topics in conformity with our definition? The answer is "both."

Initially we come up with a label to refer to the concept we have in mind. After an interim shakedown period where we see if the label really adequately conjures up what we mean, either we might modify it to better fit our idea or replace it with another better suited to the task, or the label will stabilize. After we settle on a label we're happy with we can use it, as it stands for the concept we are concerned with, to decide whether a topic falls under the definition or not.

Some additional settling or "relaxation" of the idea with the terminology may take place, but the inertia of heavy use of the label/meaning pair will increasingly help the pairing resist further change. If sufficient thoughtfulness goes into the naming then this should not present much of a problem in the future. Otherwise we can be stuck with bad terminology that we can either live with or get past by coming up with new nomenclature.

Is this all the more reason to avoid taking the term "AI" literally in the quest for a definition? Possibly. But how far can we go by taking the words "artificial" and "intelligence" literally? What if in that analysis we discover not only that the label is an apt one, that the concept that it conjures does conform well with what we mean by it in actual practice, but that in addition its definition allows us to distinguish topics which qualify as part of it from those which do not, or decide whether a software package has it or not, or whether a gadget with a microprocessor in it is intelligent or not; that the definition can "breathe," leaving room for new developments to come under it as they arise? What if, furthermore, the analysis reveals not only that we can identify topics belonging in AI but that we can begin to classify them into a taxonomy, a map of how they fit together, as well? And what if we discover in the process a whole world to explore from which we can bring back once only dimly glimpsed treasures such as a richer understanding of human intelligence; of whether other big-brained animals or "lower" animals have intelligence; of whether there is "blind" intelligence at work in evolution; of what criteria to use to decide whether intelligence is present and if so how to communicate were we to encounter life-forms from other worlds; and of where our own development might be going, the discussion coming full circle by taking into account the revolution in the technology of intelligence that is beginning?

Obviously the key to unlocking all of these answers lies not in analyzing "artificial," which just means "man-made," but in delving into the depths of "intelligence," which many attempted definitions of AI do not do, taking its meaning as a given. So, just what exactly is intelligence?

In order to answer that question we shall venture in search of naked intelligence in an effort to isolate its essence. Our path shall travel ever inward, from the outside view through the surface and on inside to the heart of intelligence, as if it were an atom and we were probing its subatomic secrets.

Along the way we shall see many paths branching off from ours, beckoning for further exploration, which we shall follow to varying distances, but not so far that we lose track of the main trail. Thus, for example, we will not claim to provide the taxonomy of topics mentioned above, but only point the way for its discovery. One final note before we start the quest: perhaps an ultimate definition, description, or treatment of the concept of intelligence exists in some dictionary or somewhere in the literature; if so, we will hopefully be consistent with it, and significantly we will be bringing the subject to life, being Galilean and investigating it firsthand. If not, it is possible that this exercise will bring us closer to that goal.

So, let us begin.

A good place to start our journey is with the dictionary definition. But two points strike us immediately: there is more than one dictionary definition, and if we happen to look in different dictionaries, even two editions of the same dictionary, we can see differences in what is given. Even a single definition in some cases is actually a conjunction of submeanings which somewhere else might justifiably be given separate status. But take heart -- we can and will deal with each of these problems as they appear, and any others that we may find as we go.

First, let us try to capture the feeling of what it is that we're really talking about. That will help us pick the definitions that are on the mark from those that are not.

Second, let's just take the ones we do like into advisement. We'll file them away for future reference, possibly recalling them later. We'll look at the question from many other points of view and perhaps we'll be able to look back and see that one of these definitions was after all on the bull's-eye.

So, what are we really talking about?

Isn't intelligence what distinguishes us from the animals, what makes culture and civilization possible? Isn't it "having brains," the ability to think? Isn't it also good judgment, common sense, the ability to choose to our benefit? What about resourcefulness, ingenuity, invention? The list goes on and on, but we're getting a good flavor of it.

We're ready to look at the dictionary definition. The following is the full entry for "intelligence" from Webster's New Collegiate Dictionary, 8th edition, 1980:

8-1a(1): the ability to learn or understand or to deal with new or trying situations: REASON; also: the skilled use of reason

1a(2): the ability to apply knowledge to manipulate one's environment or to think abstractly as measured by objective criteria (as tests)

1b: Christian Science: the basic eternal quality of divine Mind

1c: mental acuteness: SHREWDNESS

2a: an intelligent entity; esp.: ANGEL

2b: intelligent minds or mind <cosmic ___>

3: the act of understanding: COMPREHENSION

4a: INFORMATION, NEWS

4b: information concerning an enemy or possible enemy or an area; also: an agency engaged in obtaining such information

The relevant ones for our purposes are 1a(1), 1a(2), 1c, and 3. That means that we can eliminate several definitions from further consideration.

Let's look in detail at the ones we've picked, evaluate them against our preliminary intuitive feeling for what we mean by intelligence, decide if there is agreement, and if not, try to determine where the disagreement lies.

1a(1): This entry actually contains five definitions! We shall break them out and deal with each one individually. We'll give them labels for reference purposes as follows: 1a1a-e respectively.

1a1a: the ability to learn -- Yes, this is relevant, but it needs qualification -- dogs can learn, for example. So, is there a particular kind of learning that only people can perform?

1a1b: the ability to understand -- Relevant, but again, dogs can understand -- so, is there a kind of understanding unique to us?

1a1c: the ability to deal with new or trying situations -- Yes, this probably requires intelligence.

1a1d: REASON -- This one seems very much at the heart of things.

1a1e: the skilled use of reason -- Here we're getting into relative intelligence.

1a(2): Again we need to decompose the entry into a set of definitions, in this case only two, which we'll call 1a2a and 1a2b.

1a2a: the ability to apply knowledge to manipulate one's environment -- The key here is in applying knowledge and in the implied idea of purposeful activity, not so much in the manipulating of the environment (beavers, termites, and birds all manipulate their environments, don't they?).

1a2b: the ability to think abstractly as measured by objective criteria (as tests) -- this may also lie at the heart of things, if we stop after the word "abstractly."

1c: We've got two concepts here, so: 1c1, 1c2.

1c1: mental acuteness -- Relative intelligence again but with a hint of fast comprehension -- "acuteness" means sharpness, quickness (see definitions 3 and 4 of "acute" in the eighth edition), and also a note of sensitivity of perception, ability to detect subtle distinctions (definition 2 of "acute").

1c2: SHREWDNESS -- Smart in the practical sense, has to do with judgment and cleverness

3: --> 3a, 3b.

3a: the act of understanding -- This doesn't make sense as a noun meaning the same thing as intelligence unless "act" is being used in an unusual sense, as in "act" definition 2, where it is referred to as a state. Thus, a state of understanding -- still a little awkward, and understanding was already covered in 1a1b. It would be useful to see this used in a sentence.

3b: COMPREHENSION -- looking at intelligence as the ability to comprehend is fruitful, since although "comprehend" means basically the same thing as "understand," it is coming from a slightly different perspective and so contributes a feeling of assimilation of knowledge, of fitting the thing understood to a body of knowledge, to something comprehensive.

Well, instead of making things clearer, this has only multiplied the possibilities. And we're not through yet with the dictionary definition.

We're going to look at the entry for "intelligence" in the seventh edition of Webster's New Collegiate Dictionary in the same way. When we're done with that we'll move on into the next phase of our investigation, the search for a common denominator in a list of intelligent activities such as problem solving and decision-making. For now we're just developing a better feeling for the kinds of things that intelligence might be. We won't try to narrow these many meanings down to a chosen few or try to organize this seeming bedlam just yet. Much later we shall return to these dictionary definitions and look at them again in the light of what we've learned. We will then be able to detect order where now we can only see discord and diversity.

Most of the definition in the seventh edition is the same as in the eighth edition. We'll list the mapping here, and only give the entries where the editions do not match:

7-1a(1): the capacity to apprehend facts and propositions and their relations and to reason about them: REASON, INTELLECT; also: the use or exercise of the intellect esp. when carried on with considerable ability

7-1a(2) = 8-1b
7-1b = 8-1c
7-2 = 8-2
7-3 = 8-3
7-4 = 8-4

Definition 1a(1) will be split into four definitions, 1a1a-d:

1a1a: the capacity to apprehend facts and propositions and their relations and to reason about them -- First we had better redefine "apprehend" in more familiar terms. We'll use definition 3 of "apprehend" as a transitive verb in the seventh edition: "to grasp with the understanding: recognize the meaning of." So, we're talking about recognizing the meaning of facts, propositions, and their relations and reasoning about them. "Fact" here must be seventh edition definition 5: "a piece of information presented as having objective reality" or a true statement, and "proposition" must mean basically "statement." Relations between facts and propositions might include subject matter relation and consistency or inconsistency of statements. Reasoning about these things has to do, among other things, with manipulating them and reaching conclusions, arriving at truth and noting falsity. This topic might be another at the heart of the matter and of course it ties in directly with one mentioned earlier, reason.

1a1b: REASON -- Already discussed.

1a1c: INTELLECT -- let's focus on definitions 1a and 1b of "intellect" in the seventh edition:

1a: the power of knowing as distinguished from the power to feel and to will: the capacity for knowledge

1b: the capacity for rational or intelligent thought esp. when highly developed

Themes are beginning to become familiar: knowledge, rational (relating to reason), degree of intelligence. A very nice word appears here: "thought."

1a1d: the use or exercise of the intellect esp. when carried on with considerable ability -- Intellect again -- the "nucleus" of the "cell" that is intelligence? Also, degree of intelligence again.

Ok, we've seen enough of dictionary definitions for now. Please just note that we can go on quite a side tour exploring the definitions of terms used in definitions, indefinitely far.

All right, let's start the next leg of our journey. Here we'll try to close in on what seems essential to intelligence. This will lead into the search for a common denominator mentioned before. Bear in mind that both of these discussions, working in a more directly AI-oriented frame of reference, are still views of intelligence from the outside, but will prepare us for the plunge into its interior.

Let's picture a robot with a body more or less like a person's, and with abilities more or less like ours. Then let's peel away the layers of its powers until we find intelligence. Our robot can move around, but even the lowliest creatures have mobility, so that can't be intelligence. It can manipulate objects, but that's common as well. It can see and hear -- that's important, but we're still not there. It can speak and listen -- now we're getting close, but these are still not essential -- they might be considered applications of intelligence. But surely language is a certain sign of intelligence, something animals don't really have? A) our use of language may require intelligence but that doesn't mean that intelligence requires language. B) Animals do have language. The difference may lie in the usage -- more later. What about personality? No. Emotions? No. Very related but only that. Intellect? Yes. This is a concept intimately intertwined with that of intelligence, a component of our mental makeup that we can safely associate with our elusive prey.

Let's list some core concepts in AI work, some archetypal, paradigmatic themes which seem very directly related to the functioning of intelligence, as mobility, manipulation, and peripheral sensing are not, in themselves. Then let's try to isolate a common denominator from these concepts.

1. Thinking -- working with ideas, thoughts. In AI terminology we might refer to deduction, logic, reasoning, and cognitive acts.

2. Problem solving -- Working to connect a start state with a goal state.

3. Decision making -- Some decisions choose from among single actions -- we're concerned with choosing from among series of actions. We might distinguish two kinds of activity here. One features choosing the best path of actions to a goal state. The other might be a "what-if" type of function where there is no one goal state in mind. Rather we wish to pick the best path leading from the start state where paths may continue to varied outcomes.

4. Planning -- Finding a sequence of steps connecting the start and goal states.

5. Learning -- Fitting a new item to what is already known.

6. Language -- Communication with others, allows transfer of meaning, includes translation between

representations while preserving meaning.

7. Understanding -- Seeing the meaning of.

Is there a common denominator here? Yes, bridging ideas, the ability to get from one meaning to another. Are these states that we're working with ideas, meanings? Yes, because the mental models of situations which we actually work with are ideas. Let's look at "idea" definition 3c from the seventh edition: "an entity (as a thought, concept, sensation, or image) actually or potentially present to consciousness." What about meanings -- what's the relationship between ideas and meanings?

OK, your interest is piqued. This is the kind of question we will not only be posing quite often somewhat later, but will be answering as well.

But before we penetrate like that into the internals of intelligence we need to make a considerable detour to round up any stray nuances which demand addressing. For as we shall be reminded, there are a couple of important themes to be reconciled with our recent thinking about problem solving and so forth. We need to put each of the themes in their proper relationship. As we do so, we will explicitly acknowledge the threefold structure of our definition so that from here on, everything that we encounter can be conveniently categorized as belonging to either 1) the external properties; the instances of intelligence including applications of it, 2) the surface properties; what it is in general, what it is like as a whole, without looking at how it is put together; intelligence as a black box, as a subroutine, or 3) the internal properties; what it looks like inside, how it is put together, etc.

The first of these other themes to look at is that of relative intelligence, degree of intelligence. Just as "smart" can mean "smarter than (average)," so "intelligent" can mean "more intelligent than (average)," for example. We will see how this second meaning of "intelligent" relates to the other two as we go.

The second of these other themes, and therefore the third theme overall, is that of practical intelligence -- good judgment, common sense, shrewdness, and so on. A handy way to illustrate the distinction between this theme and the previous one about working with ideas is asking "Have you ever known someone who was bright but not smart?" This second theme has to do with practicality, one's actions, and resourcefulness. How does this meaning of "intelligence" relate to the one about being able to get from one idea to another? We'll see shortly. First we're going to follow a train of thought sparked by wondering about theme three.

What is practical intelligence? It brings to mind successful response to situations, making good moves, and in general behavior of an organism or system to continue its existence or even improve its lot; in a nutshell, doing what's smart.

A sophisticated formulation of this concept might be "maximizing ultimate benefit by maximizing immediate resource utilization" -- in other words, doing the best you can with what you've got. But if it's doing what's smart that we're talking about, what about the work of all of the incredibly intricate mechanisms that we see in nature whether at the organism, organ, cell, or molecular level? (We'll keep things "organic" for present purposes, although we could continue down the hierarchy of natural structures to the atomic, meson and baryon, quark and lepton, and possibly the "rishon" levels (Harari 1983).) These mechanisms are geared for continuing the organism's existence, however humble, by doing what's smart in the situation at hand. What about all the "programming" that is inherent at the DNA level on up to our very beliefs and rules of behavior? How do we interpret this in terms of intelligence?

The fascinating concept of seeing our rules, habits, customs, and ways of deciding, for example, in terms of the metaphor of computer programs is introduced by John Lilly in his work "Programming and

Metaprogramming in the Human Biocomputer" (Lilly 1973). One illuminating aspect of his approach is the idea of our ability to change our own programming. This is analogous to "PROM," programmable read-only memory, in the world of computer memory.

To extend the analogy with computer memories, perhaps there has previously only been "ROM," read-only memory, for the rules of behavior in organisms. Does this capacity of ours for self-programming point to a distinction between other species and our own? Yes. A program is an abstract structure of actions. In order to program one must be able to perceive and manipulate abstractions. We could just as easily call these abstractions concepts. The difference between them is that "abstraction" connotes abstracting, extracting some property from a set of things, which induction and generalization are examples of. "Concept" on the other hand connotes conceiving, thinking up, picturing. We have here finally hit on the essential ingredient in intelligent activity, the ability to operate at an abstract level, with concepts. We will fine-tune this definition and refine this diamond mined from the mass of our knowledge in due course. In the mean time we will look briefly at the definitions of "abstraction" and "concept" and then get back to our task of reconciling themes. Suffice it to say for now that the ability to program oneself is an application of intelligence.

"Abstraction" definition 1b, eighth edition: an abstract idea or term

"Abstract" definition 1a, eighth edition: disassociated from any specific instance <___ entity>

"Concept" definition 2, eighth edition: an abstract or generic idea generalized from particular instances

Getting back to practical intelligence and self-programming, we might think in terms of adaptability. We are highly adaptable creatures -- we can react "realtime" to changing conditions and take steps to deal with them. Looking back through evolutionary time, we can see that the flexibility for adaptation decreases the further back we go, as the time frame that that adaptation occurs in grows.

We can adjust as things are happening. Lesser creatures may not be able to change the way that they respond to events except over generations, through genetic adaptation, whereas we can modify our responses temporarily or indefinitely. This flexibility of ours can be a curse and a blessing: life is more complicated but the potential for varied experience is that much greater.

Is there intelligence at work in the adaptation and programming of lower species, albeit hidden? We can point to the information present, the knowledge of what to do, when to do it, and how, but that's really only the result of the intelligent activity, if it is that. But where then can the "intelligence" be localized? In the interaction of a species' genome, its genetic information and programming, with the changing conditions of the world. In other words, in life (living things) vs. life (the world).

"But species don't have genomes, individuals do," one might object. And, "How can the world interact with genes, which lie buried within cell nuclei, which claim is tantamount to a much-ridiculed Lamarckian point of view?" Species evolve through the successes or failures of their individual members, which live on through reproduction or die out, largely due to their genes' expressions, the creatures' characteristics, in confrontation with conditions in their world, their genetic store thereby continuing on or not, the aggregate of which, across the individuals of the species, could be considered the species' genome. As far as the world affecting the contents of genes, perhaps there is some direct effect where mutations are caused by radiation for example, but what we mean here is that through the trial and error of natural selection, what works will live on and what does not will not, resulting in a different set of genetic information being present.

Without getting into the controversial but fascinating subject of morphogenetic fields and biological development (Sheldrake 1981), we still cannot escape from some interesting conclusions arising from our current line of inquiry. If we claim that "intelligent" but "blind" adaptation has been occurring, then evolution implicitly contains intelligence. The conceptual aspect to this activity is not in the objects of its operation, physical things, nor in the operations themselves necessarily, but in the context of the activity. This is perhaps the most implicit form of intelligence.

So, going back to our idea of three themes and the threefold structure of our definition, how do things look now? Well, practical intelligence, whether realtime or over generations, is an application of intelligence, an external property of it. Relative intelligence, whether sophisticated or as primitive as blind trial and error, is a surface property of intelligence. Without referring to specific examples or applications of intelligence, or discussing the mechanics of intelligence, the external and internal properties of intelligence respectively, we can say that there are degrees of intelligence.

Practical intelligence may have been the first application of our brand of intelligence. It is still, of course, a primary use for intelligence. Perhaps this explains why it is synonymous with intelligence, as indicated by the fact that it is one definition of intelligence and is called by that name. What distinguishes our brand of practical intelligence from earlier brands? The ability to perform trial and error realtime, among other things. How might this ability have arisen? We had many of the components required already, there being a long, methodical progression of increasing powers of perception. At some point the fledgling power of intellect, "in-sight," the "mental sense" as opposed to the five senses, came into being. What other components were already in place? Things like memory, attention, will, the processing of the signals from the five senses, and the abstract properties of things, which are the sense data of that new sense. It may just have been a matter of time until our forebears, sharpening their conceptual skills in dealing with the crises encountered daily in that more primal time, found themselves with moments of relative peace, free time, and turned their new sense to wondering, asking, devising, and creating. The rest is history, so to speak.

But surely there is more to our brand of practical intelligence than the ability to perform trial and error realtime, isn't there? Yes, but that in itself is a big step forward from what preceded it. But what about performing trial and error realtime -- what does that take?

If we simply try some physical action and react to the result, that's minimally intelligent, but even that may involve an element of forethought, intention, and concept handling. Hopefully we might put more planning into the effort and be consciously testing effects. An even higher degree of trial and error would consist of trying combinations in our mind rather than in actuality. This involves simulation, modeling, perhaps some type of systematic ordering of the possible things to try. The more we know about the subject matter (problem domain) and about searching for desirable solutions (problem and solution spaces) the better we're getting, and the higher our degree of intelligence. We can utilize constraints to limit our search, but perhaps even more intelligently begin by asking the right questions and thereby be involved with a meaningful search space.

Speaking of "meaningful," it is time to return to the main path, where we will finally plunge into the internals of intelligence by first asking about the relationships of meanings, ideas, concepts, knowledge, and properties.

Let's talk about knowledge. We only know what we register, we only register what our attention is on, we can only pay attention to one thing at a time, everything we've attended to is in our memory, and a memory is impressed in proportion to the poignancy of the experience it represents. The different kinds of knowledge reflect the difference in the senses that each kind has come through, which in turn reflects the kinds of sense data able to be detected.

Items of knowledge may be associated in various ways. One way is by proximity in time of occurrence. A second is by proximity in location. Another might be by associating a series of such co-occurrences, an association of associations as it were. This latter type of association is of course the foundation of science, regularity in phenomena (Although "constant conjunction" of two things is not sufficient to prove a cause and effect connection exists between them, it is necessary for there to be any chance for such a relationship.).

If we have a sense that can see the relationships between the properties of things, and that can produce new knowledge from existing knowledge, then we have induction and deduction, a kind of perception perhaps denied animals other than ourselves. A correspondingly different kind of knowledge is stored.

It is one thing to see immediate properties of external phenomena, for example that a car is blue -- it is another thing to see properties of properties, for example that blue is a color. These latter properties only the mental sense can see. These abstractions are concepts, as opposed to simple percepts. The sum of our knowledge about some particular topic is the collection of the individual items of knowledge that we have in memory which are united by the common bond of relating in some way to that topic. That is not to say that these items need be stored together physically, forming an explicit structure. These items are stored together logically, forming an implicit structure. This structure is available upon request, called into being dynamically, any item in it giving access to the others via links of association.

If two items are related but we don't know that, the link can be created by learning of their connection. Besides being taught the new knowledge or seeing it written somewhere, such learning can take place by finding the connection ourselves through thinking.

Besides distinguishing between the two kinds of properties, first order properties which can be detected with one of the five senses, and second order properties, which are properties of properties, the relationship between properties and objects needs to be described. All properties can be considered objects, but not all objects are properties. For us, entities like concepts, ideas, and meanings are objects. From a certain perspective even processes are objects. This unusual treatment works with what will be called logical objects.

Syntactic roles such as subject, object, and predicate, parts of speech like noun, verb, and preposition, and even most nouns, and all verbs and prepositions when used together in the form of a relationship, all represent concepts.

To illustrate, if a particular noun is a type of thing, or category, as opposed to a particular thing, or instance, it represents a concept. For example, "four" is an instance of the concept "quantity," as "blue" (as a noun) is an instance of the concept "color." Interestingly, both "four" and "blue" are themselves concepts, as they can have instances. It is easier to see how "four" is abstract than it is to see how "blue" is. If we're using "blue" to refer to a particular hue that we might be seeing in the physical world, then it is concrete rather than abstract. If on the other hand we are using "blue" to refer to a type of color, a range of colors that might include light blue, royal blue, and navy blue, for example, then it is a concept. Concepts are variables to be filled in, instantiated. The roles something can fill give a description of what it is. These role or "combinability" properties of a thing are surface properties of it, properties of the thing as a whole.

This combinability property is well exemplified by relations, or relationship templates, composed of verbs and optionally prepositions and slots to be filled in by noun phrases as direct objects, or operands. For example, "giving ___ to," "thinking about," "waiting for," "hitting with ___," "escaping from," and "swearing by" are all examples of multi-term "verb" relations and, in two cases, of templates. These relations do not become part of relationships until used in a sentence, a statement of relationship, when their operands are filled in. Until that point they are just concepts; after that point they become concrete. Only certain noun phrases can fill particular slots for a given template, or participate in the relationship governed by a relation/verb. This constraint on participation is an effect of the unique nature of a relation and reflects its combinability.

Well, we've talked about knowledge, properties, and touched on concepts, but what about meanings and ideas? Where do they fit in?

Most familiarly, meanings are what symbols stand for. Although no symbol for a meaning need exist for the meaning to exist, there is a distinct connotation to the word "meaning" of being associated with symbols and expression. What is being referred to when speaking of meanings divorced from symbols? Suppose that we are about to give a symbol to some meaning, come up with a word for it -- what is it?

It's the idea of the thing being named. If an object or process in the physical world, for example, is being labeled, there is a "triangle" of items involved: the physical thing, our idea of it, and the term for it (Sowa 1984, p.11). The physical thing and the term for it are joined by way of the idea, not directly "touching."

It's a little easier to think in terms of "ideas" divorced from symbols and expression, but in the sense of "meaning" that we're using they're synonymous. Are all things seen by intellect ideas? What about images or imagined sounds? Are percepts in general, whether just coming through the system or being recalled or simulated from memory, ideas? Reviewing definition 3c of "idea," they are. Does this mean that animals, having percepts and consciousness, have intellect as we like to think of it? No. All things seen by intellect are ideas but not all ideas are seen by intellect if, for example, images and percepts count as ideas.

But here's something to think about. The following anecdote is from an article on human spatial thinking (Cooper and Shepard 1984, p.106): "One of us witnessed a German shepherd retrieving a long stick that had been thrown over a fence from which one vertical board was missing. The dog bounded through the gap, seized the stick in its mouth and plunged headlong back toward the narrow opening. Just as catastrophe seemed imminent the dog stopped short, paused and rotated its head 90 degrees. With the stick held vertically it passed through the fence without mishap. The operation that took place in the mind of the dog in the moment before it turned its head presumably was not verbal. Might it not have been a preparatory mental rotation of the stick? (And was it not by spatial visualization rather than verbal deduction that you, the reader, understood how catastrophe threatened and was averted?)"

Is this intelligence? No. Spatial reasoning with a visual simulation in a dog's mind is perceptual, not conceptual. The scene projected is a percept, and the operations on it and the percept of self-with-stick-in-mouth are percept operations. The dog may simply be doing visualize-and-test until the hold-stick-vertical-by-turning-head-90-degrees option clicks as the solution. Intellect works with percepts as well but performs conceptual operations on them. Also, some of its percepts are of logical objects such as concepts. Isn't testing conceptual? If it is done consciously, with a concept of testing in mind. The dog just does it without realizing that it is doing it, as a reaction to a visual "radar" warning. Explicit work with concepts is what distinguishes human intelligence from animal intelligence. Which brings up a fourth theme, that of the explicitness of the intelligence in instances of its appearance.

We saw "blind" intelligence, embedded in evolution, and not very explicit at all. We've seen animal "intelligence" in the sense of perception and mental processing, the forerunner of human intelligence. We will mention human-like intelligence as possible in other terrestrial species, and in non-terrestrial life-forms. This human-level intelligence quite explicitly works with concepts. We will also look at artificial intelligence and note degrees of explicitness as far as data, code, and subject matter.

But we still have much to discuss about human intelligence so that we may carry that understanding into new areas and be able to recognize whether we're dealing with intelligence or not.

Intellect sees thoughts and concepts. What's the difference? Thoughts may work with conceptual material, but we tend to think of thoughts as ideas concerning relationships, more or less statement-like, of the form "this thing, x, was/is/will be in some relationship, r, with this other thing, y." It is natural to think of the meaning of a statement -- that is just its conceptual content. Of course thoughts can be visualizations, reminiscences, and so on, perhaps containing purely perceptual content, or a mixture of perceptual and conceptual content, but it is conceptual thought that we are concerned with. That is not to say that we are not concerned with the physical world. The physical and conceptual worlds are intimately and inextricably intertwined. In fact they are just different aspects of the one world of reality.

Concepts can be considered building blocks for statements. Statements assert some relationship between operands. The operands are instantiated concepts. The relationship is an instantiated relation concept. For example, if we have an abstract relationship "x r y" where x, r, and y are not specified, we can call this a concept of relationship. If we instantiate r with a specific relation, say "being the father of," we get "x R y," "x being the father of y." You say "That's not grammatical, it's not a well-formed statement." You're right, but it's not supposed to be. It is also a concept. Until all of the variables are bound, that is, given values, the relationship is not an assertion. If we continue filling in values, say "Meister Eckhart" for x, we get "Meister Eckhart being the father of y," represented by "A R y." We could alternatively have instantiated y instead, with "German philosophy," for example. This would result in "x being the father of German philosophy," of form "x R B."

All right, suppose we complete the filling-in process to yield "Meister Eckhart being the father of German philosophy," "A R B." "Still not grammatical," you say. "Still not meant to be," is the reply. One final instantiation operation is required in order to turn this concept into an assertion. That is the operation of "state-ment." This is equivalent to appending the concept "is the case" or "is not the case" to our relationship concept. It gives the state of the relationship with respect to reality, whether affirmative or negative. Our relationship concept would finally become the assertion "Meister Eckhart being the father of German philosophy is the case." At least this is grammatical, if an awkward construction. It is an easy grammatical (syntactic) transformation to the conceptually (semantically) equivalent "Meister Eckhart is the father of German philosophy." The format of the statement might be represented as "A R B +," for example. The basis of questioning in programming languages like Prolog consists of putting a question mark where x is, for example, and getting "? is the father of German philosophy," or putting one where y is and getting "Meister Eckhart is the father of ?" Finally, we could put one where r is and get "Meister Eckhart ? German philosophy."

Just as we went through three successive phases of instantiation, filling in the relation and operand values and performing "state-ment," so we can move in the opposite direction, performing complementary abstraction operations. For example, we can abstract the concept of the relationship "A R B" from the statement "A R B +," the more general concept "A R y" from the concept "A R B," and so on to the very general concept "x r y."

An alternative conceptual representation for what we've been looking at is that of predicates. We can equivalently represent "x r y" as r(x,y), and so on. Although the predicate/operand format allows putting the operation or relation first and then stringing the variable number of operands after it, this is not how people work. Our assertions reflect the sequence of operators and operands that is natural and comfortable for us. Although there are distinctive variations on word order from language to language, for example putting adjectives before nouns or after nouns, we still maneuver roughly similarly conceptually, and not in the artificial manner of predicates.

That is not to say that predicate representation is not valuable. If it is better suited to computer processing, fine. As long as people are not forced to convolute themselves in order to operate in its manner, that's what counts. We can always translate between representations so that an appropriate representation is used in a given context. There is a long way to go in emancipating people from machine-friendly conventions of programming, but the inexorable trend toward natural language querying reflects the fact that the meaning of "high-level language" most important in the long run may be the one meaning "as close as possible to the natural way of expressing something" rather than the one meaning "most machine instructions generated" as in "most powerful commands." Let the translation into a computer-usable form be as transparent to the human users (including programmers) as possible.

So, we have seen something of concepts and thoughts, meanings and ideas, knowledge and assertions. What's next?

We're going to look at what conditions are necessary in order for intelligence to be present. Is "intelligent" data, if there is such a thing, enough in itself? How about "intelligent" operations? An "intelligent" context? "Yes" to all three.

"Intelligent" data would be concepts. "Intelligent" operations would be operations that work with concepts. Otherwise non-intelligent operations executed as part of a plan or for a conceptual purpose would also qualify, there being a conceptual context.

We saw that concepts are abstractions, "variables," in that they are things that are able to be instantiated. If some concepts are singular items like color, quantity, weight, shape, size, and so on, how do we reconcile this with other concepts in the form of relations or relationships, like "being the father of" or "Meister Eckhart being the father of German philosophy"? Each of the singular items can take the form "the ___ of," for example "the color of." This is analogous to "the father of." "Father" can be extracted from that phrase to yield a singular item that is also a concept. "Being" can be prefixed to our "the ___ of" forms to yield "being the ___ of." So we can build up constructs of concepts. This, by the way, is an example of a conceptual operation.

Here is a list of some conceptual operations in template form:

```
Abstraction    of ___ from ___
Instantiation  of ___ as   ___
Translation    of ___ from ___ to ___
Representation of ___ by   ___
Classification of ___
Composition    of ___
Decomposition  of ___
Interpretation of ___
Definition     of ___
Communication  of ___ to   ___
Creation       of ___
Perception     of ___
Generalization to ___ from ___
Deduction      of ___ from ___
Implication    of ___ from ___
```

What about conceptual context? Recalling the visualization and testing of the very aware German shepherd, we differentiate the testing operation it performs automatically, non-consciously and non-intentionally, from that performed consciously by a person with a purpose in mind. The human is testing in some conceptual context, possibly as part of a plan, program, or process. If testing were an operation that need not work with concepts, and it were not being done in a conceptual context, nor with conceptual data, then intelligence would not be at work.

So, the presence of any or all of the three criteria, operation on conceptual data, conceptual operation, or operation in a conceptual context, is sufficient for intelligence to be present. The presence of at least one of the three criteria is necessary for intelligence being present.

We can look at the arts and architecture, for example, in this new light. We see that although visual or aural mental imaging is perceptual work, it becomes intelligent activity when a conceptual element is introduced. When we express a thought, whether to others or to ourselves, it is a conceptual event. We may be attempting to communicate information, a feeling, or a stimulus for others to do with as they may, or we may combine function, distinctly conceptual, with aesthetics.

Looking at the way that we perform the tasks listed earlier, when we were looking for a common denominator, will help characterize the mechanics of intelligent activity.

The first item in the list, thinking, is such a broad topic that it really appears synonymous with intelligence. We said that it is working with ideas and thoughts. We have seen since then that work with concepts is the essence of intelligence. Therefore thinking, when the ideas or thoughts being worked with are conceptual, is intelligence. Intelligence is thus a subset of thinking. This may sound strange, but if ideas and thoughts can include perceptual material being worked with in non-conceptual ways, then this thinking is not intelligence. It is extremely similar to intelligence and could be considered a precursor to it. Remember that this visualizing, for example, as soon as it works with anything conceptual, any categories for example, becomes intelligence.

Problem solving, working to connect a start state with a goal state, is intelligence par excellence. Although states are by definition situations with variables filled in, to get from one state to another mentally is quite conceptual. For one thing we need to mentally represent the two states and any others intermediate to them which may figure in bridging the distance between them. This mental simulation or modeling will more than likely acknowledge what variables are bound and the bindings themselves. This is conceptual. Furthermore, in order to find a way to connect states we must discover a set of conceptual transformations or transitions that can cross the "space" intervening between the start and the goal with an unbroken, step-by-step chain of links.

Without going into the various strategies for bridging the gap, we shall characterize the process as breaking the problem into subproblems, and as working either forward from the start state towards the goal state, working backward from the goal to the start, or a combination of both kinds of movement. In addition, the direction of motion within any one subproblem need not conform to that in any other subproblem nor to that in the overall problem.

The general scheme that we are describing is of breaking something into parts, or analysis, on the one hand, and of bridging things, requiring extension from one or both sides to cross the intervening space, on the other hand. This joining of parts is synthesis.

Decision making is one form of problem solving. We talked about two kinds of decision making. One features a single goal state where alternative paths all meet up. The other features diverging paths with no common goal. In both cases the deciding enters in given the set of paths that has been generated. The choosing of a best path based on relative cost vs. benefit, certainly a conceptual endeavor, is the true goal. The generation of alternative paths is itself a process that we just saw in problem solving as a bridging of states. Here, each path generation is a subproblem.

In the common-goal form of decision making, the goal state marks the endpoint of the "line" extending from the start state. This defines the space to be bridged and therefore analyzed into subspans or links. In the open-ended form of decision making, synthesis plays a major role, as there is no fixed goal state, just forward projection from the start state.

What basic operations are found in cost-benefit analysis? The synthesis of individual and relative measures of cost or benefit based on the analysis of the proposed solutions in the problem domain. There is comparison and matching, search of problem and solution spaces, of the problem domain knowledge base, and of the problem solving technique repertoire.

In planning, again we have a start state and a goal state. In this mode of operation we find the bridge connecting the endpoints, using bridging techniques. If there is a choice of bridging paths we use decision to narrow down the choices to one. We then take the path chosen, and the record of the step-by-step attainment of the goal state becomes a plan of action. Turning the sequence of steps into a plan may involve assigning appropriate resources to the realization of various steps.

Learning can be seen as a form of problem solving. In it we are confronted with a new fact or a new procedure. A new fact has a better chance of being understood and retained if it can be related to what we already understand. For example, if we're confronted with a new concept and its description contains many unfamiliar terms or one unfamiliar term whose own description contains unfamiliar terms, we will have more difficulty learning the concept than when the explanation is built on the secure foundation of terms that we are familiar with. Similarly with learning a new skill or procedure -- the more that it resembles something that we're already familiar with, the easier it is to learn. Or, the more familiar that individual aspects of the skill are, the better. Patrick Winston refers to this phenomenon as Martin's Law (Winston 1984, p.401).

What we're talking about is another example of bridging a gap. The closer that the new item is to the mass of what is already consolidated, the easier it is to assimilate it into that mass.

By the way, what about learning in animals? Why isn't it intelligence? Because it is learning of tasks requiring no conceptual processing. Does human rote learning of physical, non-conceptual skills constitute anything more? Not if we are trained to simply carry out some procedure without knowing why and it is not a task that works with concepts, but only percepts. If we are submitting to this training for some voluntary, purposeful reason, that's different. Of course, if there is some "method to the madness" on the part of the trainer, then our performing the task constitutes intelligent activity on their part.

Language allows communication. Communication of conceptual material is intelligent activity. Translation of meanings between representations is also intelligent activity. What about communication between animals? Even if they do not create their language, is its use to convey messages "intelligent" activity? Yes, in the same sense that evolutionary adaptation embodies "blind" intelligence. So are we still saying that this animal intelligence is basically different from our brand of intelligence? Yes. Messages can convey perceptual information alone, like "The lion is coming." Non-conceptual communication, where just programmed reaction rather than conceptual activity is initiated on the receiving end, is merely pre-intelligent activity.

Finally, understanding, the seeing of the meaning of something, is intelligent activity when the meaning is conceptual. Can meaning be other than conceptual? All concepts are meanings, but are all meanings conceptual? No. We just saw that a message can convey strictly perceptual information. That is the same thing as perceptual meaning. Transferring an image of blue to another is not intelligence, but it is communication. Very sophisticated communication can occur without intelligence in the sense that we've developed. Does this mean that our crowning glory, the intellect, may not be all that exalted in the scheme of things? There is still no substitute for knowing how things work, besides knowing of things, if one is interested in control of the environment or even if one is simply curious about nature and enjoys learning. There is much to be said, however, for a "simple" life, one of oneness with the world, where distance between observer and observed is not inserted in order to abstract properties.

The Eastern mystics chide the users of intellect, which artificially divides up reality in categorization. These mystics claim not only far greater contentment with life than those "enslaved" by intellect, through maintaining unity with the world, but insist that their knowledge of nature far transcends any that intellect is capable of. We can scoff at this, just as we might at their associated reports of very high-level perceptual communication and information gathering as in telepathy and clairvoyance, but are we justified in doing so? Perhaps in the course of thousands of years of culture, conceptual intelligence leads full circle to a simpler but more profound life, once material needs are assured. The story of the Garden of Eden comes to mind, where we lost paradise by tasting the fruit from the tree of knowledge of good and evil. On the other hand, what if a conquering Cortez comes along? Wouldn't it be good in that case to have "worldly" but invaluable knowledge of defense?

At any rate, in broad terms the mechanics of intelligence seem to involve the teamwork of analysis and synthesis. The common-goal kind of work characterizes "practical" activity; the open-ended kind, "creative" activity. Various skills like search and matching contribute to the task at many levels. Perhaps we will be able to systematically generate the conceptual operations from a small number of primitive proto-operations.

In the area of reasoning, the operations of deduction, induction, and truth maintenance, for example, form what we call logic. This is the set of rules of conduct in the realm of assertions and truth, the transformation and transition rules of the domain. Other domains, such as that of visualizing scenes or imagining sounds have their own transformation rules, their own logics. As long as there is conceptual work going on, whatever the realm may be, if we learn and use the rules of moving among its meanings and representations, then this is intelligence, and that is logic use.

Can we use our understanding of intelligence to look into creativity, intuition, and genius? We can try. We might think of creativity as creating new structures, whether concepts, terms, artifacts, or art. Or, as taking new approaches, in problem solving, or in doing anything really. We can characterize it as seeing ahead of the norm, toward the new.

In intuition we may see ahead, down a chain of concepts, but not be aware of the individual links in the chain. This may be more common than we realize. After all, we may just do our real thinking in meanings, not words. So a verbal, methodical approach to a problem may eventually plod to the answer that intuitive thinking streaks to, the plodder sceptical of the picturer's results because the meaning traverser may not be able to detail their thinking step-by-step in words. But both get to the right answer in their own way. Talking to oneself may help thinking by providing stimuli to seeing meanings, but may be overrated as a claimed method of thinking in its own right.

Is a genius using methods wholly alien to those that the rest of us use? Or are they just a very good concept handler? Perhaps they just see very far ahead. An intriguing thought is that concepts are there to be discovered at any time, like Plato's forms ever-present. It is just a matter then of pushing the mental sense onward as far as it will go. Of course some difference in technique may make a significant difference in whether a concept is reached or not. It is said that Einstein thought in pictures, simulating the logical extensions of ideas. If we all do our best thinking in nonverbal, "pictorial," imaginative form, perhaps he had just mastered the technique, trained himself as a disciplined mental gymnast.

Concepts are not exactly pictures but they do have logical shape. They have unique properties, the ways that they can combine with other concepts, the singular signature that they fit into the world with, like molecules. They are their set of defining properties. Which things can be allowable instances of a concept is governed by this set of properties as well.

Given the components of intelligence such as intellect, attention, memory, will, the logic inherent in the orderliness of the world, and the concepts that are out there to be found, discovery follows naturally. Have other species traveled that same road? The chances are not negligible that there are minds elsewhere, possibly far in advance of our own. As we have seen that there are degrees of intelligence, whether in single areas such as the range of search strategies, or problem-solving techniques, up to control strategies, to the whole ensemble of an organism's or system's concept skills, so there may be degrees of development well beyond what we are familiar with. But even right here on Earth there may be intelligent species besides our own.

Dolphins, elephants, and whales all have large brains, even in proportion to their body weights. These brains are highly convoluted as are ours, a sign of cerebral activity (see Lilly 1961, p.170). Are they intelligent? Many believe so. Why haven't they built cultures like we have or shown other indications of intelligence then? For one thing even human cultures vary widely as to what they consider impressive. For another these other species may indeed have cultures. We've all heard of the songs of the humpback whale. The elephant and dolphin have always been our friends. If dolphins, which have been studied in some depth, have a language, which is highly possible, then it might be founded on an entirely different basis from ours. From (Lilly 1978, pp.154-156): "Recent studies on the structure of the brains of cetaceans...lead to the conclusion that the basis for the postulated language of dolphins, "delphinese," is based upon the construction by central processing of "acoustic pictures," which are the basic elements of the postulated language." Acoustic pictures!

On the one hand this may point to a proliferation of "picturing" modes across intelligent species in the universe, based on the almost limitless number of senses that might have developed in evolving intelligent species, and yet on the other hand, there is a universal constancy in the realm of concepts. Just as mathematics can be considered a universal language, logic and concepts may be even more universal, being more fundamental, the basis of mathematics actually. Whatever road an intelligence may have taken to reach the level of conceptual thought, concepts, if general enough, in other words not too parochial and tied to the local intelligence's particulars of life, may be arrived at independently by any and all intelligent species. For example, although our cultures would be different and our senses different, we could all understand sensing, communication, understanding. It is analogous to the situation right here on Earth where there are so many cultures and languages and yet we can translate freely between languages by way of the deep structure, the semantic level, the conceptual content of expressions.

The line separating us from the animals may be blurring in spite of all of our attempts to find something we can do that they can't. New studies of our primate cousins are tearing down a last bastion of defense, that we make tools and use language. From (Diamond 1984, p.60): "For too long we neglected looking into ape societies for traits considered quintessentially human, like language and culture. The recent exciting discoveries of tool use, capacity for learning artificial languages, incest taboos, and invention of new behavior in chimpanzees or gorillas were made possible by changed attitudes and expectations as well as by patient observing. What else do apes do that we have not thought to look for?"

Some of the tools that we are making now or are contemplating need to be looked at in terms of developments in intelligence.

As we have made tools to extend ourselves in a myriad of ways to date, our ingenuity in machinery of the mind is only beginning to take off. Devices to help us cope with the exponential increase in our knowledge, and to aid us in adding to that increase, are being perfected. This work will shed new light on the nature of our own intelligence, enable us to be more intellectually productive, and produce a whole new species of intelligent entities with artificial intelligence. Sentience is one thing -- sapience is quite another. There will surely be "consciousness" and sensing in our creations, but will either of us show wisdom?

Now that we have traveled from outer properties through surface properties and into the interior of intelligence and have even seen some of its inner workings, we are ready to come back out and look again at what AI is.

Is there really intelligence in AI programs, and if so, where? If AI programs perform any conceptual operations then there is intelligence. Of course, as with us, intelligence is only potential unless it is being used. Having it and using it are two different things. AI software has intelligence even if some nonconceptual operation is performed in a conceptual context.

As it turns out, manipulating conceptual data requires a conceptual operation. Does this complicate things? No, it simplifies things. We can use the presence of conceptual data as a tipoff to the presence of intelligence, but we need only think in terms of conceptual operations (whether on conceptual data or on nonconceptual data) or of a conceptual context for operations (whether conceptual operations or nonconceptual operations).

But isn't this a broad qualification? Doesn't it let in things like traditional data processing, even word processing? Yes. Traditional DP is conceptual processing. If we have data in records which have fields that contain values, this is conceptual. A field is nothing more than a property of the object that is represented by the record. The field is a variable; its value is its instantiation. The difference between traditional DP and AI is not so much one of kind but one of degree. In AI the degree of explicitness of intelligence is greater.

Using word processing technology constitutes using intelligence. Using it to capture or transfer conceptual material is a conceptual operation. The difference between this and computer processing is that this is a manual operation -- computing is unattended conceptual processing. Photocopying, although a nonconceptual process itself, simply blindly copying symbols for example, which are not concepts but merely stand for them, becomes conceptual when we are doing it for a meaningful reason, for example to help in disseminating information.

So the difference is in the explicitness of the intelligence. It is also in the degree of intelligence used, very often. Those in traditional DP can employ sophisticated search strategies, for example, but may not be as aware of such considerations as their AI counterparts because AI has focused more on intelligent techniques themselves. This is part of the explicitness difference.

In AI there is explicit interest in conceptual data, whether in the form of declarative statements containing the conceptual knowledge in a domain, production rules representing domain deduction, or heuristics, representing domain or technique methodology. There is explicit interest in conceptual operations and techniques, such as problem solving and search. There is explicit interest in conceptual subject matter, where applications explicitly concern understanding human and improving machine intelligence. There is also a transcendent level, a conceptual context, in the control protocol and metaknowledge of a knowledge system.

We can begin to categorize the various topics studied in AI in terms of the frame of reference of our concept of intelligence. Of course we have omitted such things as emotions, personality, the subconscious, conscience, and the spirit from our description of intelligence, but we were not trying to define a human being, only intelligence.

The work in motor skills and dexterity might not truly qualify as study of intelligence nor might even studies of sensory activity such as vision, hearing, and touch except where these are involved in conceptual contact with the world on input or output. Even the study of language itself, unless concerned with the semantic, conceptual, abstract level of activity, is only indirectly involved with intelligence. Only when we turn to the study of knowledge representation, logic and reasoning, problem solving, decision making, understanding and conceptual learning, and other intellectual, conceptual activities are we directly addressing intelligence in operation. A finer classification can be made by tying particular topics to specific components of the architecture of intelligence, for example memory search and retrieval, system control and span of attention, fact manipulation, and so on, down to increasing levels of detail.

At this point it will be enlightening to apply our understanding of intelligence in general, and of AI in particular, to the claims and questions of an observer of the AI field as related in a recent article in *Computerworld* reproduced in full as follows:

"FRAMINGHAM, Mass. --- Artificial intelligence (n): the capability of a machine to imitate intelligent human behavior.

"The above explanation, at least according to Webster's Ninth New Collegiate Dictionary, constitutes the strict definition of artificial intelligence. But as with any evolving and, unfortunately, overused high-technology buzzword, the term 'AI' currently sports a widening variety of confusing and sometimes contradictory definitions.

"'If you ask a number of people involved with AI for a definition of exactly what it is, you will probably get n-plus-one answers,' Henry Eric Firdman explained recently to a gathering of computer industry and business executives. 'There exists a lot of controversy and confusion about artificial intelligence today.'

"Firdman navigated the executives through a seminar here, sponsored by Rensselaer Polytechnic Institute, titled 'Artificial Intelligence and Expert Systems: An Executive Briefing.' Firdman is the president of Henry Firdman & Associates, an AI contracting and consulting firm.

"What, then, does the term 'artificial intelligence' really mean?

"'I define an AI system as one that can do three things,' Firdman said. 'An AI system can acquire new knowledge by communicating to the external world and by inferring from its current knowledge. That communication could be facilitated through sensors, input devices or even another AI system.

"'Second, an AI system can carry out goal-directed behavior if it is given a set of goals, strategies for achieving those goals and some adequate knowledge. Finally, it exhibits skill acquisition, meaning that it can improve its performance in certain problem domains based on its knowledge and previous goal-directed behavior experiences.'

"Firdman said current research undertaken in academia and industry has produced AI systems that exhibit partial capabilities in only one of those critical areas. But to date, he said, no system exists that can boast success in all three realms.

"'If a system is truly an AI system,' Firdman said, 'it will be significantly different six months after you buy it. Thus, any AI product must, by its very nature, be sold as an unfinished product, so to speak.'

"To aid his listeners in understanding artificial intelligence, Firdman unraveled some of what he labeled the popular but incorrect definitions of AI. For example, he said, many researchers describe AI as the development of computer systems capable of mimicking human reasoning and perception.

"'What is wrong with that definition,' Firdman said, 'is that we do not know much about the mechanisms of human reasoning and perception. In that light, AI is being explained through the unknown.'

"Others characterize AI as the processing of symbolic information such as concepts, knowledge, and relations. 'That definition lacks functionality,' Firdman claimed. 'It is also weak and overly broad. For example, word processing would fall into that definition because it is symbolic information processing. Also, concepts and relations are just particular cases of knowledge, and by no means all that AI includes.'

"A third camp, Firdman explained, defines AI as the branch of computer science concerned with designing systems that exhibit the characteristics associated with intelligence in human behavior, such as understanding language, learning, reasoning and solving problems.

"'What exactly does that mean?' Firdman queried. 'Is solving partial differential equations an intellectual activity? Is human vision an intellectual one? What about driving a car?'"

Mr. Firdman's three requirements can be restated as 1) declarative knowledge acquisition, or "learning that," 2) problem solving, and 3) procedural knowledge acquisition, or "learning how to." These are worthy attributes but two points must be made. First, there are learning systems satisfying all three conditions (see Michalski, Carbonell, and Mitchell 1983, esp. ch. 9). Second, learning is not a requirement for intelligent activity! It is not necessary to learn anything in order to function conceptually, given an initial set of facts and procedures. Of course learning is intimately associated with intelligence, but animals learn perceptual information and can even be taught procedural knowledge, as in dog training.

As to Mr. Firdman's criticisms of the three types of definitions of AI, let's take each definition and criticism in turn.

1) mimicking human reasoning and perception: this is a good definition as long as we qualify perception as perception of concepts. Not knowing much about the mechanics of reasoning and perception does not mean that we can't still recognize them in operation.

2) processing of symbolic information such as concepts, knowledge, and relations. Yes! Of course we'll qualify knowledge as conceptual rather than perceptual. And, relations are one kind of concept. Mr. Firdman complains that this definition lacks functionality, is weak, and is overly broad. It is quite functional if you know what concepts, knowledge, and relations are. It is only weak or too broad if you are attempting to make fine distinctions which an understanding of the internals of concepts, knowledge, and relations would be more appropriate for. The criticism continues by implying that word processing is not intelligent activity. As we have seen, it is when used for any conceptual work. That concepts and relations are just particular cases of knowledge is certainly true, but the other kinds of knowledge, raw experience and perceptual information, a la Fodor's transducers and input systems (Fodor 1983, p.42), are not the data of intelligent processing, only one source of material from which to extract patterns and concepts. If we distinguish "data" from "information," raw electromagnetic signals such as photon fluxes are data for signal transducers, and perception of objects in a setting is information from an input system. We then still need to distinguish "knowledge," or better, "conceptual knowledge," as the product of cognitive mechanisms in central processors from the information output of input systems. (One thing that seems to be neglected is that there is an internal input system besides those for the five senses. Logical objects are as much a part of our scenery as are external physical objects. We can apply similar abstraction, property extraction and processing operations to them as we can with external objects. The behaviorists deny mind by invoking behavior as real but forget that behavior includes that which one is engaged in, that which attention is on, and when we're engaged in mental activity, that's behavior.) Finally, if such knowledge processing is by no means all that AI includes, perhaps there are things that are included that rightly belong in other fields. AI incorporates expertise from a wide variety of disciplines; it is a melting pot. But some of the subjects that come under its umbrella are foster children, sheltered there possibly because they are closely related to it and partly because their proper fields are not well-established enough to stand on their own. That may change with time.

3) the branch of computer science concerned with designing systems that exhibit the characteristics associated with intelligence in human behavior, such as understanding language, learning, reasoning and solving problems. Yes! Again we'll qualify learning as conceptual learning. Mr. Firdman asks if solving partial differential equations is an intellectual activity. Definitely. Is human vision? No. Driving a car? Yes.

One final comment. Building an artificial person is not the same thing as building an artificial intelligence. Robotics is AI only insofar as conceptual concerns enter into it. Likewise with any other topic.

How do we decide if a gadget with a microprocessor in it is really "smart" or whether a software package incorporates AI or not? We saw that even if realtime adjustment is neither conceptual in subject matter nor in type of process performed, if it helps to carry out an abstract assignment, operating in a conceptual context, then intelligence is at least implicitly present. Either way, a gadget with any logic, i.e. conditional response to input, is at least rudimentarily intelligent. It is comparable to an animal, with built-in programming, "ROM." The more sophisticated that that programming is, the more "intelligent" it can be considered. One index of sophistication would be the subtlety of conditions that the logic is sensitive to. Unless abstraction or deduction, analysis or synthesis occur, however, the device itself is not "humanly" intelligent.

Similarly, any software package is at least minimally intelligent, but unless it includes conceptual operations it is not artificially "humanly" intelligent. Even within the class of programs performing conceptual operations there is a wide range of degrees of sophistication. Some might specialize in one or another kind of abstract activity, some in multiple conceptual skills, some even in many. The level of expertise in one area falls on a scale of sophistication, as does that across areas. The synthetic sum of these abilities might give a true indication of the "IQ" of a program. So the question today is not so much one of whether or not an AI package really incorporates intelligence, but is rather one of how much intelligence of which kinds, and across kinds, does it contain.

We are now ready to conclude our journey by returning as promised to our starting point, the dictionary definitions of intelligence. We'll just go right through the list we saw before, this time making pertinent remarks:

8-1a1a: the ability to learn -- Learn concepts.

8-1a1b: the ability to understand -- Yes.

8-1a1c: the ability to deal with new or trying situations -- An application, external property.

8-1a1d: REASON -- Yes.

8-1a1e: the skilled use of reason -- Degree of intelligence.

8-1a2a: the ability to apply knowledge to manipulate one's environment -- Conceptual knowledge, application.

8-1a2b: the ability to think abstractly as measured by objective criteria (as tests) -- As we said before, if we stop after the word "abstractly" then YES!

8-1c1 : mental acuteness -- Degree of intelligence.

8-1c2 : SHREWDNESS -- Application.

8-3a : the act of understanding -- Still obscure.

8-3b : COMPREHENSION -- Yes.

7-1a1a: the capacity to recognize the meaning of facts, propositions and their relations and to reason about them -- Yes.

7-1a1b: REASON -- As above.

7-1a1c: INTELLECT -- Very much so.

7-1a1d: the use or exercise of the intellect esp. when carried on with considerable ability -- Degree of intelligence.

So, in trying to define AI and thereby needing to define intelligence, we have seen four themes -- pure intelligence, applied intelligence, degree of intelligence, and explicitness of intelligence. We have seen three levels of properties -- external, surface, and internal. We have seen two kinds of conceptual activity -- explicit and implicit, conceptual operations and/or a conceptual context. We can therefore conclude our search for a definition of intelligence by calling it conceptual activity, or abstract activity. Implied in these expressions is that the activity can either be explicitly conceptual or in effect conceptual by being performed in a conceptual context. Similarly we can consider our quest for a definition of artificial intelligence completed by thinking of AI as artificial conceptual, or abstract, activity. Primarily this would cover conceptual activity in devices, but the definition could be stretched to embrace devices that aid or extend our own conceptual activity, or that help us to better understand or improve our conceptual ability.

What does the future hold in store? At least an increase in our ability to deal with the deluge of information confronting us, more and smarter tools for practical and creative purposes, a difference in the degree of our conceptual expertise. Possibly, though, a revolution in human culture is beginning. Soon we may see the implementation of a knowledge-based system for all human knowledge, facilities for enhancing our conceptual faculties far beyond anything we have even dreamed of let alone seen before, and generally not just a difference of degree in the way we live, but a difference of kind.

REFERENCES

Computerworld (1984) "Valid AI Definition Lost Amid Industry Jargon" 18:48, 28.

Cooper, Lynn A., and Roger N. Shepard (1984) "Turning Something Over in the Mind," in *Scientific American* 251:6, 106-114.

Diamond, Jared (1984) "Making a Chimp out of Man," in *Discover* 5:12, 55-60.

Fodor, Jerry A. (1983) *The Modularity of Mind*, The MIT Press, Cambridge, MA.

Harari, Haim (1983) "The Structure of Quarks and Leptons," in *Scientific American* 248:4, 56-68.

Lilly, John C., M.D. (1961) *Man and Dolphin*, Pyramid Books, New York.

Lilly, John C., M.D. (1973) *Programming and Metaprogramming in the Human Biocomputer*, Bantam Books, New York.

Lilly, John C., M.D. (1978) *Communication Between Man and Dolphin*, Crown, New York.

Michalski, Ryszard S., Jaime G. Carbonell, and Tom M. Mitchell (1983) *Machine Learning*, Tioga, Palo Alto.

Sheldrake, Rupert (1981) *A New Science of Life*, J.P. Tarcher, Inc., Los Angeles.

Sowa, John F. (1984) *Conceptual Structures*, Addison-Wesley, Reading, MA.

Webster's New Collegiate Dictionary, Eighth Edition (1980) G. & C. Merriam Co., Springfield, MA.

Webster's Seventh New Collegiate Dictionary (1967) G. & C. Merriam Co., Springfield, MA.

Winston, Patrick H. (1984) *Artificial Intelligence* (second edition), Addison-Wesley, Reading, MA.