To: Licklider @ ISI
From: Newell @ CMU-10A
Date: 16 Apr 75
Re: AI ROAD MAP MEETING
CC: Carlstrom @ ISI, Fields @ ISI, Russell @ ISI, Simon @ CMU-10A
    JMC @ SU-AI, LES @ SU-AI, CCG @ SU-AI, PHW @ MIT-AI,
    Feigenbaum @ ISI, Amarel @ ISI, Nilsson @ SRI-AI,
    Sacerdoti @ SRI-AI

Lick: Twenty-four hours has permitted some reflections on
Monday's session and the problems attendent thereto.

(1) To repeat what was generally accepted, implicitly and
explicitly, by the AI scientists at the meeting: It is
appropriate, given the current general context and specific ARPA
context, for the AI field to attempt a series of applications.
Such applications would be a good thing for AI generally, as
well as for its specific relations with DoD.

(2) Also to repeat: Application opportunities must be
discovered, verified and exploited.   Some institutional means
must be found to expedite this. For it is clear that the AI
community by itself does not have the expertise nor the
connections to find high payoff activities.   Dave Russell, at
the end of the day, strongly posited a mechanism of a Rand-like
agency with the mission of finding application opportunities,
verifying them, building a bridge to the AI Labs, etc.   There
would also exist, as an adjunct to this, an AI Applications
Technical Group (or some such title), consisting of
representatives of the various Labs, CMU, MIT, SRI-AI, SU-AI,
SU-HP, plus maybe others in related programs, such as BBN-SUS,
SDC-SUS, Amarel, etc. This group would be a prime forum and
initiation point for these applications.   I expressed some
concern that such an agent could come into being in short enough
order to satisfy the needs of the day (implying that some
temporary vehicle would have to be erected), but Russell seemed
confident that such expedients were unnecessary. It would be
better his way.

(3) It is extremely important to be sure that the payoffs of a
specific application are real.   It is too easy to get
sandbagged to have a seeming application turn to dross. Given
that ARPA is prepared to spend large fractions of its AI
community (a relatively precious resource) on producing some
specific applications, it is critical to substantiate the need
and acceptability of an application.   ARPA itself, though
inside the DoD and much closer to the application sites than the
AI community, does not itself have the expertise and,
importantly, the time to examine the situations enough to make

reliable assessements.

(4) To be concrete, on reflection I am not at all convinced that
the items on Heilmeier's list are all really in the category of
genuine application opportunities.   I cannot speak about the
ASW problem, since that is not a single problem, but many --
namely, the question of where to apply AI techniques all up and
down a complex system. But the other two are bitty problems
aimed at highly specific targets. I did not get any sense that
APRA really knew in detail whether the payoffs were real or
simply ephemeral, momentary opinions of one or two high level
people in the organizations connected with the applications.
Heilmeier's carriage-trade philosophy requires a really good
marketing and marketing research arm if it is to succeed. It
will do ARPA no good if it squanders its substance on a bunch of
irrelevant mirages.   The military scene is littered with the
dead bones of expensive solutions which were only monuments to
someone's folly.

(5) It appears that MIT-MAC is already spending substantial sums
(relative to the apparent size of the problem) on the Morse-code
problem. Surely it would not be fruitful for the AI community
to get further involved in that one. I am not quite sure why
this one showed up on Heilmeier's list, but maybe it was just
meant to be illustrative of what ARPA had no decided to do.

(6) On the language-spotting task, I need to reiterate what I
said at the meeting.   First, I do not think the AI
Labs (in counterdistinction to the SUS Labs) should take on the
problem, they simply would have to build up much of the
expertise that the SUS Labs have, which would be a genuine
duplication. Second, I consider that the SURG has not been
asked to consider taking on that task. I agree that you possibly
mentioned it to the SUSC (though I have forgotten it
completely), but since it was not brought up in the context of a
serious confrontation with the SUS 5-year goals, I do not take
it as a serious proposal. It would have been derilict of any SUS
group to take on the task, given how tightly the SUS program is
strapped down to the 5 year goals.

I do believe we can consider this task and, if it is important
enough, we can consider folding it into the present contractors
in some way. But we do have to face the potential effect on the
5-year goals and to see how to work around them. As chairman of
the SUSC, I am quite willing to go around on that issue, but I
need a signal from you or Dave Carlstom that indicates you want
to do that. This is a serious point and proposal I am making,
since I do not want to be accused later of having fiddled while
Rome burned -- of having not picked up this problem when it was
important to do so.   I would like some feedback on this
specific matter.   The problem itself seems rather
straightforward, given the current art. I would estimate half a
man year for the technical work, if done at CMU given all the
facilities. To this must be added the whole custumer interface,
which might be as much again. Much depends on details which I do
not know, of course (eg, how must it be packaged and how much
must it cost). If the problem could wait until after Nov76, you

could surely get it taken on by the SUS Labs if it were as
important as Heilmeier stipulates.

(7) Another example of a SUS-related application is the Korean
on-line communication aid, which you raised as a problem and
which I suggested a solution approach to some time ago. I do
not know what became of that. Again, it would be derilict for
the SUS community to fold that in without at least explicitly
facing the 5 year goals. All this stem, of course, from the fact
that we (read: ARPA and possibly Newell-cum-SURG-initiator)
wanted a program tightly fixed on impressive goals, and
therefore not with much slack for such things. Again, I am
willing to consider this.

(8) I cannot believe that the CBC is not on the track of an
important application. It has two things wrong with it: (1)
there is not an immediate customer eager and ready to pay; (2)
notwithstanding SRI's search, there does not seem to be one
hiding out there quite yet (though I do not know how intensive
that search was or is). Yet, it does not seem to me profitable
-- for ARPA, even on its own current terms -- to jerk that
effort up at the roots and radically redirect it. Rather it
seems to me critical to widen the scope to "Real-time operations
consulting" (namely, how to help someone carry out an operation
in real time) and to search for applications within this wider
sphere. The core of work on the CBC remains in fact strongly
relevant; and the new applications can be grafted on.

(9) What should go in the Road Map? It seems clear to me that
the Road Map for Friday has its action component defined
independently of its substantive component. To wit, the
formation of the application-finding mechanism, defined above,
will not be justified, nor require justification, from the a
statement of the current art or a statement of future scientific
goals. However, this application proposal will differ from all
other such attempts by the promise, implicitly extracted at the
meeting, by the AI Labs to enter into such an application-search
wholeheartedly.

The substantive mode must perhaps still be there by Saturday.
You don't have much to work with, in terms of what was generated
before and during the meeting. Thus, I would attempt to get the
action component to stand in for the rest. Let me discuss each
of the substantive components a little, and then come back to
this.

(10) When a set of the worlds best scientists, being asked about
the their very own scientific domain, becomes tongue-tied and
produces answers unsatisfactory in a first year qualifier, then
the conclusion is not that the science doesn't exist, it is that
the question was posed wrongly or the situation inhibited
adequate response. You asked us, I think, to do something
under constraints that communicated: (1) that none of our prior
attempts was to be considered satisfactory -- that something new
and different was required; (2) that we adopt a form of
specification of results and of expectations that is foreign to
CS and AI, and largely foreign to science (namely, to state in

advance the content of the scientific results to be expected up
to several years in the future, so that the questions would only
be whether or when the result would be attained. This pre-empts
the science and leaves us tongue-tied).

For instance, in the CMU proposal I have just finished writing a
statement about the basic scientific questions of AI and the
high level propositions that characterize what we have found out
in AI. Apparently that is to be discarded as not adequate or
appropriate to the task -- and I am to find yet another
statement, different from that, that is to be adequate to the
new (yet identical) task.

For instance, Nils has just finished writing a paper (IFIPS 74)
devoted to a summary of what AI has done and what areas it has
worked in. Apparently that is to be discarded as not adequate to
the task -- and Nils is to find yet another statment, different
from that, that is to be adequate to the new (yet idenical)
task.

Let me strongly suggest, for instance, that as far as
characterizing the present state is concerned, you take a copy
of Nils IFIPS paper and underline in red the items in the
bibliography that are done in the ARPA AI Labs, and on the many
charts that draw a map of the area, circle in red these same
items. This will give (1) a direct picture of the coverage and
scope of the field of AI that ARPA has given birth to; and (2) a
direct picture of the extent to which ARPA is responsible for
these results and for the important ones.

Let me further suggest that you put in front of Nils paper the
first section on AI goals from the CMU proposal, as giving a
high level coherent picture of what AI as a science is striving
to achieve and what in global terms it has found out.

These two items answer only the question: What are the results
in AI in its own scientific terms. They do not answer it fully,
but they will do as well as what you can put together in yet one
more attempt in a few hours.

(12) I do believe that several additional descriptions of AI
scientific results are possible that will appear to be more
satisfactory to upper ARPA than the two above items (if, indeed,
anything is satisfactory). I cannot carry out these
descriptions in the time available, indeed I think it would take
a couple of months of very hard work (maybe more). But I can
sketch and illustrate one part of it (which is indeed based on
past efforts to systemitize).

Progress in AI proceeds in terms of increases in scientific
knowledge about the various components of the intelligent agent,
components that are defined functionally. A standard division,
which corresponds in part to Nilsson's core areas, is:

> Recognition and description (Perception)
  > Vision
  > Speech

> Language
> Representation
> Problem Solving Methods
> Control Structure
> Assimilation & Accomodation (Learning)

Within each component one can describe a series of structures (or mechanisms) that are possibilities for this component. The discovery of each such structure and mechanism is an advance for AI and a result. Verification, of course, is required; it comes, usually, from incoroporation in several total systems. Knowledge about each mechanism grows with experimentation and theoretical sharpening. Such knowledge, again when verified experimentally, constitutes scientific results for AI.   It consists mostly of statements of adequacy or sufficiency in specific task environments.

Thus, the statement "what are the results of AI" at a given date is a listing of the various mechanisms (usually described by conventional technical names), plus the associated statements of adequacy. This list grows over time, and it, rather than a parametrization of how good are the systems that can be produced constitutes the core transferable knowledge of AI. This core is indeed transferrable, precisely because it consists of the abstracted mechanisms which have been shown experimentally to be useful in several task environments.

I cannot produce the lists of results for the total field, mostly because they have not been extracted, labelled and organized in this way.   I can do it for one subpart, that of problem solving methods. Here, much that we know can be given by specific methods (analagous to the methods of numerical analysis). A fairly good list is:

> Generate and test
> Hill climbing
> Heuristic search
   > Search stragegy:
      > Depth first, Breadth first, Best first,
        Progressive Deepening
   > Evaluation
      > Evaluation functions, level of aspiration,
        duplication avoidance, external limits
> Matching
> Hypothosize and match
> Means ends analysis
> Substitute & eliminate
> Range restriction
> Abstraction planning

To find a short way to say what we know, eg, about Hill Climbing, takes more energy than I have at this wee hour. We do know the major things to beware of (Multi-modality, Mesas, Ridges, Cliffs), we do have some empirical things to say about when Hill Climbing seems to work and when it doesn't. We do have a way of classifying the refinements of the method (as simplifed models of the hill, which are used to predict the

optimum hill-step to take). And so on.

The existence of this list implies a large kit of tools
available to be used in applications, and indeed, when Ed
Feigenbaum says they used "standard AI" in Dendral, he means
that the techniques in Dendral pretty much are drawn from the
list of such own and characterized methods.

To plot the success of the Problem Solving Methods component
over time is to watch this list grow and/or the amount of
knowledge about each such component grow.

I do not know how much each of the other components can be so
characterized, though I expect it could be pushed quite far. But
not tonight!

(13) Future goals of AI must be in general to extend the
mechanisms and structures of each area and to show that they are
adequate to wider and more difficult problems. The discover of
a new method is not to be asserted in advance as a goal, if it
could then the new method would have been found! Sometimes one
knows enough about a method or structure to specify as a goal
that it is to be explored. To find the true scope of the range
restriction method, is such a (small) goal.

Goals, in the sense that you (read: ARPA) want them, are only to
be associated with systems or with instruments (Physics has such
goals for the energy of interactions its accelerators will
reach, or the resolution of microscopes). There has not been
any difficulty, as far as I know, in determining how to
parametrize the structure and performance of specific narrow
classes of systems (such as SUSs or Dendral-like systems) when
the task comes close enough to feasibility to make it worth
adopting a system as an AI goal. It can surely be done for
various other specific classes of systems, though it will not
cover, thereby, all of AI's goals.

[I have run out of gas here -- I believe more can be said about
stating AI pure goals, but it just escapes my fog-bound mind. I
move on.]

(14) Applied goals, I believe are to be stated in one of two
ways. These provide other ways of describing what AI has done
in terms of how it can contribute to these goals.

One way is how we started out to do it at the meeting. An
applied system is posited (ie, a total military system, such as
an ASW system). Then, within that some points of AI application
are found, from which one attempts to derive the AI that might
make a difference. This is a form of systems analysis, and one
that can lead to a backward chain of available relevant research
results and, as well, of still needed research -- methods,
structures, knowledge, experimentation, , etc. needed to do the
job. I think we should do a substantial amount of this, and I
believe quite satisfactory road-map results would come out of
it.   Unfortuneately, I believe that the effort per complex
total system must be a summer study-group sort of thing, ie,

about what we put into the SUS initial report, But this is
exactly what is to be done by this Rand-like agency (plus some
of us).

The second way is to specify applied technologies, The two
prototypic examples at hand are the notion of a SUS technology
and (more pertinent) the natural language front-end technology
that we all were talking about. One can take the development of
such a technology as a goal and describe both what existing AI
already provides and what new research is needed to get it.
This can be carried out much more within the AI community,
though some sense for what is really required to make specific
applications go is important, But again, it takes a fair sized
effort to lay out such a technological alternative, We could
commission such explorations. (It would depend, I guess, on
upper ARPA being prepared to consider such expansions.)

(15) We did assert rather strongly that there have been a number
of civilian applications of AI, eg, in management science, in
design, etc. Ferreting these out and asking whether any of them
could be applied to military systems would be an additional
important task for this applications organization (along with
some of us), This, of course, is yet one other way of stating AI
results.

****

I am really to the end of my rope tonight (this morning) and I
will send this out after Herb gets a chance to look it over, I
am willing to work on expanding or modifying any piece of this,
I am pretty much around from here through the weekend.

A.N,

P.S, Recall that I am expecting feedback on the Language
Spotting issue.