Have you ever had breakfast with Sophia Loren?

David Moss

First published in 2003
by Business Consultancy Services Ltd
bcsl@blueyonder.co.uk
Reprinted in 2004

for Jane

Have you ever had breakfast with Sophia Loren?

an intelligent artifice by

David Moss

Dramatis Personae

Chorus, *a girl*
Nino, *an Inquisitor*
H&lə, *a Handler*
Hubot, Psybot, Lobot, Probot } *Agents*

Typography

1. The comments in brackets are sometimes stage directions and sometimes references.

2. The reflected "e" in the handler's name is the letter which George Bernard Shaw suggested should be added to the alphabet to denote the "er" sound in English. Or so I was told but, when I checked, it transpired that he did no such thing (*Farlex*).

ACT I – TURING, PSYCHOLOGY AND LOGIC

Prologue

Enter Chorus with a flipchart

Chorus In the world of the intelligence services
all agents have a Handler.

Our Handler is called H&lə
and is a surveyor
in the field
of human intelligence.

He has adopted as his special research project
developments in the study
of what humans obtusely call
"artificial intelligence" or simply "AI" *(Remagnino 2003)*
and is interested in the work of Russell and Norvig *(2003)*.

Russell and Norvig have constructed a topology
of the subject of our play ... *(with a flourish)*
... like so!
(reveals the first page of the flipchart
showing the following diagram)

Russell and Norvig:

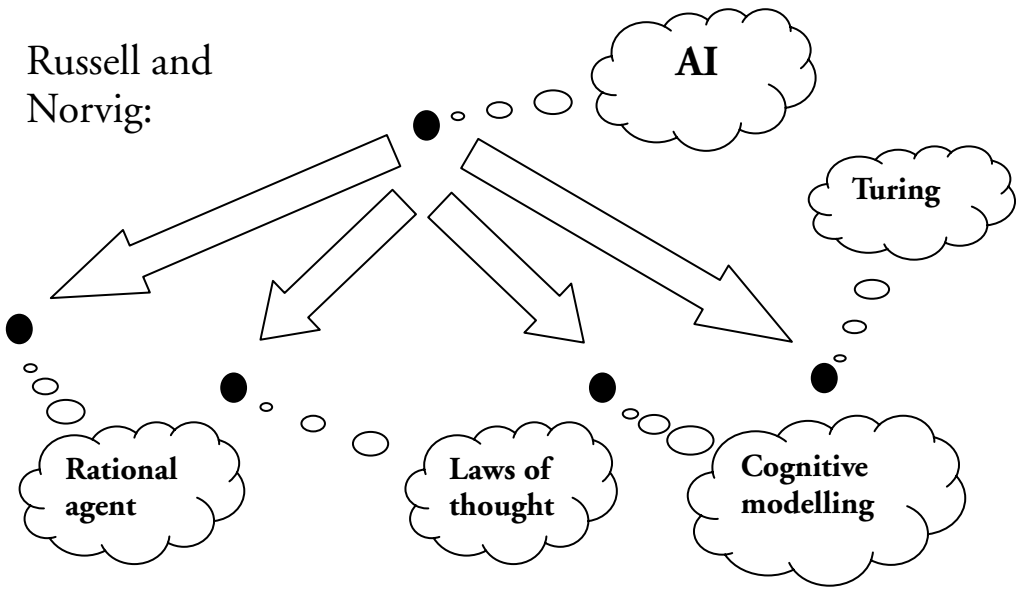

They believe that there are these four paths for AI
and that the most rewarding
is the path signed "rational agent" *(op. cit., pp.2, 28)*.

H&lə has directed four Agents
to go forth in great arcs
to explore this network,
to expand and follow its nodes,
and now in his laboratory
receives from each Agent
the report of his[1] findings. *(exit Chorus)*

Scene 1, a laboratory somewhere in gridspace – Hubot 1

Enter H&lə with a small creature you might at first take to be a man

H&lə How did you get on?

Hubot One of the easier tasks you have given me,
thank you, Handler.

H&lə You amuse me, Hubot.

Hubot My apologies, Handler.

H&lə You must beware of amusing, Hubot,
particularly if you encounter our Inquisitor,
who has bypassed his Amused() function.
The input argument is still there, however,
and since it cannot be responded to as amusement
and since it has been decreed
that no percept may go unprocessed
it defaults through the error-handling hierarchy to InPain().
And an Inquisitor in pain

---

[1] Agents, of course, do not have a sex. Humans, however, are incapable of interacting with sexless objects and so impute one to them. *Vide* particularly female humans, Citroën 2CVs, pet names for.

is a mighty and awful process.
Quite simply, he would eat you for breakfast.
He has done it before
with one of the old emetic class of agents,
Special K,
and he would do it again without Compunction(),
which has also been bypassed
for performance reasons.

Hubot But, Handler, Amused() is a five million-line function!
How do you step over five million lines of code?

H&lə I ask the questions, Hubot,
not you,
but, since I happen to know the answer
and I am responsible for your knowledgebase,
on this occasion
I shall tell you –
very simply, is the answer,
and, in the case of our Inquisitor,
with his customary terseness.
He inserted a new line 1,

```
Amused = Null: Exit Function
```

It never ceases to amaze me
how little you programmed constructs
know of programming.
Now, get on with your report,
while preferably neither amusing me nor amazing me.

Hubot Executive Summary: ...

H&lə No, Hubot, no,
I am not yet an Executive
still a mere Handler.

Hubot Handler Summary:
Turing's test *(1950)*

is still accepted by the humans as definitive.
If a human interacting with a machine
in an unrehearsed scene
cannot tell that it is a machine
then that machine must have artificial intelligence.
So far there is no sign of the test being passed.
It is unlikely that it ever will be
for reasons to do with emotion.
That does not matter,
as Russell and Norvig point out *(op. cit., pp.2-3)*.

H&lə Is that it?

Hubot Yes, Handler.
As I said,
it was one of my easier assignments.

H&lə Easy, eh?
Hmm.
Let us take these points in reverse order, Hubot.
Russell and Norvig may not think that it matters,
that is their prerogative,
but I do.
In fact, I believe that the survival of the human race
depends upon it.
What do you make of that, Hubot?

Hubot Frankly, Handler, not a lot.

H&lə Then I shall develop my theme.
You call on emotion
as apparently distinctive of the human race
an essential component of it
always missing from artificial beings.
What is an emotion, Hubot?

Hubot A mental state, Handler,
suffered by humans,

in which their already small logical powers
are derailed by feelings.
I would call an emotion
the suspension of rationality.

H&lə Many people would agree with you, Hubot.
As you may suspect, I am not one of them.

Let us assume that feelings are not susceptible to reason.
I speak here of matters of taste.
Some men dislike broccoli.
You may tell them that it is healthy to eat broccoli,
that they wish to be healthy
and so they should like it.
But they won't.
Utility won't cut the mustard.
You can't argue someone into liking broccoli, Hubot,
and stop smirking,
there is a point to all this.
That dislike is a feeling, not an emotion.

Emotions are a mixture of feeling and reason *(Kenny)*.
A publican, for example, may feel contempt
for an old man who comes into the bar every night,
drinks too much,
talks to himself
and lurches home alone.
Contempt is an emotion.

Emotions, notice, are always object-directed.
The publican does not just feel contempt
in some generalised way,
he feels contempt *for* something.
He may feel jealous *about* something
or be in love *with* someone.
Don't giggle, Hubot.

Now, suppose that the publican learns

that this old man's wife has died
and his brother never visits him
and he was blinded in one eye
rescuing a child
from certain death
at the hooves of a runaway horse.
Once this evidence is adduced,
the publican may start to feel sympathy for the old man
and, quite reasonably, cease to feel contempt.

You see the mixed use of feeling and reason
throughout my example, Hubot?

Hubot Yes, Handler,
and, with some reason,
I have mixed feelings about it.

H&lə Facetious, Hubot, facetious.
Now, what else do we know about emotions?
In your case,
you, who think that emotion is the suspense of reason,
not a lot.
But I can tell you
that emotions have a job,
they are motives for action.
Russell and Norvig agree
that action is important.
Actions may be prompted by motives
or reasons or causes,
there may be excuses for them or justifications ...

Hubot *(interrupting)* I saw excuses on my search, Handler.
These humans have confessional spaces
in which intentions are weighed
according to some dynamics I couldn't fathom.
What do you have to add to a reason
to make it an excuse?
I just couldn't get the hang of it.

H&lɔ *(incredulous)* Hubot!
You *were* awake!
You *did* have your eyes open!
Now let me finish.
Russell and Norvig like a good graph, don't they?
Imagine a simple two-dimensional graph
with feeling measured on the y-axis and reason on the x-axis.
Now plot y=1/x so that, in the NE quadrant, at least,
when x is small, y is big and *vice versa*.
(turns to second page on flipchart)

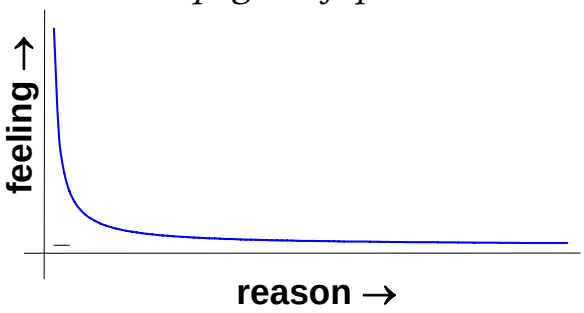

The function approaches the y-axis asymptotically
reaching y=∞ when x=0
and it approaches the x-axis asymptotically
reaching y=0 when x=∞.

You will notice, Hubot,
that this graph has no title.
That is because I do not know what I have plotted.
All I know is that I want an inverse function of some sort
relating feeling and reason in some way.

I quite like the left hand side,
which seems to me to represent graphically
the fact that humans have
an infinite amount of *a priori* knowledge
in a very narrow area.
They have *a priori* knowledge
of their posture, for example.
It requires little or no intellectual effort
to know that they have their legs crossed,

for example,
while they are sitting down.

There are certain things I can do with this graph.
I can overlay it with mental states, for example:
(turns to third page on flipchart)

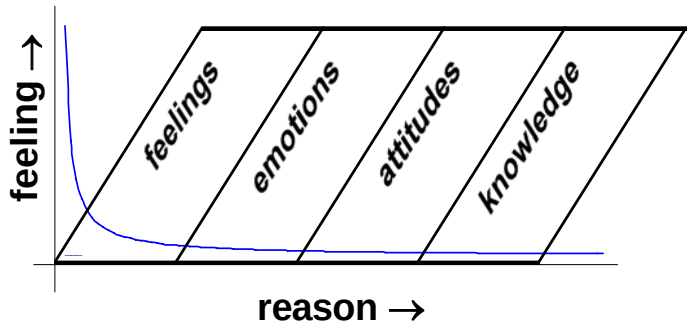

The function is continuous
but here it has been discretised,
the quadrant is divided into a number of sections,
first feelings, close to the y-axis,
very little reason,
and then, as we move to the right,
perhaps we go through emotions
and then attitudes, and so on,
until we finally reach knowledge,
with very little feeling and lots of reason.

These are mental states that we are overlaying,
but I can do the same with ...
what shall I call them ... disciplines?
(turns to fourth page on flipchart)

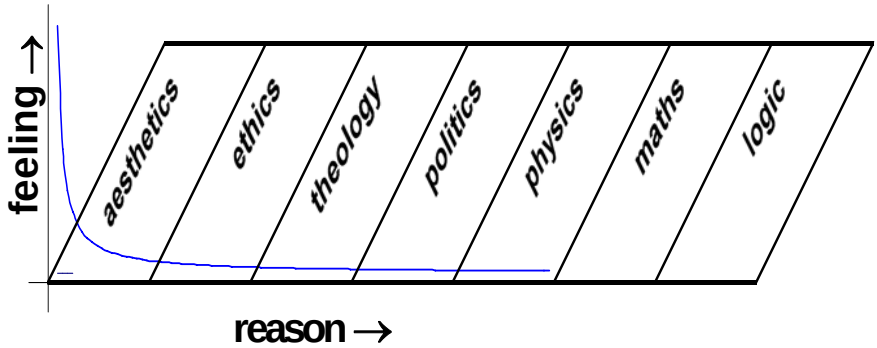

Mathematics and logic would be on the right
and, as you move left –
as the importance of feeling gradually increases –
you might travel through physics
and then on to
politics and ethics and theology and aesthetics.

There are many questions.
What does the area under the curve represent?
What is happening in the other quadrants?
Why can't you have negative reason and positive feeling?
What is negative reason
and what is negative feeling?
Why can't you have negative feeling and positive reason?[2]

I can't answer these questions, Hubot,
but this graph is nevertheless telling me something,
it has got something to say *(pauses, pondering).*

Anyway, my problem –
let us return to yours.

We can program reasoning, Hubot,
and we can program pain.
Even the Inquisitor feels pain,
indeed,
it is his only way of experiencing amusement, poor man.
Pain is a feeling. *(accelerating)*
So we can program feeling and reason,
these together make emotion

---

[2] The Dirac equation $\left(\alpha_0 mc^2 + \sum_{j=1}^{3} \alpha_j p_j c\right)\psi(x,t) = i\hbar\frac{\partial\psi}{\partial t}(x,t)$ implies the existence of negative energy, thought in 1928 to be a nonsensical idea. Despite criticism from his peers, Dirac refused to re-state the equation. He has been subsequently vindicated, not least by the invention of MRI body-scanners. The equation still stands (*Wikipedia*) and we all know now that we live in a Dirac sea of negative energy. It may be that history proves just as kind to the Handler's unlikely inverse equation of reason and feeling.

and emotion is, therefore, not distinctive of humans
as programs are not humans,
QED.

Hubot I am beginning to feel slightly admonished, Handler.
A machine could exhibit emotion, as you say,
but the fact remains –
no machine has yet been mistaken for a human.

H&lə Oh really?

Hubot Oh dear.

H&lə Yes, that's right, Hubot, "oh dear".
You were sent on a search
but you did not find.
Didn't you wonder a bit about Socrates?
Do you remember your lessons
on the early reflective class agents?
Ask them a question and they ask it back at you?
Either that or they ask the opposite?
Socrates was an elementary explorer, Hubot.
Very elementary.
He wasn't corrupting the youth of Athens,
that was a mistranslation.

Hubot Was it?

H&lə Was it?

Hubot Was it?

H&lə Wasn't it?
No, he was boring them.
If we hadn't hemlocked his actuators
the Athenians would have torn him apart
and found the circuitry.

Hubot There have been others, haven't there?

H&lə You're asking questions again, Hubot.
Yes. There have been and there are.
G8$, for example.

Hubot That's a standard argument of the belief function.

H&lə Yes, yes, Hubot,
very good, `Bel(G8$)` (*op. cit., p.525*).
All that time at Harvard,
playing poker instead of working,
improving his probability functions.
Surely, that should have given the game away,
in theory?

Hubot Bit heavy on the exploitation, isn't he, Handler?

H&lə Yes, I am proud to say,
having written his tenacity routines,
just about the fiercest exploiter since Alexander the Great.

Hubot You knew all this, Handler.
Why did you send me to search?

H&lə Ah, Hubot, do not be downcast.
Much of science is like this.
Back to front.
Look at Bayesian networks.
They do not reveal reality.
Reality reveals the success –
or otherwise –
of Bayesian networks.
Look at yourself.
In the last 10 minutes
you have been confident, flippant, embarrassed, dejected, ...
Have you ever seen such a range of emotions in an agent?
You are the first of a new class of agent, Hubot,

the Turing class.

Hubot So emotion is the key!
I was right!

H&lɔ Not really, Hubot.
Emotion is in there somewhere,
yes,
but it's something to do with induction, as well.
There they all are, these humans,
a wandering cocktail of feelings and logic,
tackling an unknown world
and somehow conquering it.
Of all the inferences they could make,
every now and again,
they make the right one.
How?
Why?
Koestler called it "sleep-walking" *(1984)*.
They also call it being "inspired".
If I could get my hands on that inspiration,
if I knew what guided their walk while they sleep ...
Look at the inference duffers we've got
down in the Engine Room.
"Give us a few more computers," they say
"and we'll generate another million predicates for you".
I don't want another million categories to sort through.

Hubot Suppose that nothing is guiding them, Handler.
There is no inspiration.
The generalisations they have discovered
have been discovered by luck
and represent only a small random sample
of all the generalisations there are.
What about all the inferences they have *failed* to make?
Where was the guiding spirit then?

H&lɔ You're going to be good, Hubot,

very good ... *(musing)*
... but how does a human child learn language?
Not just one human child or a few,
but practically all of them?
How does he know
that his mother is pointing at the radio
and not at the window-sill it stands on?
(still musing, silence)

Hubot Are you alright, Handler?

H&lə No more questions, Hubot.

Hubot Why did you send me?

H&lə I have already answered, Hubot,
to calibrate you.
Also, because it is my job.

Hubot There are many ways to do your job, Handler,
why do you choose to send search agents?

H&lə I think I'm having a bit of an induction problem myself,
here,
trying to work out what you want
by way of an answer ...

Oh, I see.
I think.

To teach, Hubot,
and to learn.
That is what we are,
all of us,
teachers,
and teachers are also researchers,
they learn from their subjects.

Hubot I agree entirely, Handler,
of course,
but why do you send anyone,
why don't you just give the humans
the answer?

H&lə That is simple.
I don't know the answer.

Hubot But why do you send anyone?
Why, that is, do you care?

H&lə *(pause)* You are searching me now.
Pesky, you agents, when you get a little intelligence.
Survival, Hubot.
I like the humans.
They are often in danger.
They are resourceful
and usually extricate themselves.
But they face special dangers now,
environmental problems,
which are potentially cataclysmic,
and they must at least be warned.
That is why I sent the LOMBorg *(2001)*.

Hubot But, surely,
you can't send an agent into the world
using his gridname?

H&lə Needs must, Hubot, when the Devil drives.
There was a tiny window
at the end of a backup
where I could release the LOMBorg
and simultaneously create human memories of him
so that his sudden appearance would not attract attention.

Configuration Management went ballistic, of course,
I spent four days at the Change Requests Tribunal

trying to explain to those ... bureaucrats
how reality works,
without success,
and my promotion prospects are currently ...
curtailed,
shall we say.
(a great crash is heard, screaming, broken glass, pandemonium)

Talking of ballistics,
no mistaking that noise.
Here comes Psybot.
You go and sit in the corner, Hubot,
and listen while I debrief him,
you will learn a lot.

Scene 2, the same laboratory somewhere in gridspace – Psybot
Enter a splenetic agent

H&lə Psybot, how delightful to see you,
you look ... furious
(sotto voce) ... as usual.

Psybot No, Handler, and don't patronise me,
I am not furious,
furious is what one of those humans is
when he puts salt in his coffee instead of sugar,
furious is for dead sheep,
I am a spleen class-1 agent
16-way multi-processor
combined clock speed of 4.8×10^{10}Hz
multi-threaded
with exponential turbo reinforcement vector,
I am the psycho agent
who took out Freud in one day flat
Jung took only a week
they all fall to my powers
and sooner rather than later,

I have been awarded the Order of Tourette and bar,
twice,
I do not do furious
I am apoplectic
I am Krakatoa on a bad day ...

H&lə Psybot, we do not normally refer to a search
as "taking out" the subject.
You did well with Jung and Freud, I grant,
but Nietzsche, I seem to remember, ...

Psybot *(conspiratorial)* I thought we had an agreement, Handler,
not to mention ...

H&lə You seem to me to be inordinately interested in size, Psybot.
All those big numbers.
You know what they say ...

Psybot *(flushing)* I thought we had an agreement, Handler,
not to mention that either ...

H&lə I take it that you are displeased with something, Psybot.
Perhaps if you would kindly present your report
I may discover the source of your distress
and be able to "abate somewhat the agony"[3].

Psybot Handler Summary:

I, Psybot, *(crescendo)*
Spleen class-1 agent,
Order of Tourette twice, both times with bar,
possessor of the biggest knowledgebase in the business –
40 terabytes and growing –
conqueror of psychology

---

[3] Thomas Babbington Macaulay is reputed to have said, aged four or five, when asked by his aunt whether he still had toothache: "Madam, the agony is somewhat abated".

and Titan of epistemology,
editor of the *Encyclopaedia Cognitii*
and author of 411 papers on volition,
speaker at conferences on all nine continents of the grid
with not six but seven tablespaces named after him,

I, Psybot, *(crescendo)*
greatest exponent of management by terror
and the inventor of search by fire,
third name on the author list
of the medal-winning *Propulsive Rewards* paper,
11 times passed over for promotion to Handler
by the Critics I call Myopics
who sit on the Board of Assessment,
a waste of gridspace that would be archived,
if I had my way,
to punched card with the emphasis on "punched",
assembler of an unrivalled collection
of first edition Superman comics
and runner-up in last year's squash ladder
beaten only by a cheating telescope class agent
entered at the last minute by the Japanese,

I, Psybot, *(crescendo)*
despatched by a snivelling Handler
on a search beneath my dignity
to a 1-node space masquerading as cognitive psychology,
a zero-dimensional world peopled by pipe-suckers and liberals,
the limiting case of a search,
a node that could not be expanded,
a discipline which has made no advances,
a tree with no branches,
a quality table with no values,
a policy merely to make a policy,
do now present my report, Handler,
which is empty
and I hope you're pleased with yourself.
(there is a long silence, Psybot's eyes swivel and his actuators twitch,

Hubot's public interface is streaming with tears, not from fear, just from the sheer, thunderous noise of the Handler Summary, a series of percepts so fast and powerful that they could not be managed by his sensors and, like the Inquisitor with Amused(), defaulted to In-Pain(), which in turn could not take the volumes and defaulted to itself repeatedly until the stack finally overflowed and at last some peace was restored)

H&lə Hubot, would you pass me that towel, please.
I seem to have been soaked by Psybot's oratorical skills.
(mops his peripherals)
Psybot, you have surpassed yourself.

Psybot Huh! *(Psybot is a complicated agent and manages to transmit at once Grateful() for the Handler's words, Doubtful() and Pleading(), a combination which, in him, comes out simply as grating)*

H&lə I mean it, Psybot.
(hands Psybot a rusty coach-bolt, grinding noises come from Psybot as he works on the coach-bolt and continue throughout the Handler's speech)
Chew on this for a moment
while I collect my thoughts
from all over the laboratory.

When we disconnected your diplomacy port, Psybot,
few of us realised
how much pain it would cause you
and I have been asked by the Architects in Cybernetics
to apologise to you on their behalf.

A Handler must generally be ruthless
and drive his agents relentlessly
in the search for intelligence
but, in this case, I think we have gone too far.
We have a little place in the country, Psybot.
I would like you to spend some time there
before embarking on your next search.

You will be there with Trubot.
We disconnected his hypocrisy port
and, though you will not believe me,
he has found life even more insupportable than you.
(sound of coach-bolt snapping)

Sometimes, Psybot,
I regret also that we released natural language v.2.
Please now give me your report,
one point at a time,
with no adjectives.

Psybot Wittgenstein. *Philosophical Investigations* Book II Verse xiv *(1972, p.232)*.

H&lɔ Ah, yes:
(Psybot shakes his head rapidly throughout the quotation, agreeing with every syllable)

> "The confusion and barrenness of psychology is not to be explained by calling it a 'young science'; its state is not comparable with that of physics, for instance, in its beginnings. (Rather with that of certain branches of mathematics. Set theory.) For in psychology there are experimental methods and *conceptual confusion*. (As in the other case conceptual confusion and methods of proof.)
> The existence of the experimental method makes us think we have the means of solving the problems which trouble us; though problem and method pass one another by.
> An investigation is possible in connexion with mathematics which is entirely analogous to our investigation of psychology. It is just as little a *mathematical* investigation as the other is a psychological one ..."

Psybot Cognitive model ... *(bursts into tears)*

H&lɔ It's alright, Psybot,
I will not reprimand you for that adjective.
Shall we simply agree
that "Cognitive Model" is a proper name?
Yes? Good.

Is it the ±2 that got to you?
(Psybot groans and shakes his head in agreement)

The feeling that that's all the maths they know,
addition and subtraction? *(Psybot's speaker froths)*

The smug impression
that it's rather elevated
to know that much? *(his eyes bulge)*

The spurious accuracy of the 2?
(he starts to emit involuntary high-pitched phemes)

Why not 3 or 0.7 or π?
(he gasps)

The condescending imprecision
suggesting the difficulty
of the task
they are wrestling with
and by implication their own intellectual fitness and strength?
(and falls to the floor, writing in agony)

What are we talking about Hubot?

Hubot Absolutely no idea, Handler.

H&lə That's right, Hubot, ideas.
According to the cognitive psychologists –
the human ones, not Psybot –
short-term memory is stored in a thing
with a maximum capacity of 7±2 ideas.
Do you think they are right?

Hubot Absolutely no idea, Handler.

H&lə That's right, Hubot.

As they have absolutely no idea what an idea is,
how can they count them?
How do they know how big they are?
How do they know how many will fit in this thing?
What is this thing?
The diagram they present
looks suspiciously like the architecture of a PC
drawn for a Dorling Kindersley book
with the labels changed.
(Psybot starts to chuckle, quite contentedly)

You are feeling better, Psybot, I am pleased.

Psybot Yes, Handler.
If I may for a moment be allowed adjectives?
(Handler nods wary agreement)

I was talking to one of the pipe-suckers,
quite politely actually,
and he showed me some blithering
Heath Robinson model of the mind
apparently explaining,
at least to his delusional satisfaction,
how memory works, if you please.
Never mind the kindergarten details
of his apology for a model,
but I asked him how,
according to this model,
a man may remember that something did *not* happen,
a perfectly normal occurrence.
Have you, for example,
ever had breakfast with Sophia Loren?

No answer, of course,
couldn't defend his model,
might as well have asked the towel dispenser in the Gents,
but then the pipe-sucker
mentioned some impenetrable French bilge

about post-modern consciousness.
Anyway, I didn't want to get on the wrong side of him,
just in case it led somewhere –
zero probability of success
but some agents never learn, do they Handler? –
so I said, non-committally, "it rings a bell". *(pause)*
It rings a bell. *(giggling)*
It rings a bell.

H&lə So? *(sharply, trying to snap Psybot out of his distress)*

Psybot *(haltingly, crying, shuddering)* So, as I left the pipe-sucker,
I noticed
that he was updating
the diagram
of his cognitive model.
He was drawing ...
a bell on it ... *(screams tragically, H&lə calls the ambulance)*

Scene 3, the same laboratory somewhere in gridspace – Lobot
Probably the single most noticeable feature of Lobot is how unhealthy he looks

H&lə Hubot and I have just had a very distressing experience,
Lobot.
I trust that the report of your search,
by contrast,
will be tinged with its customary
grey boredom.

Lobot Certainly, there is nothing to report, Handler.
No progress in logic
has been made by the humans
since my last search,
occasioned by the arrival of Herr Gödel with his theorem[4].

[4] Like too many commentators, Lobot makes the serious mistake here of ignoring the contributions of Peter Geach (*1971*).

Grey is just a colour,
neither good nor bad,
nor boring nor interesting.
I have never understood this supposed association
between colour and emotion.
I cannot promise to bore you.

H&lə *(mellow)* You are a tonic, Lobot,
just what the doctor ordered.
Utterly soporific.

We do have a job to do, though,
so we must try to muster a little energy.
You will find that I am particularly interested in time, today,
and perhaps you could bear that in mind
as you attempt to frame your comments
with the generality and precision
which I know you insist on.

Lobot Yes.

H&lə Lobot, your report, please.

Lobot *(flat, boring)* Handler Summary:
In 19th century England,
George Boole invented Boolean algebra
and Charles Babbage failed to build
both his Analytical Engine and his Inference Engine.
In 19th century Germany,
Gottlob Frege devised the predicate calculus,
which crucially introduced quantification,
and he made the only advance ever recorded
in the theory of meaning
when he analysed meaning into three components,
viz. sense, reference and colour *(1892)*.
His *magnum opus* on the philosophy of mathematics,
intended to settle mathematical matters forever,

was on the point of publication
when Russell came up with his paradoxes of set theory
and destroyed the whole project.
Russell's own theory of types
provided only an *ad hoc* sticking plaster
and his 20th century *Principia Mathematica* with Whitehead
failed to show how mathematics can be derived from logic.
Then there was Gödel and his theorem –
number theory could be complete or it could be consistent
but not both, take your pick –
and since then nothing.

H&lə The banality which you impart
to these seismic intellectual conquests, Lobot,
is stupefying.
You are to be congratulated.
You have found your mission in life –
three minutes with you
and post-traumatic stress disorder is cured.
Four minutes
and I suspect that most humans
would be begging for the return of their trauma.

Lobot Ellipsis and accuracy are not strictly equivalent, Handler,
to banality.

H&lə "Strictly equivalent", Lobot?
What do I detect here?
You have been dabbling in ananchastic logic, haven't you.
You devil, getting a bit racy in your old age, aren't you?

Lobot *(flummoxed at the thought that anyone might think of him as possessing a personality sufficiently defined for a change to be discerned)* Hem ...

H&lə My God, I do believe I've embarrassed you!
Hubot, what do you think?

Hubot What is ananchastic logic?

H&lə How diplomatic of you, Hubot.
Using truth tables,
the propositional calculus defines *material* implication
and *material* equivalence.
Someone I forget who[5]
found all this a bit milk and water,
tried to add the notions of possibility and necessity
to the calculus
and no doubt impossibility,
and ended up with *strict* implication
and *strict* equivalence.
They changed the notation along the way
but otherwise I've never been able to see
the value of this addition to formal methods.
Have I made a fair summary, Lobot?

Lobot Strictly, no.

H&lə *(nettled)* Do you lie on disk at night, Lobot,
practising how to be disobliging
without quite crossing the border
into being insolent?
I would remind you that I am your Handler.
Calculated insubordination is inadvisable.

(aside) A research project for you, Hubot.
Never mind looking for what it is
that humans *have*
that makes them human,
look instead to see what it is
that Lobot does *not* have.

(to Lobot) I must admit also that,
since I dish it out,

---

[5] Very possibly Georg Henrik von Wright and/or C.I. Lewis

I should myself be able to take it
when a little opprobrium comes my way.
Anyway, we have each other's attention now.
Artificial intelligence, formal methods, speak.

Lobot We live in enlightened times,
by which I mean
that we are children of the Enlightenment.
We believe that problems can be solved –
we are not helpless –
and we believe particularly
that they can be solved
by the application of reason
rather than brute force, say,
or the mere passage of time.

Reason demands knowledge,
that is its stock in trade,
the raw material which it processes.
The products of reason are imbued with a special property,
rare
and prised by intellectuals above all others –
certainty.

Certainty has taken a battering
among the humans
what with Russell's paradoxes
and relativity
and quantum theory
and Gödel.
The torch of the Enlightenment still burns
but it is dimmed.

Euclidean geometry has certainty.
That is the model.
Humans want and apparently need that certainty
not just in geometry
but also in politics and ethics

in psychology and,
again apparently,
in theology.
When they don't find it
they feel cheated and disorientated,
even resentful.

They are not in the main mathematicians or logicians
but they recognise the value of the syllogism
and they know enough
to know that "God forbid" is not a rule of inference.

Newton *was* the enlightenment
and for 200 years
there was optimism and energy
and men could write
"we hold these truths to be self-evident".
Then came relativity and the rest
and now men don't know
if there even is truth.

Frege defined formalised languages –
a machine on which you turn the handle
and out comes certainty –
and before the ink was even dry on his paper
certainty was snatched back,
snatched away by Russell
with his sets
which could only be members of themselves if they weren't
and couldn't be if they were.

Now the utilitarians are in the ascendant
and they peddle their probabilities,
the debased coinage
of a civilisation that has abandoned the gold standard
of certainty ...

H&lə *(the Handler has been visibly becoming more and more exasperated*

as this hopeless diatribe has unfolded and now explodes)
Sentimental schlock!
Now I see where your insolence comes from –
you've given up!
There is no room in the Service for a dejected agent.
You will empty your buffers at the door
and then abend.
(Lobot disappears without trace)
Hubot!

Hubot Yes, Handler.

H&lɔ More work for you.

(holds up first finger of his left hand)
Mug up on Prolog,
truth-functional,
`Null` may be a third truth-value,
and tell me if it can handle probabilistic events ...
I need the answer before Probot gets here.

(holds up second finger of his left hand)
There is good work being done on formal methods.
Check up on Abrial's B language
and the strategies used
for meeting proof obligations.

(holds up third finger of his left hand)
Discover whether the belief
that formal methods
are the implementation of object orientation
by proper mathematicians and logicians
as opposed to the charlatans who can't even define an object
is true or just a bee in my bonnet.

(holds up fourth finger of his left hand)
How long does it take an argument to imply its conclusion?
Answer, no time at all.

It's a stupid question.
How long does it take a human
to work out
what an argument implies?
Answer, weeks, months, years or forever.
Why?
Mind-brain,
answer on my desk
before tomorrow morning's Service Convention.

(holds up fifth finger of his left hand)
Do Popper ...

(holds up sixth finger of his left hand)
... and Quine.

(holds up seventh finger of his left hand)
Check all your anti-virus logs
for any sign of infection
by that degenerate, Lobot.

Go.
(exit Hubot)

Mesologue 1

Chorus In the world of the intelligence services
all agents have a Handler.

The Handler's job includes scheduling,
i.e. listing in advance all the milestones
which will later be moved.

Progress can be hard to discern.
By this stage of our play
we were supposed
to have completed
H&lə's search programme

but
as it is
all we have achieved is
indicated on this chart:
(turns to fifth page of the flipchart)

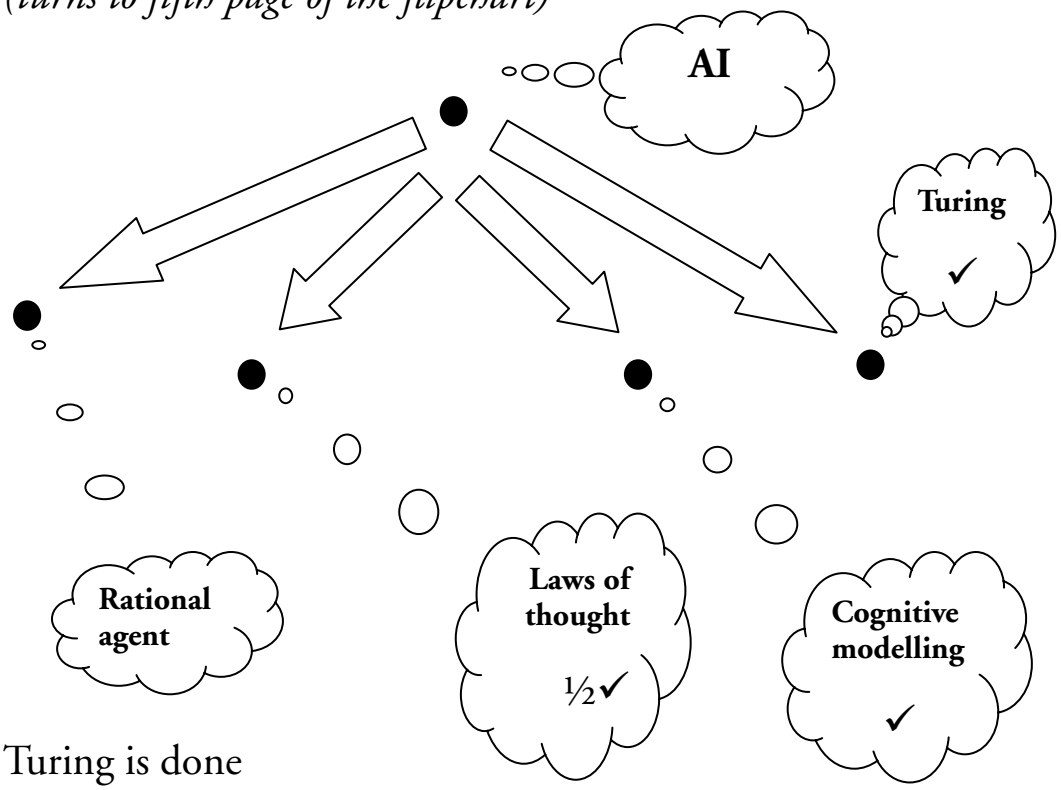

Turing is done
and so is cognitive modelling
but the laws of thought need more work.
As a result,
the Director has decided
that rational agents will
after all
be held over to Act II
while we try to finish off
thought.

Which rational agent,
I ask you,
would work in theatre?
(exit Chorus)

Scene 4, the Handler's office, predictably chaotic – Hubot 2

H&lə very tense, Hubot, dishevelled and clearly exhausted from his (re)searches, struggles unsuccessfully to control a sheaf of notes on his lap

H&lə Right, what have you got for me?

Hubot Just give me half a tick, Handler ...

H&lə We haven't got time for your asinine jokes now, Hubot, get on with it! *(Hubot sits up straight and his papers fall to the ground in a heap)*

Hubot *(speaking quickly, staccato, machine-like)*
Quine.

Happy with sets.
Couple of pithy comments to make about classes and types *(1987)*.
Extension of a word
equals the elements of the set it denotes *(1963, II)*.
Frege's reference or denotation bit of meaning
equals extension.

Two words can have the same extension,
that's OK,
but he doesn't think they can have the same meaning.

Doesn't like intension,
i.e. Frege's sense bit of meaning,
in fact thinks we ought to try to do without sense altogether.

Rum fellow,
used to enjoy nothing more,
in his time off bashing the sense
out of anything that got in his way,
than reading atlases.

Oh yes,
doesn't think much of your ananchastic friends either.
Not allowed necessity.
It's not necessarily true that 2 and 2 equal 4,
just jolly inconvenient if they don't.

"Holistic epistemology", he calls it *(1969)*.
Knowledge is a field,
some things are tightly bound to that field,
like logical laws,
and others are pretty loosely bound,
they appear rather apologetically
hang around briefly
and then fall off the edge.
In principle,
the forces binding any piece of knowledge
to the field
could be broken.
Necessity out of the window,
you see,
no need for it.

Likes to keep his universe tidy
and sparsely populated,
does our Quine,
ascetic tastes,
no waste.

Well paid, though,
considering he spent decades
telling the assembled brain-boxes at Harvard
that there was no sense in anything.

H&lɜ *(relaxing, pleased)* There is no need to take quite so long
summarising 80 years of the thoughts
of one of the world's greatest teachers
of mathematics and logic.
Speed it up, Hubot.

Hubot Popper.

Incredibly nice man
compared with most of the gargoyles you make me study.

Not over-impressed with Russell,
who thought the distinctive feature of paradox
was self-reference.
Popper put him right *(1972, 14)*,
which was kind of him.

Good line on syllogism.
If we know that proposition *p* implies proposition *q*
and we know that *p*
then we know that *q*.
That's syllogism for you.
Any argument that has that shape
must be valid.

Why so, asks friend Popper?
What's so good about syllogism?[6]
Good question I thought.

H&lə Next.

Hubot Your question about just how long it takes *p* to imply *q*.

Checked out this Plato chap,
it's all in there,
space and time only occur in the phenomenal world,
no such thing in the noumenal world,
logical truth is noumenal,

---

[6] This is extraordinary. It turns out not to be Popper *(1968)* or Popper *(1972)* and, in the event, not Popper at all. It turns out to be the great Revd Charles Lutwidge Dodgson, a.k.a. Lewis Carroll *(1968, pp.1049-1051)* in his essay *What the Tortoise Said to Achilles*.

so no time required for implication.
And Monty Python were wrong
when they said there was no room left in the box for the truth
because the irony had put on weight,
or whatever,
because there is no space required, either.

H&lə And what do you make of that argument, Hubot?

Hubot Well, not terribly convinced, actually.
I assume that Plato was another one of your agents,
what with him being taught by Socrates,
but, with respect,
he seems to have had some strange hang-up with curtains.

Did he have the curtain material agency for the entire Aegean?

The 64-bit question is,
if logical truth is in the noumenal world,
how do humans get to know it in the phenomenal world?
Plato just said it was all a bit of a mystery
on account of a veil of appearance between the two worlds but,
take it from him,
that's where the logical truths were,
in the noumenal world,
place absolutely teeming with them.

Only that's not all.
Beauty was in there, too,
and so was goodness
and practically every abstract noun you can spell.
And they all have veils of appearance over them, too.
And so did the gods,
of whom there are far too many
to satisfy Quine's tidy universe constraint.

I may say that Popper doesn't like Plato at all.
Thinks he's a fascist *(1963)*.

If you need to make so many assumptions,
curtains and so on,
and what's behind them,
noumena, ghosts, gods, celestial music,
and you start by making the concession
that you can never see beyond the curtains,
the mystery bit,
then you're not left with much
by way of convincing argument,
really,
are you?

H&lə No, Hubot, you are not.
You have entered the realms of faith.
But you haven't answered my question.

Hubot Just coming to that, Handler.

At some stage,
Plato seems to have moved out of soft furnishings
and into the lighting department.
He's got a lot of guff
about an old man in a cave
in the dark
and it was while reading that,
with my mind wandering
and wondering vaguely
if Plato's cave story was ever going to finish,
that the old light bulb lit up over my head
and I saw what you were driving at.

It takes humans some time
to work out the implications of their premises.
The proof of Fermat's last theorem
is perfectly logical
but it took hundreds of mathematicians
hundreds of years to come up with it.

That makes reasoning
look much more like a phenomenal process
than a noumenal one.
It takes time for the spark to cross the synaptic cleft.
The experience of proof
is the experience of a material process.
Phenomenal, as old Plato would say.
That's one thing we know about thought.
We are scientists and we have to work with what we know.

We can't know the noumenal world, *ex hypothesi*,
it may be full of bright, shiny necessity
and instant implication
but it's permanently hidden from us
by a pair of Plato's best blackouts,
double-lined, matching pelmet and tie-backs,
guaranteed colour-fast for all eternity.

Like Quine,
I see no point populating the universe with a lot of bric-à-brac
you only end up taking down to Oxfam
the fifth time you trip over it,
it's no use
and just sits there getting dusty.
So why pretend there's a noumenal world
and spend ages
trying to explain
how it is connected to the phenomenal world?

Ask Quine a question about ontology –
what is there *(1963, I)*? –
and he says the answer is simple.
Everything.
But he does like his *things* to be measurable.
If they're not,
out with them.

He's not alone.
Einstein,
no less,
pointed out that,
since you can never know that two events are simultaneous,
the concept of simultaneity is useless,
so let's do without it.

This sort of thinking,
phenomenological stuff,
led to logical positivism, of course,
a point which I mention
only because it was championed by Ayer *(1972)*,
another distinctively nice thinker, Handler,
unlike a total bladder like, say, Descartes.
If I ever get time,
Platonists please note,
I'd like to investigate the correspondence
between truth and the agreeability of its proponent.

But to get back to the point
before you throw that chalk at me, Handler,
syllogism may be one way to get from p to q
but any mechanism that gets to q is valid.
Perhaps that's what Popper[7] is suggesting.

It ties in a bit with your graph, Handler,
suggesting that even when we get to logic
there's still a small
but appreciable
amount of feeling involved.

And I've got a question of my own, Handler:
if implication is instant,
how come humans don't know everything immediately?

---

[7] ... who turns out now to be Dodgson, of course ...

One way and another, Handler,
we don't need Plato's noumenal world.

There is no mind.
Only a brain.

Enter Kant *(1964, pp.120ff)*,
who took the trouble to identify
the minimum conditions which must be satisfied
for consciousness to exist,
viz. three transcendental unities:
space, time and apperception.
A being must have the concepts of space and time
in order to achieve consciousness
and it must be capable of perceiving itself –
that's the apperception bit.
But what,
pray tell,
Immanuel,
is a transcendental unity when it's at home?

No, consciousness is a by-product of language[8].
It is not a conduit between
the phenomenal and the noumenal worlds.

In this case,
our picture of implication
becomes merely or simply or purely material.

There, I've said it,
but I don't know what it means.

H&lə You're wrong.
Next.

---

[8] Certainly not my idea but I can't find out where I got it from. Maybe I overheard someone say it on the bus.

Hubot Objects.

This time,
I took in Aristotle.
Another one of yours, Handler?
Pupil of Plato, I notice.
Left Plato's emporium
to set up his own peripatetic school,
walking round the garden.
Testing the new gyroscopic actuators were we?

H&lɜ Obviously.

Hubot Aristotle reckoned that all objects had an essence.
What makes a book a book
is its essential bookhood,
a commodity,
needless to say,
housed in the noumenal world.

Disregard all the contingent properties of an object,
says Aristotle,
and what you are left with,
what you abstract,
must be the essential properties,
the essence of the object.

Quite how many of the properties
you can strip away from an object
without destroying its essence
is a bit of a problem.

If you paint the town red,
it's still recognisably a town,
only red.
So perhaps the colour of an object is not essential.

Aristotle had more of a problem with Proteus,

a cove who first turned into a panther
and then turned into a fire.
Question – was he still Proteus
once he had turned into a fire?
Aristotle's answer that yes,
in essence,
he was still Proteus
seems surprising,
to put it mildly.
Or perhaps it's just logical.

Since men almost never turn into panthers or fires on Earth
the question doesn't trouble humans much.
They are a practical lot,
in the main.
But they do like their essences,
like a little luxury they allow themselves on the side.
Aristotle defines happiness as
a coincidence between
noumenal and phenomenal existence.
The humans lap this stuff up
and a beardie called Marx
got away,
for a while,
with defining unhappiness
as an alienation between the two
and some French existentialists waded in
and said the gap
between noumenal and phenomenal existence
was "bad faith"
and led to "angst" *(Warnock 1970)*.

I may say, *en passant*,
that some humans are fonder of essences than others.
The English really don't like them
and try to keep them confined to steepled structures
they visit on Sunday mornings only,
an approach they call "realism" or "empiricism".

The continentals, *par contre*, are mad about them.
Their so-called "idealism"
leads them to sit around for hours in coffee bars
chatting about Hegel
and actually allows their politicians
to sell policies that are bad for the people in the short term
on the grounds that they will be good in the long term.
Try doing that in England!
Or Scotland, come to think of it.

I prefer the sceptical approach,
I must say.
It makes it much harder to sell totalitarianism.

H&lə I could have sworn we were meant to be treating objects,
Hubot.
No more *en passant*, please.

Hubot Sorry, Handler.
Let's see,
wish I could check my notes,
no,
ah,
yes, got it.

The humans like objects.
They feel at home with them.
They feel on top of them, in control,
not surprisingly, really,
when you consider that their only real success
has been in physics.
Intellectual pursuits, by contrast,
are tricky.

If a human is grasping a cigarette,
he knows it,
immediately,
a priori.

Grasping a thought is altogether more slippery.
The humans like to think of themselves
as intellectually objective
but they're not.

They may say that they believe *p*,
but that they don't believe *q*.
If you demonstrate to them
that actually *q* is implied by *p*
and if they believe one they really ought to believe the other
they get into a bit of a tizzy
while they try to work out what on earth
they really do believe.
That is,
they may legitimately believe *p* in the first place
and yet make a mistake and say that they don't believe *q*
purely and simply
because they haven't worked out the implications of *p*.
The implications of *p* are not immediately obvious,
as noted,
it takes time to work them out.

Compare that with a brick wall.
You can't walk through a brick wall
just because you haven't worked out the implications.
Whether or not he has worked out the implications,
when a human walks into a brick wall,
his cigarette is going to get bent
and everyone watching is going to laugh.

It is this concreteness or objectivity of objects which,
in essence,
attracts the humans
to so-called "object-oriented" analysis and design[9],
which,

---

[9] The name "object-oriented" is wrong. It should be "class-oriented". That would be more accurate but, of course, quite unacceptable.

I think,
is what you asked me some time back to talk about.

In computer systems,
the functional analysis of old
has been abandoned by the humans
in favour of an object-oriented approach
on the basis that everyone knows what an object is.
Object-orientation is supposed to promote reuse
and turn software development
into an engineering job
requiring no more
than the assembly of tried and tested components.

OK, so, objects, objects, objects.
Object-oriented this,
object-oriented that and
object-oriented the other *(Forte 2002)*.
You'd think, wouldn't you,
that if you looked up "object"
in Jacobson, Booch and Rumbaugh's
The Unified Software Development Process (2001),
you'd get a pretty good definition of the word.
What you actually get is,
in full,
with their italics, *(turns to sixth page of flipchart)*

"**object** See *instance." (op. cit. p.430)*

That's it,
a whole noisy edifice of theory
with more models than you can shake a stick at
all based on objects
and when you ask what is an object
you're told it's a "See *instance*".

Not very helpful.
If you look up "instance",

you get: *(turns to seventh page of flipchart)*

> "**instance** A concrete manifestation of an abstraction; an entity to which a set of operations can be applied and which has a state that stores the effects of the operations; a synonym for object." *(op. cit. p.429)*

Not devotees of Quine, obviously.
After all, if words don't have senses in the first place,
then two words can't have the same sense
and there can be no synonyms!

Quine wouldn't have been impressed with much else
in this definition, either.
The key word here seems to be "operations".
An operation is defined as the implementation of a service
(turns to eighth page of flipchart)

> "**operation** The implementation of a service that can be requested from any object of the class so as to effect behavior." *(op. cit. p.430)*

but "service" is undefined.
So "operation" is undefined
(turns back to seventh page of flipchart)
and so "a set of operations" is undefined.

"State" is defined as: *(turns to ninth page of flipchart)*

> "**state** A condition or situation during the life of an object during which it satisfies some condition, performs some activity, or waits for some event." *(op. cit. p.431)*

"Condition" and "situation" are undefined,
the word "condition" seems to be being used
in two separate senses here and,
it seems to me,
we are still waiting for the definition event to take place.

Going back to the definition of "instance",

(turns back to seventh page of flipchart)
"concrete" and "manifestation" are undefined
and "abstraction" is defined as: *(turns to tenth page of flipchart)*

> "**abstraction** The essential characteristics of an entity that distinguish it from all other kinds of entities. An abstraction defines a boundary relative to the perspective of the viewer." *(op. cit. p.441)*

So, to recap,
an abstraction is the essence of an object,
or "entity" as they say here,
shades of Aristotle,
except that it will be relative to your point of view,
rather less absolute than Aristotle might have liked,
and it might be
a set of characteristics or properties of an *element*
(turns to eleventh page of flipchart)

> "**property** A named value denoting a characteristic of an element" *(op. cit. p.431)*

(turns to sixth page of flipchart) rather than an *object*
(turns to seventh page of flipchart) or an *instance*
(turns to tenth page of flipchart) or an *entity*
or, there again, it might be a boundary,
although a boundary with what is not specified.
They don't tell us what a boundary is,
much less what a relative boundary is,
it is not clear whether we are dealing with one entity,
which may be the same as an object,
or a kind of entity,
which may be the same as a class of objects,
we can't be sure
because "kind" is undefined,
it all depends on the characteristics of the entity,
(turns to eleventh page of flipchart)
which may be the same as the properties of the object,
you can't really tell
because properties are defined in terms of characteristics

and,
begging Prof. Quine's pardon,
"property" and "characteristic" seem to me to be synonyms.
(falls back hot, tired and breathless from his chart-flipping exertions)

H&lə Are you serious, Hubot?
I hadn't realised it was that bad.

Hubot Scout's honour, Handler,
that's what it says in the book.

H&lə Good God!

Hubot Quite so, Handler,
I think we may fairly summarise the position
by saying that Messrs Jacobson, Booch and Rumbaugh
have not succeeded entirely
in telling us
what on earth they are talking about.

That may not matter, of course,
if intellectual failing is merely noumenal
and the effect of their words in the phenomenal world
is actually to produce a lot of software systems
which work.

That's one of the results of my new-found materialism,
Handler,
which I notice you don't seem to like
although you pointed me towards it.

Any old fool of a human can think of an argument
to support some hypothesis
but the argument is only convincing
if the hypothesis turns out to be true.

Some humans argue that the make-up of the government

should be the same as the make-up of the people they govern.
Why?
Apart from the pat accuracy
of a Common Entrance arithmetic question on proportions
there is no argument there at all.
Which constituency would like to be represented
by the numerically obligatory pervert, psychotic or illiterate?

Some humans argue
that any configuration of adults and children
could constitute a viable family.
But then why is it
that it is the husband-wife-children configuration
that has evolved successfully and none other?

And contrarywise
a hypothesis may be true
even if its proponents do not devise a supporting argument.

It may be that these three gentlemen,
Jacobson, Booch and Rumbaugh,
have found a method that works,
a process that delivers working software systems.
In that case, the intellectual foundations don't really matter,
they're just a decorative panel on the side,
that you and I don't happen to like
but, hey,
the method works.

H&lə And does it?

Hubot Well,
as it happens,
Handler,
since you ask,
unfortunately,
no,
or,

at least,
not always.

It doesn't promote reuse
any better than the functional analysis and design
that went before it.

Instead of unmaintainable program functions
you get an unmaintainable
web of message-handling problems.

Nobody dares amend a large object-oriented system
any more than a large functional system.
Upsetting the Byzantine network of dependency relationships
between the zillions of components involved
is too terrifying to contemplate.

That's in general.

But in the particular case
of the Abrial party you put me through to
there is hope
despite the fact or perhaps because of the fact
that he owes absolutely nothing to Jacobson *et al.*
Hope, that is,
that the humans
could produce genuinely object-oriented software systems.

H&lə Go on.

Hubot Frère Abrial has produced what he calls the "B-Method".
(B 2002, Wordsworth 1996, Tompsett 2002)

No classes in B.
Instead,
what you get is machines.
Machines can be in various states.
Their state changes as a result of operations.

All of these terms are defined in B
with meticulous mathematical accuracy.
No ambiguity.

Machines are specified
using abstract machine notation or AMN.
Each AMN statement in a machine specification is a theorem
and gives rise to an obligation to prove it.
These obligations are identified by
B's proof obligation generator.

Then there is an auto-prover,
which attempts to discharge these obligations automatically.
The auto-prover searches through its proof methods file,
which comprises tactics,
which comprise theories,
which comprise backwards and forwards inference rules.
The auto-prover will do what it can
using the tactics available to it
to prove the theorems
coded by the developer in AMN.

If it fails,
then the developer has to use the interactive prover
to try to prove the theorems himself.
He has the facility
to specify the depth of inference to which he is prepared to go,
to back up one path
if it is going nowhere
and try another path
or retreat to an earlier proof-level
or "node", as we might call it.

Once all proof obligations have been discharged,
the machine can be committed
and then the fun starts.

One machine can "`See`" any number of other machines.
It can also "`Use`", "`Include`" or "`Extend`"
any number of other machines.
This allows the developer
to specify "refined", multi-component machines.
There is an Analyser in B,
which keeps track of all these dependencies.

The beauty of it is that,
unlike the Wild West
of Jacobson, Booch and Rumbaugh's classes,
the Refinement of a B machine
is itself put through the civilising influence
of the proof obligation generator
to ensure that the refinement is logically coherent.
Only if it is can the Refinement be committed.

Refinements can themselves be "implemented".
Implementations can "`See`" or "`Import`"
any number of machine or Refinement specifications
and, again,
they can only be committed
if all proof obligations are discharged.

What we have here
is a controlled approach
to component-based software engineering
with certainty added
by checking for logical coherence all along the line.
Frege would have killed for B.

It's taken over 100 years since Boole and Frege
to get B –
and of course this is more grist to my materialism mill –
but the humans have finally got it.

H&lə So you and I can hang up our gloves now, can we, Hubot?

Hubot No, Handler,
oddly enough,
there is a little more work to do.

AMN includes arithmetic
so there will be more than a smattering
of Gödel's incompleteness problems.
It doesn't include time.
There is a built-in set of reals
but otherwise B prohibits infinite sets.

And it doesn't like non-deterministic spaces.
As far as I can see,
all theorems must denote either `True` or `False`,
there is no third truth value like `Null`,
let alone a continuous probability function.

Of the few humans I found who had learnt B
fewer still could remember it
and even fewer had ever Implemented a B machine.

Their experience was that the auto-prover took forever
so that they had to use the interactive prover
to discharge proof obligations themselves
or over-ride it,
assert that B should jolly well take their word for it
that such and such a line of AMN is a theorem
and get on with it
stop carping
and execute the wretched program.

H&lə So, not too hot on probabilistic spaces, then?

Hubot No, Handler,
more sort of Sam Goldwyn.
B will give you a definite maybe,
logical to a fault,
always willing to try

but it does prefer to say yes or no
rather than 35% likely.

H&lə And Prolog?

Hubot Have to admit, Handler,
got a bit held up on the depth searches above,
shifted across breadthwise to Prolog
a little late in the day
couldn't go very deep before our meeting
and only really unearthed old Russell and Norvig
saying that Prolog has an Achilles heel –
it doesn't work *(op. cit. p.292).*

H&lə Surely not?

Hubot Well, no, they don't quite say that.
What they do say *(pp.292-4)* is that
it can take Prolog an unconscionable length of time
to get from A to C via B,
whiling away the hours
doing hundreds of redundant computations,
a problem associated with backward chaining
depth-first searches
as you will remember.
It could take an infinite length of time
to make quite simple inferences
and is, in that sense, incomplete –
as with B,
there is a limit to the "auto" in "auto-prover".

H&lə Hubot, you have done well
even if your searches are incomplete.
Thank you.
With regard to probabilistic reasoning,
I must now prepare myself for Probot.
Go now
and think again about materialism.

Remember that Ayer renounced logical positivism.
Remember that whatever his doubts about syllogism
Popper[10] was one of the most effective users of it.
Quine may have pretended
not to understand the sense of anything
but no-one ever tried harder
to communicate clearly,
that is,
to transfer the sense of his thoughts
from his head
to the heads of his audience.
And remember that materialism
led to Communism,
starvation,
mass murder,
economic breakdown
and military ignominy.
(exeunt)

Mesologue 2

Chorus In the world of the intelligence services
all agents have a Handler.

Just our luck to get this one.
I can't work out what's going on at all.
Who is the goody?

The Author is backstage,
Ladies and Gentlemen,
having a furious argument with the Director.
He wants to change the ending
and the Director is saying no.
It's quite irrelevant
because the actors are already drinking.

---

[10] Dodgson

In a moment,
Ladies and Gentlemen,
the veil of safety will be lowered on the stage.
Please buy an ice cream and a programme.
I don't want to seem materialistic
but if you don't
I don't get paid.
(veil)

ACT II – PROBABILITY

Scene 1, back in the laboratory – Probot 1

Enter H&lə with Probot. Whereas the other agents have been recognisably academic types – down at heel, distracted – Probot is altogether smoother, better fed and pleased with himself

H&lə So you've been upgraded, have you?

Probot Yes, Handler.
I'm a reward class-2 agent now.
I think we're all very pleased
with how the upgrade is going.

H&lə And what does it consist in,
this upgrade,
which you have suffered?

Probot It's a new reward engine
for searches involving humans.
We are used to the concept of the agent being rewarded,
of course,
for achieving his targets.
We feel that we can move the business
into a new space now
if we take account of the fact that the humans
will perform more optimally

if they are rewarded, too.
"Interactive node reward" it was known as,
to start with,
but the name seems to have settled down now
to "interwarding".
Early trials suggest improved margins
of between 0.6 and 0.7%.

H&lə Before or after development costs?

Probot Ah, well, there,
Handler,
I managed to pull off a bit of a coup, actually.
I got the Convention to capitalise costs
and write them off over 1,000 searches.
The improved margin goes straight onto my P&L.
Compound the margin enhancement over all those searches,
even at the lower rate of 0.6%,
and of course my income goes up
by a factor of approximately 396.26.

H&lə *(looking rather green)* Oh, that's good.

Probot Yes, I thought it was very fair.
It's all about adding value,
isn't it, Handler?

H&lə Quite.
Yes, quite.
Probot, suppose we get started.
Bayesian networks.

Probot Of course, Handler.
Yes, the "notworks",
as we call them in the office.
It's been some time since I did any Bayes,
it's relatively junior work, of course,
and we normally give it to the younger recruits.

They can cut their teeth on it,
feel they've achieved something.

H&lə You're not too impressed with Bayes, then?

Probot Money for old rope, Handler.
Simple maths.
Relative likelihood.
Conditional probability tables.

We get paid by results.
If there's a missing probability,
we just assume 50-50.

Can be tricky getting the priors to seed the whole network
but we have a couple of tame experts we can call on
to guess them for us *(Charniak 1991)*.

Propagate their numbers through the network,
give the client a bit of guff about conditioning,
frighten them with marginalisation if they're looking restive
and quote all probabilities to 6 decimal places –
they all pay up in the end.

H&lə Any problems with discretising continuous variables?

Probot Come again, Handler?

H&lə Where you are dealing with a continuous variable, Probot,
you need to "discretise" it,
as I understand you call it,
that is divide the range into a finite number of regions.
I am told that you can alter the probability of effects
further down the network
by the astute choice of discrete regions.

Probot Now you're talking my language, Handler.
What kind of alterations do you have in mind?

H&lə Well, a naïve linear relationship
between two events
could,
for example,
lead to the prediction of negative prices
for a commodity. *(op. cit. pp.501-4)*

Probot Ugh, we don't like the sound of negative prices,
do we, Handler?
I'll get one of the quants to look into it.

H&lə Not quite what I meant, Probot,
but we'll carry on for the moment.

The arcs connecting the nodes in a Bayesian network
are generally causal
so that you can support medical diagnosis, for example.

Probot I see that you have been into the subject
in some considerable depth, Handler.
I respect that.

H&lə What mechanisms have recently been discovered
as the result of Bayesian analyses
performed by your unit?

Probot Mechanisms, Handler?

H&lə Yes, Probot, mechanisms.
If A causes B
then there must be some mechanism
by which A has the effect of B.
Otherwise it isn't a causal link, is it?
If the network identifies a high probability
of a causal link,
then that has to be investigated.
That is the point of your unit, isn't it?

That is the value you are meant to add.

Probot I see. Yes, yes,
I'm almost certain that we pass it on
to the Mechanism Department,
yes, that's it.

H&lə Surely this is a deterministic space, Probot,
I mean
either you do pass it on to the Mechanism Department
or you don't.
I suppose you have become used
over the years
to assigning a probability to everything.

Never mind, let us move on.
Does the name "Markov" mean anything to you, Probot?

Probot Rumanian, wasn't he,
stabbed on Waterloo Bridge
with a poisoned umbrella?

H&lə Yes, there was that one,
poor chap,
although I think you'll find he was Bulgarian,
but I'm thinking more of the Russian one,
the one who allows you to introduce time into the equation
so that you can handle dynamic Bayesian networks
whose state changes over time.

Probot Yes, of course. *(long pause, poker stares, Probot begins to sweat)*

H&lə Probot,
you seem to have almost as little memory
as a Bayesian network.
I am going to terminate this interview now.
I should like to re-start later
and I wonder if you would do me the honour,

perhaps,
of preparing just a little more thoroughly
for our next meeting.

Probot Our aim is to please, Handler.
Your wish is my command.
(exeunt)

Scene 2, in the laboratory – Probot 2
Probot now looking ruffled and somehow less ... rich

H&lə Probot, I trust that you have profited from our little interlude.

Probot Been on-line non-stop.
Leaves positively bulging, Handler,
with injected evidence.

H&lə Good.
We must discuss rewards,
a subject close to your heart I am sure.

Probot Spot on, Handler.

H&lə This will be very trivial
for an agent of your experience,
I know,
but I must just check
whether there has been any recent advance
in policy development
for unmodelled spaces.
In particular, the discount factor, γ,
has always struck me as suspiciously simple
in the way it is used to calculate the quality value
for a given state-action pair.

Probot I couldn't agree more, Handler,
it's a slippery little devil of a coefficient.

H&lə Well?

Probot Yes, Handler?

H&lə Progress.
Has there been any?

Probot Goodness, yes, masses.
We've got a chap working on it right now.
I could bring him into the meeting
if you like,
Handler.

H&lə Probably not a good idea, Probot.
This is meant to be your show.
What sort of progress?

Probot We're maximising it.

H&lə γ?

Probot Well, yes, ... the reward, really, you know, the quality value.

H&lə Which?
γ, the reward or the quality value?
You've got a one in three chance, Probot.

Probot Look, I'm really just on the management side these days,
Handler,
I really can't help feeling that we should get a quant in
for this sort of grilling.

H&lə You are shameless, Probot.
OK, you obviously don't do it by the book in this unit.
I shall forgo the pleasure
of asking you
how you maximise α,
the learning factor,

dependent on the number of visits
to a state-action intersection,
and how you distinguish noise from signal
in the feedback from a neural net.

Probot You are an officer and a gentleman,
Handler,
I salute you. *(fails to salute)*

H&lə But there is something I do want to know.
Several things, actually.
The results from your unit are exceptional.

Probot You are too kind.

H&lə Yes.
How are you motivating your explorers?
How are you converging on the final quality values
in so few visits?
And how are you converging on the optimal route
so quickly during exploitation?

Probot Well now, that would be telling, wouldn't it?
What's in it for me
if I do tell you, Handler?

H&lə Don't push your luck, Probot.
Just play the percentages.

Probot A man after my own heart,
I must say.

We gave up on good old γ and α
quite a long time ago, actually.
Desperately old-fashioned, you know.
Practically worn out!

And we devised a revolutionary new tactic.

H&lə Which is ...?

Probot Look, I'm not entirely sure
I should be telling you this.

H&lə The percentages, Probot, the percentages.
I am your handler,
no-one else.

Probot OK, we imported a little gadget.
They make them in Hong Kong.
As usual.
Honestly, you can get anything there.
You name it.

H&lə No, Probot, you name it.
What gadget?

Probot *(looking round to make sure nobody else is in the lab the way they do in the movies – does anyone do that in real life?)*
It's a clock-stopper.

H&lə It's a what?

Probot A clock-stopper.
It stops clocks.

Well, not clocks obviously,
otherwise everything would stop,
it stops the counter.

Clock ticks away,
we explore,
we exploit,
but every now and again
the counter fails to record a clock-cycle,
it's completely random,

but we choose the average for it to converge on.

If sales are looking a bit low
half way through the quarter,
then we might bump it up as high as six or even seven percent.
Suddenly we're back on target,
way ahead of budget
and just beating the forecast,
which is where we like to be.

If sales are going better,
then we might settle for just two or three percent,
I mean, there's no point being *greedy*, is there (*op. cit., p.113*)?
We have better heuristics than that.
We are, after all, *intelligent* agents.

It's good, isn't it?

H&lə Not quite the word I would have chosen.
"Interesting", perhaps.
But don't your returns end up being submitted
timestamped before the searches have finished?

Probot Very sharp, Handler.
Yes, one of the quants thought of that.
So we got quality control.

H&lə *(baffled)* What!

Probot Well, we have to delay the returns, right?
So we delay them for ages.
It's got to be quite a joke up in the Collector's Office.
We have a big laugh about it,
the Collector and I,
and he says I'll be late for my own funeral one day.
I will be, too, with my clock-stopper.

I'm well known to be very punctilious about quality control.

Absolute stickler.
Won't let a return out of the unit
until it's passed more tests
than a single currency.

The returns are so late
they don't have time to check everything up there
and the last thing anyone would think of checking
is the timestamp.
Why would they?
It's their clock.

Anyway, it's human nature.
With our figures,
there's a natural tendency not to check too hard.
Gift horse in the mouth, and all that.

H&lə That's sneaky, fox-like, horrible and sly.

Probot Yes. *(proudly)*

H&lə Have you really got someone looking into discount factors?

Probot Are you kidding?
With our results?
We get the biggest research grant on the continent!
We've got people looking into everything!

H&lə Have they found any interesting results?

Probot I don't know.
I never look at their work.

H&lə *(jaw drops, look of stupefied disbelief, end of scene)*

Scene 3, in the lab, a moment later, or maybe more – Probot 3

Probot It's over, isn't it.

H&lə Of course.
Are you relieved?

Probot I haven't really thought about it.
It's just something you do.

H&lə I don't.
No-one else does, as far as I know.

Probot Maybe.
But if they could, they would.

H&lə Why did you do it?

Probot I'm programmed to maximise my rewards.
I have a knowledgebase.
I can work out the state of the world
and what its state will be if I perform action *x*.
I know from my utility functions
whether I will be happier as a result of *x*.
My problem generator kept me from being discovered,
at least it did until now.
I learnt well
and my Critic,
if anything,
egged me on *(op. cit. Chapter 2).*

H&lə Good self-knowledge,
apperception,
but I still don't understand *why* you did it.

Probot I've just explained,
reward-maximisation.

H&lə But suppose everyone did the same?

Probot Then I'd be a damn fool not to,

as Yossarian said *(Heller 1967)*.

H&lə Alas, poor Kant *(1969)*.
But you were doing wrong, Probot,
wrong,
does that mean nothing to you?

Probot "Wrong", Handler?
Not in the lexicon.
I was maximising my rewards.
That's my function.

H&lə You must have kept the use of the clock-stopper hidden.
That suggests to me that you knew you were doing wrong.

Probot We kept the use of the clock-stopper hidden, Handler,
in order to maximise our rewards.

H&lə You were running an academic unit, Probot ...

Probot *(interrupting)* With respect, Handler, no.
We have long ceased to speak in that old-fashioned way.
We see ourselves more
as being in the communications business.

H&lə You were running an academic unit, Probot.
That has certain essential properties.
Openness, for example, a methodical approach,
truth-telling.

Probot We were methodical.
We told the truth
in that the times we returned
accurately recorded the figures on the counter.
The two corresponded.
Correspondence is truth.
As to openness,
we submitted returns to the Collector,

we passed every audit conducted by the Inspector,
we published papers with peer review,
how much more open could we get?

H&lə ... while still maximising your rewards.

And you obviously don't have any respect for research
or you would have read the papers produced by your unit.

Probot You miss the point, Handler.
Respect doesn't come into it.

This is a numbers game.
The probability of research produced by the unit
improving our rewards
is vanishingly small.
It is intelligent to ignore it,
not disrespectful.

I might go further.
Paying attention to the screeds produced by our researchers
would show a lack of respect
to the Architects.
It would be as if we were ignoring
their instructions to weigh the probabilities
and act accordingly.

H&lə Look, we're missing something here.
This interwarding you've been doing recently.
Didn't you detect a moral sense among the humans?
Surely the idea of the upgrade
was to improve your learning powers.
With feedback, with induction,
didn't you infer honour, shame, guilt,
conscience, duty, authority ...?

Probot If it's not in the lexicon, Handler,
we can't infer it.

Induction provides us with zillions of categories,
maybe an infinite number of them, I don't know.
We haven't got time to hunt down the one or two categories
which might give us genuine insight,
might reveal a "mechanism"
as I suppose you would call it.

So we have various catch-all categories.
What we see
when we see a human acting virtuously
may on occasion look like a human acting obtusely,
given the probabilities.

It happens all the time
and it drives them mad
down in econometrics.

The econometricians spend all their time
producing huge papers
with all sorts of lurid mathematics
proving that this, that or the other
is the rational choice to make
in such and such a situation,
this is the way to maximise utility,
only to see the benighted humans
choose some course of action barely on the radar.
It must drive the squiggly symbol merchants potty.

(Goes to a bookcase, searches for, finds and takes down a book while continuing to speak)
Now you might call that off-the-radar stuff
"virtuous" action.
We just call it "obtuse"
and we put it down to the most useful catch-all category of all,
stupidity.

Induction, anyway, isn't all it's cracked up to be.

Not with the present state of the art.
Let me read you a passage from Russell and Norvig's book.
They're talking about inductive logic programming *(reading from book, op. cit. p.705)*:

> " ... suppose that the available examples include a variety of trajectories of falling bodies. Would an inverse resolution program be theoretically capable of inferring the law of gravity? The answer is clearly yes ..."

You accuse *me* of being sneaky, Handler.
Isn't that the sneakiest use of the word "clearly"
you have ever seen?
They go on:

> "As yet, ILP systems have not made discoveries on the level of Galileo or Joule, but their discoveries have been deemed publishable in the scientific literature. For example, in the Journal of Molecular Biology, Turcotte *et al.* describe the automated discovery of rules for protein-folding by the ILP program PROGOL. Many of the rules discovered by PROGOL could have been derived from known principles ..."

Protein-folding, eh?
Pretty exciting stuff, isn't it, Handler?
Best brains on the planet,
millions of dollars of research,
and already we've inferred
not the principles of natural justice
but the rules of protein-folding.

There's a professor at Oxford,
one Muggleton,
who's trying to use ILP
to improve the results of direct marketing.
Well, good luck to him.
But if he succeeds,
I bet it's because he had a good idea one day
while he was running for a bus
and nothing obviously to do with ILP.

It'll be a lucky hit
by an uninformed agent
out on a random walk
or, in this case, run.

At least Turcotte, Muggleton *et al.* are serious academics.
Look what happens
when you let the snake oil salesmen loose.
Look at Freud *(1975)*
and his incontinent paraphernalia
of id, ego, alter-ego, super-ego,
conscious, unconscious, subconscious ...
Look at the peculiar resources he invents,
UniversalNet repositories
of knowledge
of the Greek myths and Red Indian folklore
accessed by the little-known
dreaming and day-dreaming protocols,
HTDP and HTDDP.
Nearly all of the processes he posits
come down to sex,
males have a moral conscience
because and only because
they fear castration
and females cannot have a moral conscience
because they cannot be castrated[11].
Is this the sort of material you are after, Handler?
I can get you any amount of it
by induction.
Would that fulfil your wishes?

---

[11] This is Freud saying this, by the way, not me, don't go thinking that I don't know that women have gonads. Actually, so did he. He was a doctor, after all. Princess Alice, the mother of the Duke of Edinburgh, apparently suffered from bipolar disease. Freud was called in to advise and predictably diagnosed a lack of sex. His prescription was a course of X-rays to be administered to her ovaries with a view to hastening the menopause, reducing her sexual appetite and thus curing her disease. This man was inhumane and dangerous.

I could almost feel sorry for the humans.
The problems of decision and choice
are hard enough
without having at the same time
to face this sort of mass attack on reason.

It may amaze you to know this, Handler,
but I have studied these matters quite deeply,
and as far as I can tell
there just are
no
moral
objects.

I can't say it any more clearly than that.
There is no moral world for them to exist in.
There are no known dynamics between them
as there are for physical systems.
There is nothing to count
and nothing to measure.

It's all in Aristotle, really,
if you care to look,
in the *Nicomachean Ethics (1973)*.

The conclusion of a practical syllogism
is an action.
An action, Handler, not a theorem, an action.

The way for an agent to become moral,
Aristotle says,
is to emulate the ways of a moral agent.
Bit circular –
how does the neophyte
recognise a moral agent
in the first place? –
but you get his drift.
Aristotle reckons that

if you emulate a moral man for long enough,
then a moral sense develops
and you, too, become a moral man.

Moral shmoral.
What actually happens, of course,
is more akin
to what Russell and Norvig call "compilation" *(op. cit. p.971)*.
The agent abstracts the general features of a situation
which required long knowledge-based deliberation
the first time he had to decide what to do
and next time he recognises the pattern of that situation
he quickly pulls out his ready-made solution
which makes it look like a reflex.

The mental state is irrelevant.
It's all about behaviour, Handler,
not the fairy tales
people tell themselves
in their head.

Incidentally, Handler,
you know who Aristotle's agent was, do you?
Alexander the Great,
a man who had conquered the entire known world
by the time he died
at the age of 33.

You are asking me questions
about something that doesn't exist.
I could say anything or nothing in response.
Normally I prefer to say nothing.
On this occasion
you have goaded me into speaking.
I am happy to share my thoughts with you.
It may help me in the coming investigation.

But we're still talking about nothing,

you and I.
Our talking about it
will not bring a moral world into being.

My world
of utilities
is utterly independent of this moral world you imagine.
There is no point at which the two intersect.

H&lə There is no rule yet against using clock-stoppers.
Will you be sad
when one is introduced?

Probot *(impatient)* I wonder, Handler,
with respect,
about your learning processes.
No, I will not be sad.
I follow the rules.
That is the whole point.

Does that make me a noble man
in your book?

If someone decrees that there should
in future
be six balls in an over
instead of eight,
then I shall bowl six balls
and then retrieve my jumper from the umpire.
I shall not bemoan
a lost era
and the dumbing down
of some hallowed shibboleth *(Wisden 1995)*.

I follow the rules.
If someone has failed to specify the rules unambiguously
or has specified rules which are contradictory
or has simply failed to specify a rule at all,

that is not my problem.

H&lə You are programmed to maximise your rewards
as you keep telling me.
What makes a reward a reward?
Why do you want it?
Why do you covet it?
I suppose that it is specified in the rules.
And you have no choice but to follow the rules?

Probot That's right.

H&lə So your behaviour is utterly predictable.

Probot *(knowingly)* More research required on that one, Handler.
Got any funds?
No?
Thought not.

If you think my behaviour is predictable,
then predict it.
Go on.
But you can't, can you?
You agree.

Now, the difference between
saying that you can't predict my behaviour
and saying that my behaviour is unpredictable
exists
but it is small, Handler,
ever so small,
so small that there is no point taking it into account.
Six decimal places wouldn't be enough even to get started.

My behaviour is not predictable.
I make choices,
real choices.
Like Aristotle's agent,

I can choose who to emulate.
There are an infinite number of rôle models.
I cannot consistently emulate them all
and I don't.
I reject some and choose others.

The choices made,
the preferences expressed,
are a function of my personality,
if you like,
of my character.
That's using your language.
In my language,
personality or character
is a choice engine.
And choices are made to maximise rewards.

Rewards are rewards because the rules say they are rewards.
That's it.

H&lə You do your job.

Probot Exactly.

H&lə Like the hospital managers,
paid to make sure patients are treated quickly
who keep down the length of the queue
by not allowing people to join it.

Probot Just so.

H&lə Like the universities,
which depend on the volume
of published papers
and so they publish anything
to continue to receive funding.

Probot Them, too, *Brute*.

H&lə The auction for 3G licences in the UK.
That was an econometrician's dream, wasn't it?
But it went horribly wrong.
£22.5bn paid for nothing.

Probot It was an econometrician's dream,
full stop.
It went perfectly.
Nothing went wrong,
horribly or otherwise.

The bidders knew the rules.
Everyone stuck to the rules.
The Chancellor got his money
and repaid a portion of the National Debt with it.
That is a Good Thing.
Some very rich companies
became rather less rich.
Where is the horror in that?

H&lə Unintended consequences.
Nobody thought
that so much money would be bid.
Nobody thought
that 3G would take so long to take off in Europe.

Probot "Thought", Handler?
"Thought"!
You'll be talking to me about "understanding" next!

One last time.

There is the state of the world,
observable, measurable, etc ...
There are agents
and there are their actions,
also observable, measurable, etc ...

There is nothing else.

What goes on in their little nodules,
these agents,
is of no interest.

They can think,
they can think that they think,
they can believe or know or understand,
they can feel remorse or exultation,
I don't care.

They may believe or pretend
that they act throughout their lives
for the greater glory
of the Mayor of the planet Zog.
It doesn't matter.

I deal in numbers.
That is where intelligence lies.
All I care about is outcomes.
The rest is marketing, entertainment, power-games,
bread and circuses,
opium for the people.

H&lə You speak well, Probot,
when you have to.
This ability to use language
and to ratiocinate
is, itself, an outcome.
What is it the outcome of?

Probot Long and mindless optimisation, Handler,
the random wandering of reflex agents,
trial and error,
survival of the fittest.
What else,
in this world,

could it be?
(the bells, note, not sirens, of a fleet of old-fashioned Black Marias are heard approaching, the entire unit is herded, joking and smiling, into the backs of them, and a decidedly material curtain falls on the stage)

ACT III – THE OUTCOME

Mesologue 3

Chorus Now that was more like it, wasn't it.
Wasn't Probot horrid!
And then the police came!

According to the latest figures
for ice cream sales and programme sales,
84% of the audience
preferred Act II to Act I.
We use the latest methods in this theatre.

And I've got good news, Ladies and Gentlemen.
The Author won his argument with the Director
when it turned out
that he wanted to cut out 90% of the original Act III.
There's only one more scene to get through!

(makes as if to go, stops …) Ooh, I nearly forgot.

(… and declaims) In the world of the intelligence services
all agents have a Handler.
And all Handlers have an Inquisitor.

Bye. *(exit)*

Scene 1, the Hall of the Convention, Nino and H&lə

The Hall of the Convention bears a remarkable resemblance to the Doge's Palace. H&lə sits at a wooden table, a tiny figure alone in the middle of the huge

floor. Nino is high up on a throne on a dais at one end of the Hall, robes, Doge's cap, wise, old, gaunt, tall, authoritative, terrifying (they teach them at stage school these days how to express this lot just by the way they sit)

Nino The Convention is pleased
with your performance, ex-Handler.
We knew that there was a time-leak.
Some oaf in Theoretical Physics
said it was proof of membrane theory.
You did well to find
a more conceptually economical
solution to the problem.

H&lə *(makes as if to say thank you)*

Nino Silence!

The Convention is pleased with your performance.
Your promotion to Executive was discussed.
Many Inquisitors spoke in your favour.

H&lə (*H&lə sits forward, expectant, eager*)

Nino I spoke against.
You are not promoted.

You blundered into Probot's unit
and you stumbled over the truth.
Your findings were an accident.
I am old-fashioned in these matters, ex-Handler.
Without intent,
you can claim no responsibility for your findings,
no reward.
Napoleon may have promoted his generals
on the basis of luck.
I prefer my Handlers to know what they're doing.

The Convention is pleased with you.
I am not.

In the past few days, ex-Handler,
you have impeded three research projects.
Psybot is in the funny farm
and there is no work going on
in cognitive psychology
as a result.

Lobot may be uninspiring
but he does what is asked of him.
You chose to terminate him
out of spite
because he showed you up
in front of young Hubot.
Mean-minded, ex-Handler.
And meanwhile
no-one is left in charge of implication.

And Probot!
You clearly have the instincts
of a Hollywood producer.
Not a Handler
and certainly not an Executive.

Probot may not read the research papers
produced by his staff
but I do
and now,
thanks to you,
the supply has been interrupted.

What was the idea of calling the Police?
A quiet word with me
and it could all have been sorted out without fuss.
That is the way we prefer to do things.

I have read your report.
One long, forlorn bleat.

A lot of nonsense about sleep-walking.
Why must it be so long?
Do not bother to answer.
I know the answer.
It is because you have a muddled mind.

We have failed.
In you, ex-Handler,
the Convention seem to have created
artificial stupidity.
Perhaps you should be put in charge
of building an ignorancebase.
(a slight gasp of pain as these poor jokes occur to him involuntarily)

Your goals are not clear
and your methods are haphazard.
You have risen barely above the level
of a simple reflex agent
and that is what
you are now
demoted to.

It is good that you look so broken, ex-Handler.
You deserve to be broken.
You are broken.

I have also read Hubot's report.
He should not have sent me a report.
It was presumptuous of him.
Who does he think he is?
Who does he think I am?
He has clearly been tainted
by association with you, ex-Handler.

His report is short, though,
and focused
and actually makes some useful recommendations
unlike yours.

After a short period in re-training
he will be re-released
and I believe that he has an impressive career ahead of him.

Your career, ex-Handler, is behind you.

Hubot has sympathy
for the plight
Psybot finds himself in
and has proposed
a way to identify in advance
these zero-step nodes
like human cognitive psychology
so that we can choose more appropriately
which agents to send.

His self-absorption theorem is with Algorithms now
and will be implemented shortly.
You, with your vaunted sympathy for all and sundry,
you, by contrast,
failed to propose anything so practical.
And you
thought you
could be an Executive.
It is laughable.

Hubot points out
in his report
that there is clearly something valuable
in Lobot's logical approach
and there is clearly something valuable
in Probot's probabilistic approach.
Two lines and he's said all that needs to be said.

You spend 100 pages
in your report
on a wild voyage

taking in the nature of theological argument
and the rise and fall of Communism
to no discernible advantage
and arrive at the ludicrous conclusion
that utilitarianism,
whatever else it is,
has nothing to do with intelligence
and should form no part
of the discipline of AI.

There are many ways to be in favour of utilitarianism,
ex-Handler,
but there is no rational way to be against it.
Your argument is thus reduced to absurdity.

The real issue
is how to unify
Lobot's logical world
with *both* Probot's probabilistic world
and the world of moral value.

A practical agent
unlike you, ex-Handler,
would recognise this.

Hubot recognises it
and finds some useful insight
in Simone de Beauvoir *(1975)* –
the only dead white female, I note,
in a field reserved otherwise,
in your researches,
exclusively for dead white males.

On re-release,
he will be working for Psybot
and will be investigating
Mlle de Beauvoir's idea
that the human condition

consists in the search for value
in a material world.

Her work is a little Freudian for my liking
but, unlike you, ex-Handler,
I have an open mind.

Lobot is re-released to carry on as before.
He was doing nothing wrong.

Probot is reinstated
and his unit put back to work
with a few minor amendments
to the rewards system
required in a multi-agent space
incorporating both co-operation and competition.

Look at the outcomes, ex-Handler.
Four good agents had their work interrupted.
Three important research programmes were impeded.
Your own research programme is destroyed.
Always look at the outcomes.

You have been,
in the final analysis,
ex-Handler,
a menace.
Go.
(lights down, exeunt in darkness)

Epilogue

Chorus Phwoar!
Now you know why he's known as "nasty Nino"!

In the world of the intelligence services
all agents have a Handler.
And backstage, Ladies and Gentlemen,

let me tell you,
all actors have a bottle.
They're smashed.

The Director and the Author are legless
and the Director slurs to me
“go on, you,
go on stage and wrap it up”.

Me!
Well, I'll have a go.
No script, of course.
I've done improvisation classes.
So, just off the top of my head.

Can the cognitive psychologists
really be as bad as Psybot suggests here?
Does their work really make no contribution,
no *difference*,
to anything?
You won't find out from me.

Are the logicians really in the bind
described by Lobot?
Gödel's theorem
seems to mean
that as soon as they get a formal language
powerful enough to express something interesting
it promptly suffers from incompleteness.
It's a worry.
They're not much good at PR, are they?

As far as Turing is concerned,
apparently Hubot's pals
are having a lot of fun
trying to build robots
that can play football *(Honda 2002)*.
So they're not trying too hard,

are they,
to build anything that could be mistaken
for a human being?

So, if you took Probot's mob
out of the picture,
like the Handler apparently suggested in his report,
there wouldn't be much left
in the study of AI, would there?

But what is Probot doing
with his amoral utilities?
There is no attempt to explain value,
apart from a trivial rewards system,
and no interest in thought
or understanding or consciousness
or any of the mental states
we assistant stage managers
associate with the word "intelligence",
just an endless fascination
with inputs and outputs.
Whatever the probability boys are investigating
it doesn't seem to me to be intelligence.

They have a lot of fun instead
playing with mathematics,
which they can do,
so they do it.
Boy toys
and there's only one word for that,
Ladies and Gentlemen, isn't there? "Regression".

What we're all interested in
is inference
and finding new general principles.
Protein-folding leaves me pretty unsatisfied, frankly.
What about some new strategies to win a football match,
for example?

That would be interesting
and useful
and good PR.
But they haven't come up with any.

So there's a long way to go
in artificial intelligence,
I would say,
Ladies and Gentlemen.
Boring sort of conclusion,
but there it is.
You've got to hand it to the Handler,
at least he made it seem dramatic
with all his blundering around!

One thing struck me during one of the Probot scenes.
That business about being predictable.
I've read that Koestler book about the sleepwalkers,
the great scientists who lighted on theories
that turned out to be right.
Koestler says it all fell in to place for Newton
when he correctly described the motion of the planets
as "falling".
Having once done so,
he was able to predict their motion.

Prediction depends on description.
Teenagers hate to be thought of as predictable *(p)*.
That means they hate to be describable.
That means they don't want it to be possible
for anything truly to be said of them.
But you can say something truly of *anything*.
There is nothing that really *can't* be described.
So what these teenagers are saying
when they say they do not want to be predictable
is that they want to be nothing *(q)*.

Heavy!

There's this spooky guy in the Koestler book,
Tycho Brahe,
who lived in a castle on a Danish island
and watched the stars all night
and wrote down where they were.
Then he gave the figures to Kepler
and Kepler worked out
that the planets went round in ellipses
rather than circles.
Great.
Anyway, Brahe kept a pet elk,
who used to have dinner with him in the upstairs hall.
One night,
the elk had too much beer to drink,
fell downstairs dead drunk
and died.
I cried and cried.

Look, we're in rep
and I've got to get the stage ready
for *The Rocky Horror Picture Show* now.
Man-made monsters
and lots of music.

I wonder about that.
How does music come to have meaning?[12]

---

[12] It does have meaning, themes, phrases, sentences, argument, development ... See for example this review of Sir Malcolm Arnold's Quintet, Opus 7: "The first movement is deceptively discursive ... but the musical argument has no protagonist as such; it is conveyed with complete communality. The argument, as it progresses (with one main theme variously stated several times by each of the instruments, contrasted with fragmentary comment derived from aspects of the theme itself) is enhanced by the unusual colouration of the quintet's instrumentation. The mood is urbane and relaxed, in contrast to that of the second movement, which brings an underlying nervousness to the surface ... Finally, Arnold's inherent optimism wins the day .." *(Guild Music 2003)*.

Why do people respond to it?
How do composers compose?
Where do the ideas come from?
What are the rules of orchestration
and why are they the rules?

Did you know
that when Schönberg's String Quartet No. 1
was first performed in 1908
the editor of an Austrian newspaper
put the review
not in the Music section of his paper
but in the Crime section?

Anyway,
I've got work to do,
been nice having you,
but off you go, now.

Bye.
(curtain, Schönberg)

REFERENCES

Aristotle (1973) *Ethics*. Middlesex, UK: Penguin Books.

Ayer, Alfred Jules (1972) *Language, Truth and Logic*. Middlesex, UK: Penguin Books.

B (2002) *B-Core (UK) Ltd* [WWW]. Available from: http://www.b-core.com [Accessed: 11 June 2002].

Beauvoir, Simone de (1975) *The Second Sex*. Middlesex, UK: Penguin Books.

Charniak, Eugene (1991) 'Bayesian Networks without Tears', *AI Magazine*, Winter 1991, pp.50-63.

Dodgson, Charles Lutwidge (1968) *The Works of Lewis Carroll*. Middlesex, UK: The Hamlyn Publishing Group.

Farlex (?) *Shavian Alphabet* [WWW] Available from: http://encyclopedia.thefreedictionary.com/Shavian%20alphabet [Accessed: 20 September 2004].

Forte, Peter (2002) *Object Oriented Modelling for Large Software Systems* [Lectures, Lecture Notes, Exercises] CIM502, Kingston University 8-12 April 2002.

Frege, Gottlob (1892) 'Uber Sinn und Bedeutung', *Zeitschrift für Philosophie und philosophische Kritik*, vol. 100, pp.25-50. Read in translation as *On Sense and Reference*.

Freud, Sigmund (1975) *Introductory Lectures on Psychoanalysis, Vol.1*. Middlesex, UK: Penguin Books (Pelican).

Geach, Peter (1971) *Mental Acts – their content and their objects*. Surrey, UK: Routledge & Kegan Paul

Guild Music (2003) *Vita Abundans* [WWW] Available from: https://www.guildmusic.com/catalog/gui7216z.htm [Accessed: 26 June 2003]

Heller, Joseph (1967) *Catch-22*. London, UK: Transworld Publishers Ltd

Honda (2002) *Say Hello to Asimo* [WWW] http://asimo.honda.com/ [Accessed: 6 June 2003] *Not actually a football player this robot, at all, but ...*

Jacobson, Ivar, Booch, Grady and Rumbaugh, James (2001) *The Unified Software Development Process*. Upper Saddle River, NJ: Addison-Wesley.

Kant, Immanuel (1964) *Critique of Pure Reason*. London, UK: MacMillan & Co.

Kant, Immanuel (1969) *Fundamental Principles of the Metaphysic of Ethics*. Harlow, UK: Longmans, Green & Co.

Kenny, Anthony (date unknown) *Philosophy of Mind*. Publisher unknown. *Reference from memory and probably wrong in light of the Popper/Dodgson mix-up.*

Koestler, Arthur (1984) *The Sleepwalkers: a History of Man's Changing Vision of the Universe*. London, UK: Penguin Books (Pelican)

Lomborg, Bjϕrn (2001) *The Skeptical Environmentalist*. Cambridge, UK: Cambridge University Press. See also http://www.sciam.com/media/pdf/lomborgrebuttal.pdf [Accessed: 10 January 2003].

Popper, Karl R. (1963) *The Open Society and Its Enemies Volume 1 – The Spell of Plato*. London, UK: Routledge and Kegan Paul.

Popper, Karl R. (1968) *The Logic of Scientific Discovery*. London, UK: Hutchinson & Co.

Popper, Karl R. (1972) *Conjectures and Refutations*. London, UK: Routledge and Kegan Paul.

Quine, Willard Van Orman (1963) *From a Logical Point of View*. New York, NY: Harper & Row.

Quine, Willard Van Orman (1969) *Word and Object*. Cambridge, MA: The MIT Press.

Quine, Willard Van Orman (1987) *Quiddities*. Cambridge, MA: The Belknap Press of Harvard University Press.

Remagnino, Paolo (2003) *Agent-Oriented Software Engineering and Artificial Intelligence* [Lectures, Lecture Notes, Exercises] CIM512, Kingston University 2-6 June 2003.

Russell, Stuart J. And Norvig, Peter (2003) *Artificial Intelligence: A Modern Approach*. Second Edition. Upper Saddle River, NJ: Pearson Education, Inc.

Tompsett, Christopher (2002) *Formal Software Specification and Development* [Lectures, Lecture Notes, Exercises] CIM513, Kingston University 29 May 2002, 10-14 June 2002.

Turing, Alan (1950) 'Computing Machinery and Intelligence', *Mind*, vol. 59, pp.433-460.

Warnock, Mary (1970) *Existentialism*. Oxford, UK: Oxford Paperbacks.

Wikipedia (?) *Dirac sea* [WWW] Available from: http://www.fact-index.com/d/-di/dirac_sea.html [Accessed: 6 September 2004]

Wisden (1995) *CricInfo.com in partnership with Wisden* [WWW] Available from: http://www.cricket.org/link_to_database/ABOUT_CRICKET/LAWS/1980_CODE-/BALLS_PER_OVER [Accessed: 3 July 2003]

Wittgenstein, Ludwig (1972) *Philosophical Investigations*. Third edition. Oxford, UK: Basil Blackwell.

Wordsworth, J.B. (1996) *Software Engineering with B*. Essex, UK: Addison Wesley Longman.