# In the Eye of the Beholder?

## Another look at Cognitive Systems

Nello Cristianini

University of Bristol

# Cognitive Systems, Artificial Intelligence, and all that …

- The past decade has seen significant technological progress in computational pattern analysis, machine learning, probabilistic data-driven AI.

- Can we use this to go back to the task of understanding and designing cognitive systems?

# Putting all Together…

- Data driven approaches to AI are the success story of the past decade, they are behind Google, Amazon, and much more
- Robotics has been progressing in very interesting directions
- Our understanding of biological systems has been improving

- **Can we be part of the real excitement?**
  **Putting all this together, to address the big questions of AI**

**REMINDER TO SELF:**
**you did not start this career**
**because you wanted to become a statistician**

# Modelling Cognitive Systems

- Biological systems are made up by many interacting parts (metabolic networks, ant colonies, nervous systems, …)
- Robots are often made up by smaller parts (swarms, subsumption architectures, etc…)
- Web systems are made up of many interconnected things

- How does learning and decision making take place, without a central representation?
- Can we *model* that with our current technology?
- Can we *build* intelligent systems following these principles?

# “Pascal” Tools → Impact on AI

- The question: how can PASCAL approaches (pattern analysis, statistical learning, probabilistic AI) contribute to our understanding and design of real cognitive systems?
- We will hold **a 6 months Thematic Programme** on this topic…
  [see monday’s talk]

- Today: will just try to suggest some informal ideas, to prepare for the discussion on monday

# A First Question

- Let us spend some time to clarify the scope of this research programme.
- **What do I mean by cognitive behaviour?**
(much less than you probably do)

- Is cognitive behaviour inherently of a different type from mechanical, reactive, stimulus-response behaviour?
- Is cognitive behaviour found only in people?

- We need to draw a line somewhere…

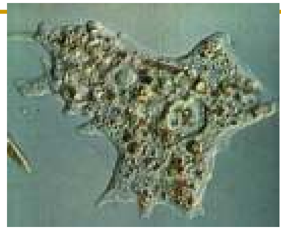

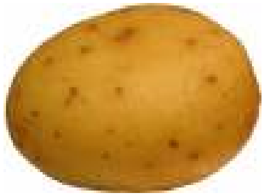

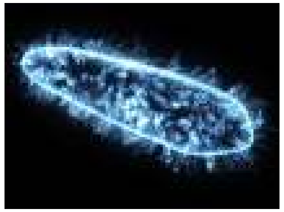

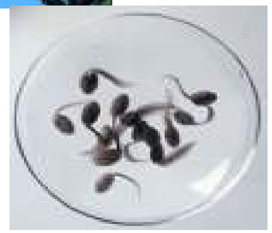

Who *knows*?

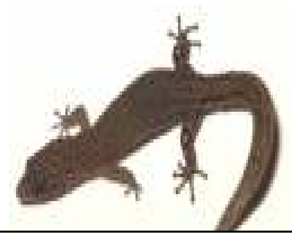

6 obvious slides will allow me to frame the discussion and establish some terminology

# Obviously (1/6)

- The occupation of increasingly complex ecological niches requires the creation of increasingly sophisticated behaviour.
- Organisms that can rely on a constant source of energy, may not need to develop much decision making capabilities, and can ‘hardwire’ the assumption that energy will be there in a given form
- Organisms that depend on multiple sources of energy, need to be able to recognize and exploit opportunities, by “doing the right thing at the right time”.
  - The benefit: they can enter a niche that is inaccessible to ‘simpler’ competitors.
  - The cost: they need to be able to make decisions between alternative “actions”. The alternative actions may be hardwired, the choice is based on information from the environment.

# Obviously (2/6)

- Regularities, invariances, in the environment may be incorporated into the design of the organism. (NOTE: these are just *assumptions*, extinction follows in the niches where they do not hold)

- Predicting (implicitly) some aspect of the environment seems to be a fairly essential feature of life
- Three equivalent viewpoints:
  - the enzymes that can be produced by a cell dictate what energy sources can be accessed
  - they represent an assumption on the kind of nutrients that will be found in a given niche
  - they dictate which niches are accessible

# Obviously (3/6)

- Three equivalent viewpoints:
  - the enzymes that can be produced by a cell dictate what energy sources can be accessed
  - they represent an assumption on the kind of food that will be found in a given niche
  - they dictate which niches are accessible

- In evolution purpose (function) and effect of an innovation are often confused
- Teleological language is used (inappropriately?) when we say that the “function of fins” is to enable swimming

- This kind of language shapes our perception of cognitive systems too

# Obviously (4/6)

- Once all simple niches are filled, there may be opportunities that are not easily accessible, as they need more complex assumptions
- For example: there may be some niche that is accessible to an organism that can switch from photosynthesis to respiration to fermentation, in the right conditions…
- This environment would appear random to many simple cells, might appear manageable and predictable to a more complex one
- Information processing of environmental information gives an obvious advantage, and is found in all organisms

# Obviously (5/6)

- In an evolutionary arms race, it is obvious that one of the dimensions of the race will involve better decision making, better use of information, possibly better adaptability …

- We will call this the *cognitive dimension*, and is just one of the possible ones

  (together with others involving better use of materials, or better immune capabilities, time of replication, etc. etc)

# Obviously (6/6)

- It is clear that cells spend enormous amounts of energy in ‘software’ to deal with their environments
- And it is known that once a simpler niche is conquered – perhaps by a parasite – entire pathways are lost, and genomes are greatly simplified…

# My working definition…

- For the purposes of the thematic programme and of this talk, I will use this working definition of cognitive system (one that perhaps covers nearly every living thing too)…

- **Cognitive systems: exploit information about the state of their environment in order to make decisions, select actions and pursue their goals…**
- **Their *behaviour* cannot be understood only based on the environment, one needs to consider their goals, possibly their history, in order to understand and predict it**

- The ***capability of making informed decisions by exploiting environmental information and background knowledge (assumptions)*** seems to be crucial…

- I am sure human philosophers will have many counterexamples…

# Cognition: not human / possibly mechanical

- Philosophers have been trying for a long time to draw a meaningful line between human-like cognition and all the rest of the world…
- Let us ***focus on non-human*** cognition instead…

- One very debate direction has always been:
can a *mechanism* be cognitive ??

- It has often been assumed that - nearly by definition – mechanical reactive systems do not display cognition.

- How can you predict that 2 eels will meet next year in the Sargass sea? How about 2 robots?

[can behaviour generation / explanation be a theoretical focus for us?]

# Alan Turing, 1947

- National Physics Laboratory, UK, Summer 1947

NATIONAL PHYSICAL LABORATORY

INTELLIGENT MACHINERY

"Intelligent Machinery".

I propose to investigate the question as to whether it is possible for machinery to show intelligent behaviour. It is usually assumed without argument that it is not possible. Common catch phrases such as 'acting like a machine', 'purely mechanical behaviour' reveal this common attitude. It is not difficult to see why such an attitude should have arisen. Some of the reasons are

A Report

by

A.M. TURING

# Alan Turing, 1947

- The report remained unpublished, however forming the seed of the celebrated 1950 article on MIND, where the Turing Test for intelligence is introduced…
- Among other things, it touches on machine learning, connectionism, and genetic algorithms
- But what is more important for us: the ideas in the 1947 report show us what was Turing’s basic insight about intelligent behaviour…

The extent to which we regard something as behaving in an intelligent manner is determined as much by our own state of mind and training as by the properties of the object under consideration. If we are able to explain and predict its behaviour or if there seems to be little underlying plan, we have little temptation to imagine intelligence. With the same object therefore it is possible that one man would consider it as intelligent and another would not; the second man would have found out the rules of its behaviour.

A first, not human-centric,
version of his famous test.
A bit like hypothesis testing:
can I explain the behaviour
in simple mechanical terms?

# Looks Intelligent to me…

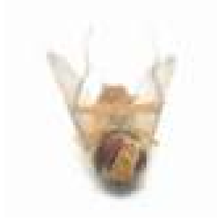

- If I see an insect pretending to be dead, I may think it is a rather clever thing…
- Should I change my mind when I see it (and all its friends) doing the same thing day after day, in the same way, even in the wrong context?
- What we know is that that behaviour is often useful to achieve a specific goal, if used in the right occasion

# Let us be inclusive…

- The fact that we find it hard to draw a line between cognitive systems and non-cognitive ones, is POSSIBLY that we insist in drawing it in the wrong place
- We want to rule out ‘lower’ life forms, but keep ‘higher’ ones, and then we have a problem in making a coherent distinction…
- It is perhaps much easier to decide that cognitive behaviour includes, among other things, life
- So a cell does show cognition
- And so does a plant
- And for increasingly complex niches, there is increasingly sophisticated behaviour

- This would make various parts of biology *directly* our concern

# The Eye of the Beholder

- If we can easily predict / understand the behaviour of a system in causal / mechanical terms, we may believe that there is little intelligence behind it

- What if we just find it *easier to describe* it in terms of goals, motives, intentions… ?

# 1- Information Processing

- "Informed" decision making
- It seems that we consider as a feature of cognitive systems the capability of **using information about the environment when deciding actions**.
- Processing environmental information to determine actions separates clocks from thermostats, for example.
- This implies some degree of "**awareness**" of the state of the immediate environment

Awareness:
knowledge of events and objects in the system's environment, and use of that knowledge to inform behaviour

# Not very restrictive…


- However this can still fall within a stimulus-response framework.
- A plant can open its flowers in response to certain stimuli. Is this a first step in a continuum of cognition?
- Besides:
  information processing takes place in many ways
  (chemical, mechanical, electronic, …)
  one could say that a dynamical system can process information…
  (more on this later)

# About Symbols

Signal Transduction **pathways** found within cells (or plants) effectively couple stimuli with responses,

But the coupling takes place via a chain of intermediaries, that have no other necessary property than to trigger each other… they are not used because of their absolute properties, but because of their relation to each other (and to the I/O signal)

- Lac Operon can produce enzymes to digest lactose
- A repressor of its transcription always produced by gene lacI
- when lactose is present, a lactose metabolite called allolactose binds to the repressor, causing a change in its shape
- and repression is blocked

**SHAPE OF LACTOSE REPRESSOR *represents* a state of the world**

[We will see later that the concentration level of a given chemical in a cell or plant can effectively be considered as the 'internal symbol' for some state of the world… or set of states…]

[Furthermore, long chains of intermediaries create the opportunity for MODULATIONS of the signal, memory, information integration and even generalization …. As we will see later]

# 2-Internal States

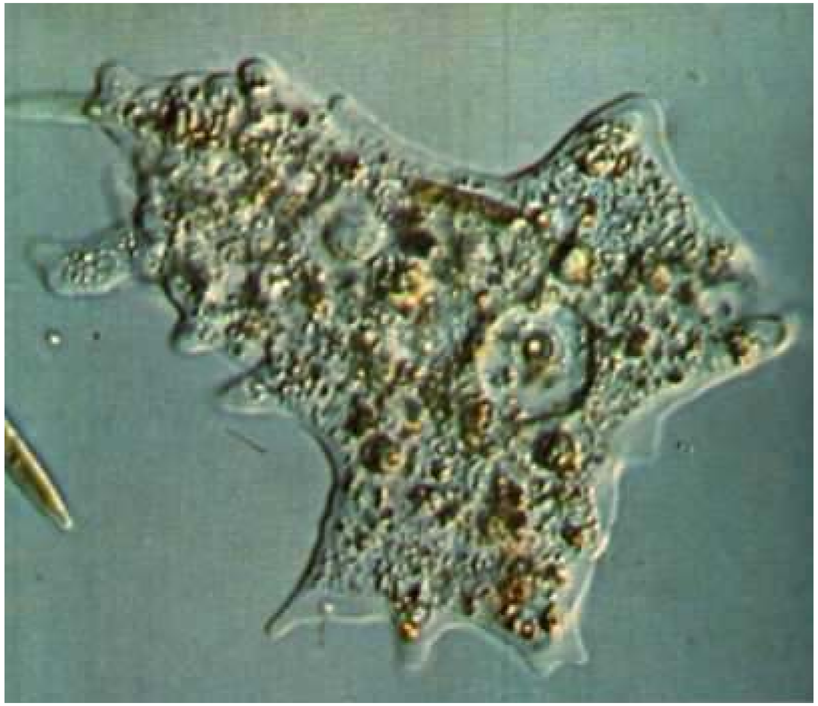

- It seems also that using information about past states of the environment marks another step along a continuum. (A sudden change in temperature can trigger reactions in many organisms, as opposed to a gradual increase.)
- Using past information implies having an ‘internal state’, that is not reacting purely to the environment, but also to its own state. Two systems might behave differently in the same conditions, depending on their history, or past experience.
- For example: [Habituation? Also bacterial **chemotaxis** is based on short term memory]. This also makes it very hard to predict the system’s behaviour, for an external observer. [how to implement state: concentration of a chemical for example, in metabolic networks]

# What Cells Think…

CELL COGNITION

- Cells can have a very rich behaviour. They can
  - seek food, escape danger, exploit opportunities.
  - know which food to eat first, and which enzymes to produce to digest a given food.
  - produce repair proteins when they sense there is DNA damage.

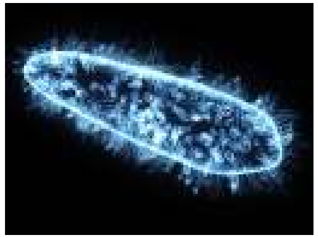

- The bacterium *E. coli* is 1 micron long and contains millions of proteins of about 4000 different types.
- New proteins are constantly produced, to respond to changes in the environment.
- Deciding which protein to produce and how much, is the result of a ‘computation’.

# Chemotaxis

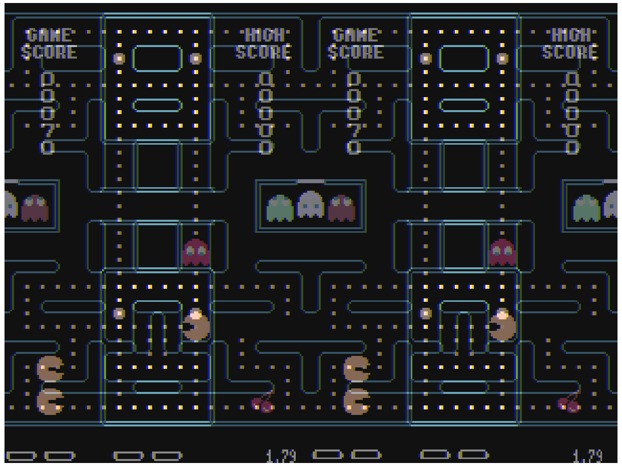


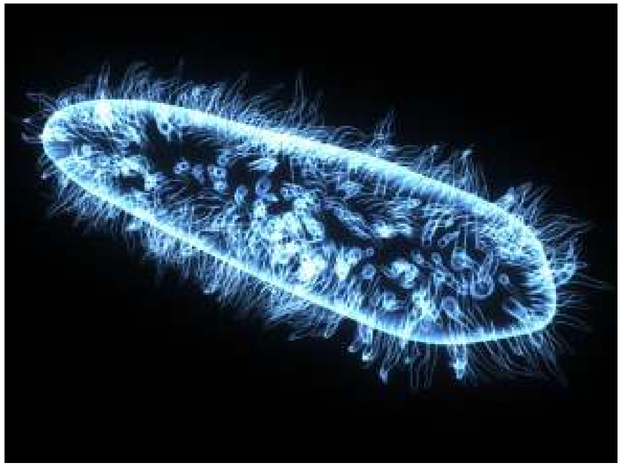

- Simple mechanisms can do very interesting things…
- Chemotaxis: seeking nutrients, escaping toxics…
- Simple pathway in bacteria…
  (implements a randomized gradient ascent strategy)

# What Cells Want…

- “Every cell’s dream is to become two cells”.
- They need to survive, fiercely competing for resources, exploiting opportunities and avoiding dangers…

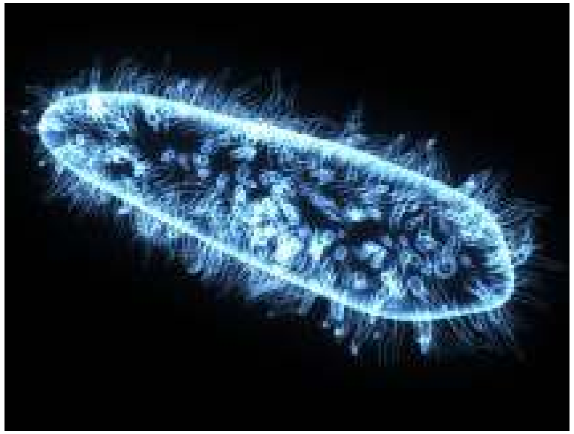

- All they **do** is activating the right ‘actuators’ or produce the right ‘chemicals’ at the right time.
- Both actuators and chemicals in cells are proteins, or are controlled by proteins. Also sensors (receptors) are proteins too.

# Cognition in Cells

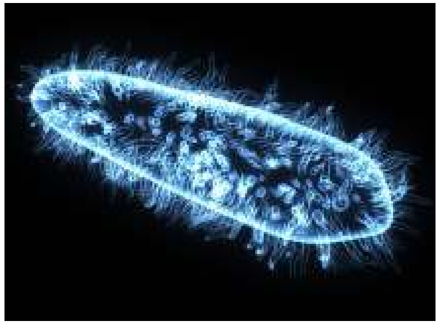

- Simple pathway:
  - Lac Operon can produce enzymes to digest lactose
  - A repressor of its transcriptior always produced by gene lacI
  - when lactose is present, a lac metabolite called allolactose binds to the repressor, causin change in its shape
  - and repression is blocked

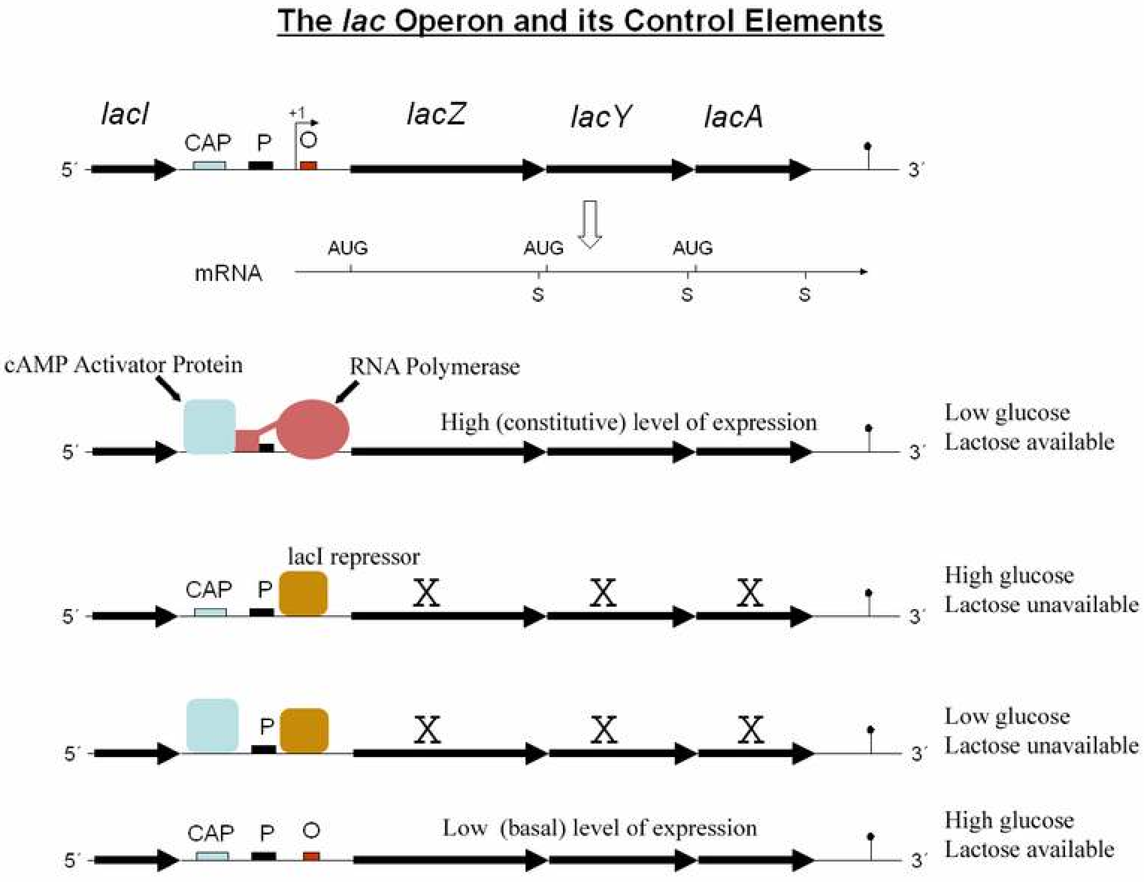

# Cognition in Cells

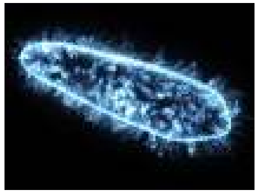

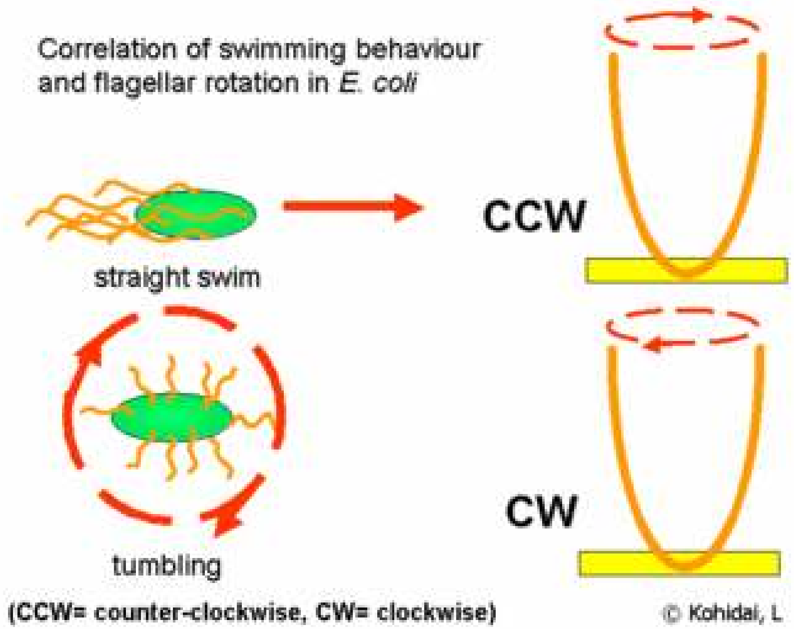


Randomized algorithm for gradient ascent:

- Flagellum can spin in 2 directions, one induces random stumbling behaviour, the other forward motion
- Time spent in forward motion depends on "quality of the direction"
- Quality of the direction:
  compare previous concentrations with current concentrations
- Net result: brownian-like motion towards the concentration peak

# Bacterial Chemotaxis

- Chemical gradients are sensed through multiple transmembrane receptors, called methyl accepting chemotaxis proteins (MCPs), which vary in the molecules that they detect.
- The signals from these receptors are transmitted across the plasma membrane into the cytosol, where *Che proteins* are activated.
- The Che proteins alter the tumbling frequency, and alter the receptors.

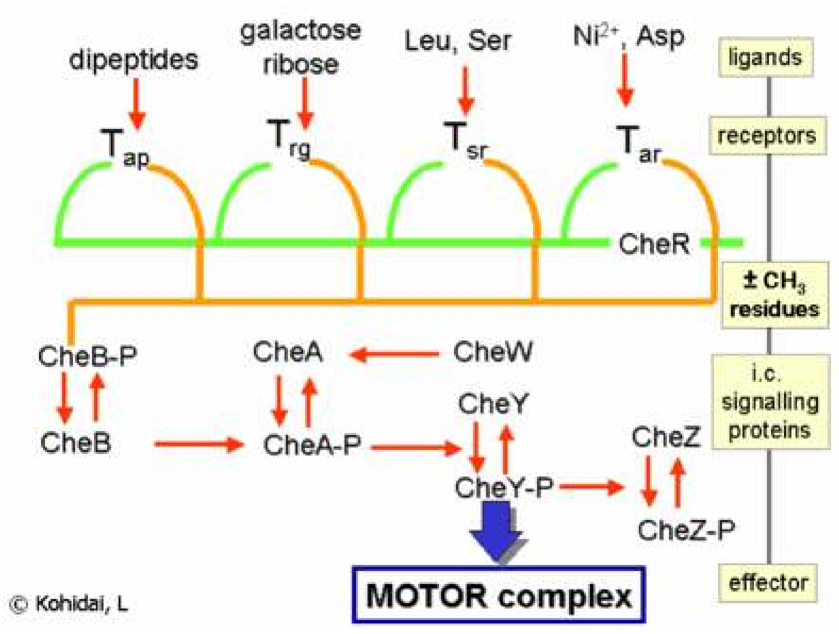


DETAILS

•The proteins CheW and CheA bind to the receptor.

•The activation of the receptor by an external stimulus causes autophosphorylation in the histidine kinase, CheA, at a single highly conserved histidine residue.

•CheA in turn transfers phosphoryl groups to conserved aspartate residues in the response regulators CheB and CheY

•CheY induces tumbling by interacting with the flagellar switch protein FliM, inducing a change from counter-clockwise to clockwise rotation of the flagellum.

•Change in the rotation state of a single flagellum can disrupt the entire flagella bundle and cause a tumble.

From Wikipedia

# How Cells Think…

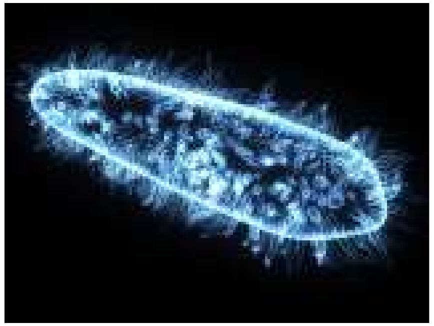

- What can a cell sense? Light, temperature, nutrients, toxics… chemical signals …internal parameters, DNA damage, … some cells can keep time with an internal clock …

- CERTAIN PROTEINS do not directly act on the environment: they simply control (trigger or inhibit) the production of other proteins. They are called **transcription factors**. In *E. coli* there are about 300 transcription factors, in a system with 4000 proteins.

- One way to interpret this is that most of the variation in the state space of the cell depends on these 300 parameters.
(in reality we should consider immediate (un-mediated) stimulus-responses, and also long pathways involving many TFs – but this is a rule of thumb…)

- They MODULATE the transcription rate of genes, hence controlling the production of the various proteins of the cell.

See for example Uri Alon's book

# Cognition in Cells

- CELLS do have internal representations, they do not just respond directly to stimuli.
- Concentration levels of various chemicals, and other quantities too, represent internal states
- They are a coarse representation of their internal and external environment, sometimes integrating multiple signals into the same state and possibly representing complex situations.
- [eg: compare with ‘low battery’ signal in laptop; this triggers decisions, vs – say – automatic adjustment of screen light to room lighting in laptops, etc…]

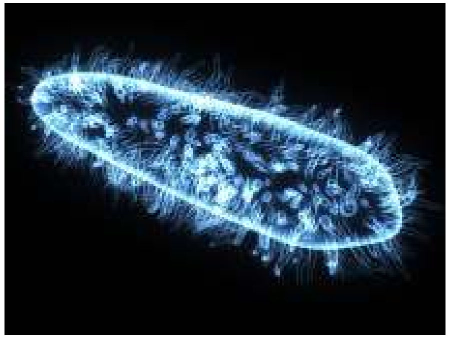

# How Cells Think…

- SIGNAL REPRESENTATION:
  signals can be changes in concentrations of certain chemicals (a sugar molecule entering the cell and binding to the TF, as in the lac operon), changes in shape of certain proteins (once a TF changes shape, it can become active or inactive, propagating the signal), chemical bonds (methylation), …

SIMILARLY IN PLANTS, COLONIES, SOCIETY, … etc:

- (…in plants signals can be concentrations of hormones, in animals perhaps electric potentials in nerves, but the concept is the same.
- In collective processing by a colony, the signal can be both in the state of members, and in the external marking – a strategy used by swarm robotics too.)

- **Is amount of money in a market a signal in the same sense?**
- **And pheromone markers in insect colonies?**

# 3-Information Integration

- It seems that the integration of multiple stimuli, in order to make a more complex decision, marks a further step.
- For example a security flood-light will be turned on just when there is **movement and it is dark**.
- Information integration can make the behaviour of a system highly unpredictable, **especially in a complex environment**.
- Full interfacing with environment needed in artificial systems

- May also create more effective and customized behaviour, in such an environment…
- Plant growth … (think like a plant)…

# Clever Plants

- Viewpoints…
- What do plants do?
  - Development same as behaviour for them
  - How do they decide what branches to grow?
  - A combination of light, water, soil, internal clock, etc… etc…
  - Integration of information
- About Ottoline Leyser's work very relevant to us …
  (also: rosalind franklin award of royal society)

# Clever Plants

- How do plants decide which branches to expand?
- A mix of hormonal signals, integrating information about other active meristems, soil situation, light situation
- It is mediated by genetic information (targeted mutations can change the decision algorithm)

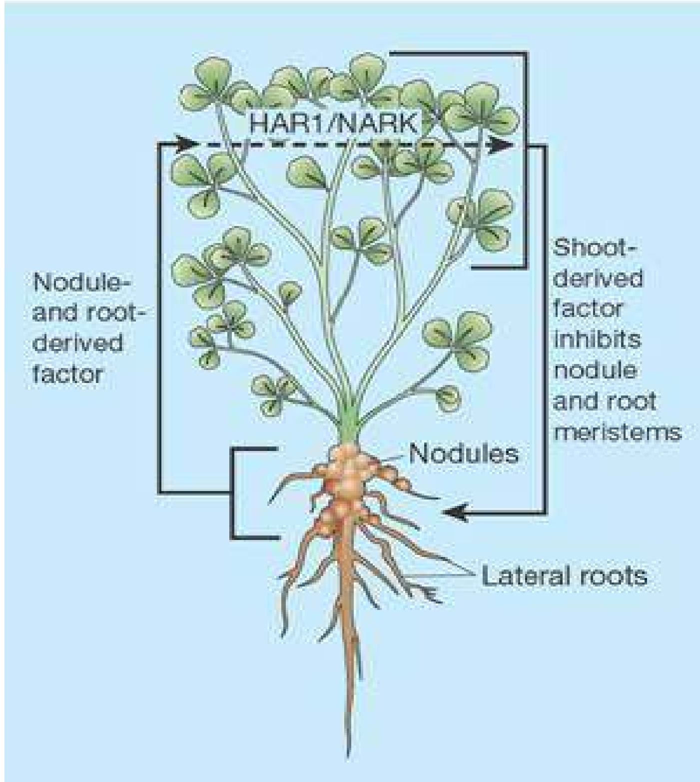

# Insect Colonies

- COLLECTIVE DECISION MAKING IN INSECT COLONIES
- How do social insects decide that the majority of a colony prefers a nest over another? Strategies: accumulation of chemicals until a threshold is exceeded….

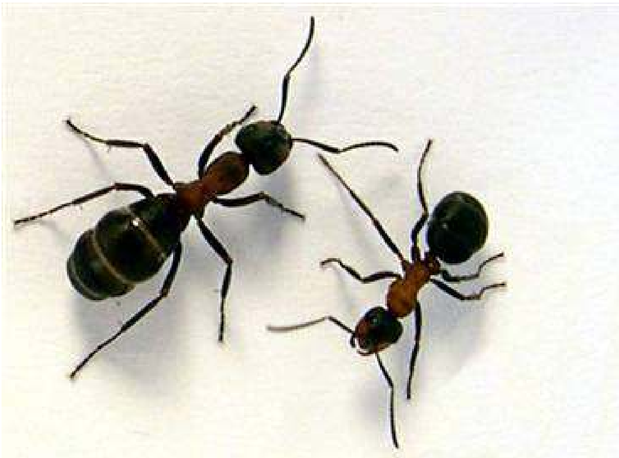

- **Stigmergy** is a method of indirect communication in a self-organizing emergent system where its individual parts communicate with one another by modifying their local environment.

- Stigmergy was first observed in nature - ants communicate to one another by laying down pheromones along their trails, so where ants go within and around their ant colony is a stigmergic system. Similar phenomena are easily seen in many (all?) eusocial creatures, such as termites, who use pheromones to build their very complex nests by following a simple decentralized rule set.
- Each insect scoops up a 'mudball' or similar material from its environment, invests the ball with pheromones, and deposits it on the ground. Termites are attracted to their nestmates' pheromones and are therefore more likely to drop their own mudballs near their neighbors'. Over time this leads to the construction of pillars, arches, tunnels and chambers.

# 4 - Adaptation and Learning

- Being able to change stimulus-response coupling as the result of experience can become very abstract, when complex internal states are allowed.
- Rather than a simple numeric value (concentration of a chemical), one can use very sophisticated descriptions of past experience. Possibly a system can be composed of various adaptive parts, whose interactions depend on the state of each of them.
- This leads to adaptation and learning, and makes the system even harder to predict. Now it integrates multiple information sources, processes them at the light of its own history, and makes decisions about actions.

  [Example of learning from biology: **crayfish escape behaviour**].

# Summary: on Creating Complex and Purposeful Behaviour

- The key is for the system to achieve its goals, and outperform its competitors
- This is achieved – in some systems – by exploiting information about their environment, and assumptions / knowledge about it
- This information is processed in order to inform behaviour, and decide actions
- There is a clear market opportunity for this kind of system: an evolutionary NICHE

# Again: Teleological Language

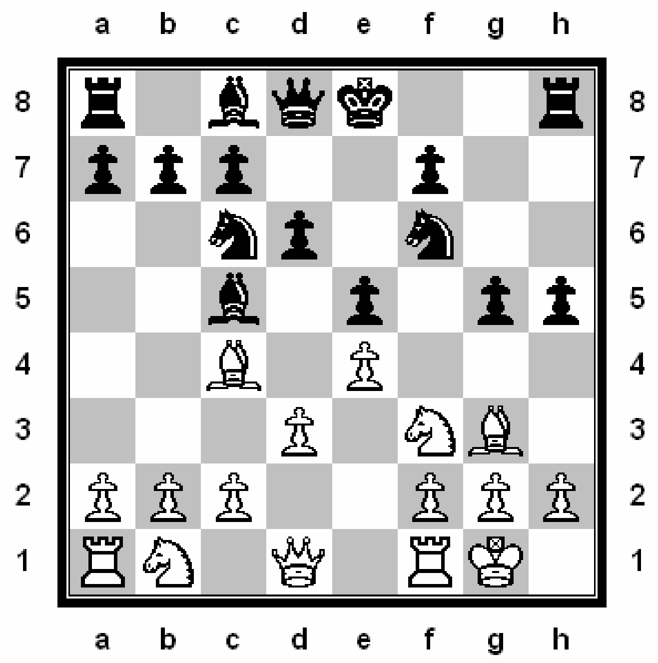


- The interaction of sensory integration, memory, learning, and adaptive scripts, takes place in a highly complex environment.
- The behaviour of the system is very hard to predict, or even describe, in purely causal language.
- It turns out to be much more efficient for us to start using a teleological language: trying to hypothesize a goal, or a purpose, and describe the behaviour of the system in terms of achieving goals, perhaps with constraints.
- “Mind”-language is introduced when purpose is introduced.

It is possible to do a little experiment on these lines, even at the present stage of knowledge. It is not difficult to devise a paper machine which will play a not very bad game of chess. Now get three men as subjects for the experiment A, B, C. A and C are to be rather poor chess players, B is the operator who works the paper machine. (In order that he should be able to work it fairly fast it is advisable that he be both mathematician and chess player). Two rooms are used with some arrangement for communicating moves, and a game is played between C and either A or the paper machine. C may find it quite difficult to tell which he is playing.

(This is a rather idealised form of an experiment I have actually done).

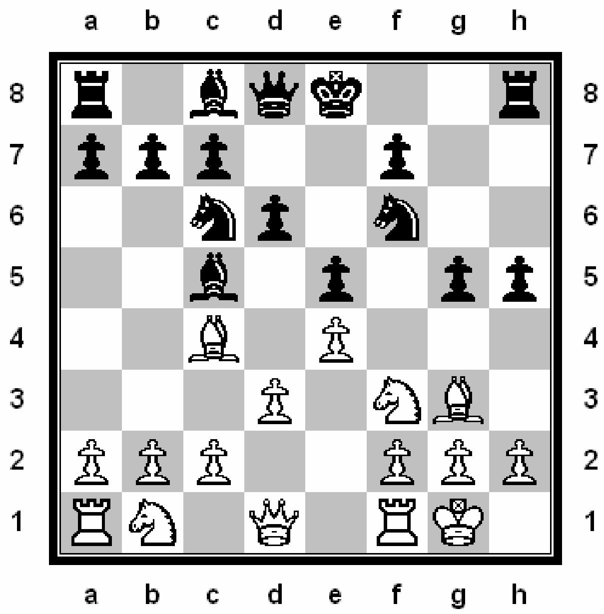


# Repertoire of Actions…

- The notion itself of an action can be very complex.
- Certain systems can perform a complete script, based on a single stimulus (crayfish escape).
- They may choose a script out of a portfolio of several different options (see also irreversible sporulation decisions in yeast).
- **Combine this with triggering scripts based on extensive integration of information and memory…**

# Intelligent Behaviour in Computers?

- Why not?

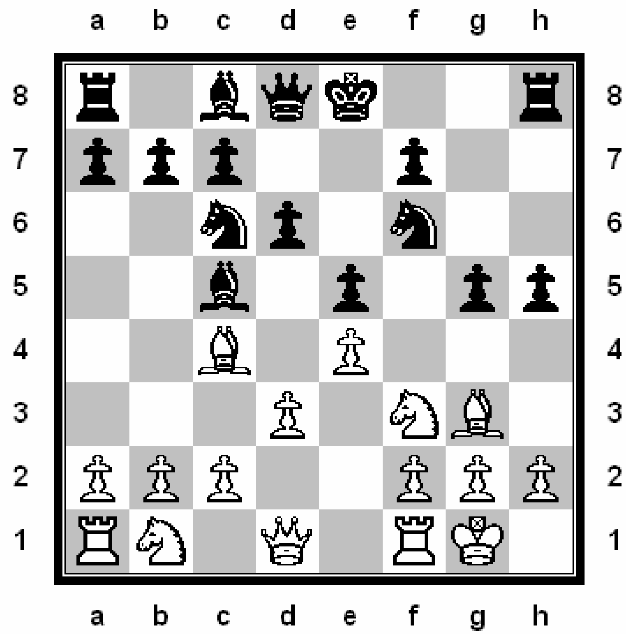

# On Cognitive Machinery - 2008

- What would be today’s counterpart of Turing’s (paper-implemented) chess match?
- Something that would appear to require intelligence, according to his definition.
- Something that is best described in ‘mental’ terms: goals, motives…
- Something that is aware of its environment, learns from experience, has memory, integrates multiple information sources, has a goal and uses assumptions to pursue it ….

- mmmm

# Making Suggestions


It does seem that a recommendation system like Amazon would have passed some form of Turing Test 50 years ago, when competing with a human in giving book suggestions to experts.

# Answering Questions

(and making suggestions, and finding customers, and …)

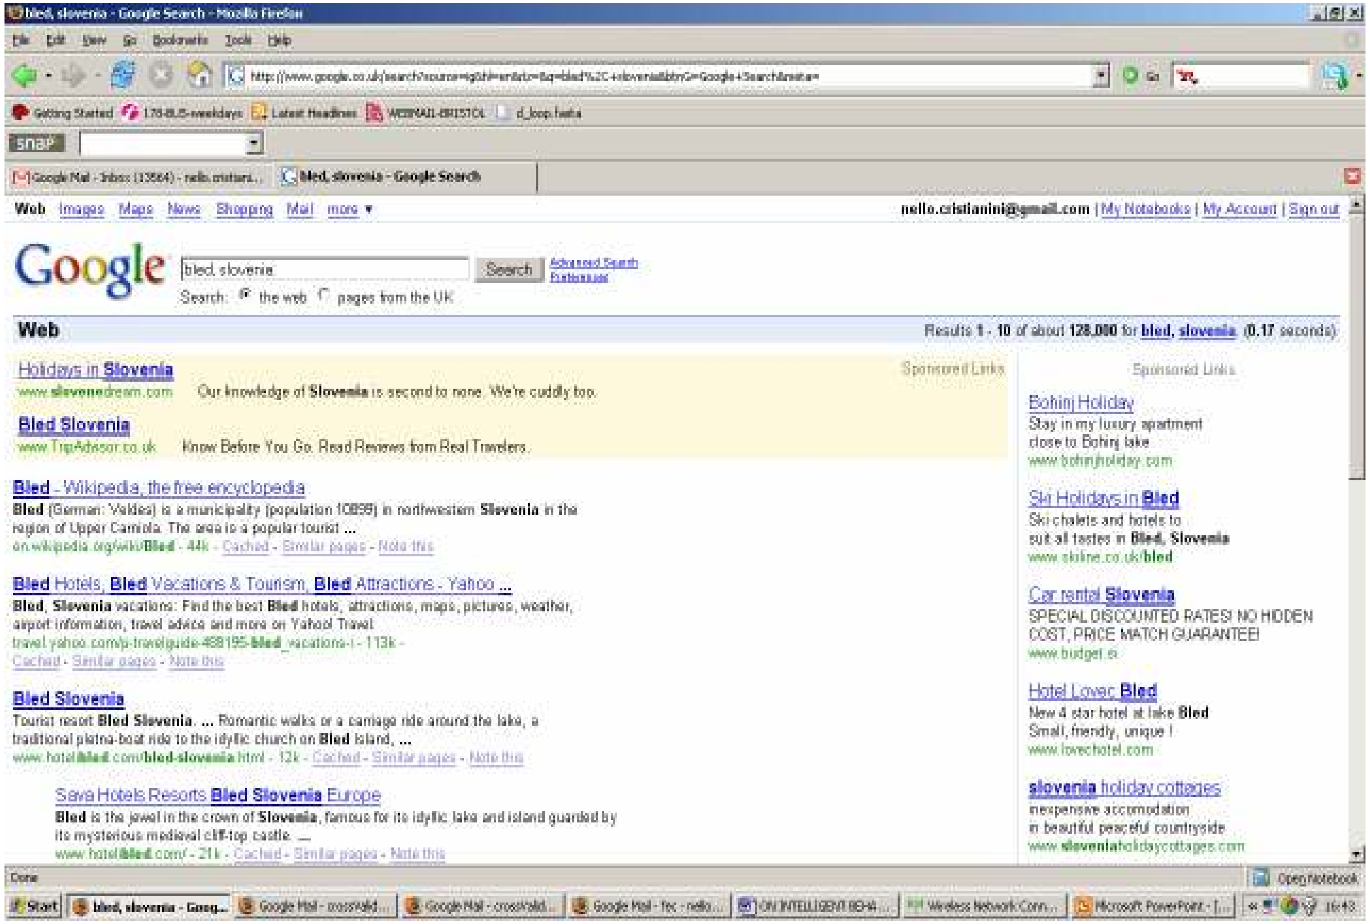


Take Google: it integrates various forms of information; it models the user; it adapts its behaviour over time.

It is closely interconnected with its very complex environment.

Why should we not consider it as a cognitive system?

[at least as much as a plant?]

# Intelligent Behaviour in Video Games

# Cognitive Robots

- Certainly reactive robotics from the 1990s can produce highly successful systems that can be considered as cognitive in many ways.
- Subsumption architecture: pathways connect sensors to actuators
  various intermediate steps
  but no central deliberating stage

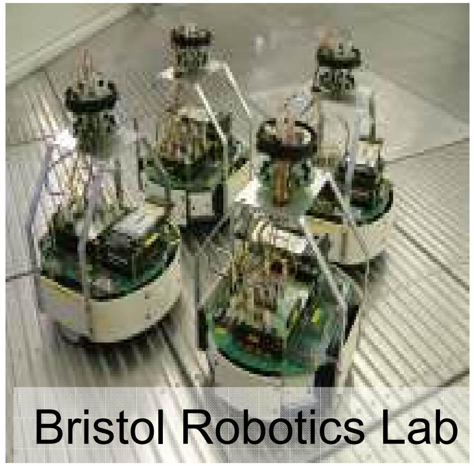

# Reactive Robotics

- Rodney Brooks
  [*intelligence without representation*]

SENSOR → COGNITION → ACTION

vs.

SENSOR → ACTION

[Observer (cognition)]

# Brainless Intelligence

- Distributed systems like ant colonies, cells, plants, reactive-robots…
- Complex decision making is achieved
- Learning is achieved

- Game theory, decisions, dynamical systems, control theory, reinforcement learning, … all in a parallel and asynchronous fashion…
- Can we work on this too?

# Cognitive Pathways

- A common aspect of all these systems is the **modularity of the design**, the presence of **pathways where sensory information is gradually integrated**, processed… and turned into a decision
- Distinction between pre-processing and inference disappears

- Can this be a good target for analysis?
- **Information integration pathways**, signal transduction pathways that create integration, memory, adaptation, generalization ? And complexity?

- The loose integration of various modules, to form cognitive pathways, where information is increasingly processed and abstracted, until a decision is made, seems to be not only a feature of existing systems, but a (necessary?) result of evolutionary processes.

- From cells to organizations, careful combination of robust components in a ecological network seems to produce the kind of behaviour that we are most prepared to call 'cognitive'.

# 5- Planning vs. Reacting

This is a different but key aspect…
→ ***REASONING***


- It also seems that we value the capability of a system to predict the **long term effect** of its own actions.
- It is unclear where the line is drawn here, in biology.

- If we want to add this layer to the definition of cognition, we need to consider systems with some form of symbolic reasoning (or those capable of memorizing the final outcome of past actions).
- A chess player would be able to internally simulate the possible evolution of a match. Is this where we want to draw the line?
- This was clearly the attitude of AI in the 70s and 80s. And **problem solving is obviously a part of the story**..
- A system like **Expedia** would probably have been considered as a cognitive system, suggesting paths, connections, solutions to achieve a stated goal (to fly from A to B and back, under some constraints).

# Cognitive Systems on the Web

- Is the web the obvious playground for this kind of intelligence?
- Cognitive behaviour is defined within a given environment, just like fitness in evolutionary biology.
- The **best possible interfacing with an environment** today can be found not via sensors, but via web. A web system can be aware of an enormous amount of high quality information and use it to determine its behaviour.

- In the future, **sensor networks will play a similar role**, and will require similar solutions.

# New Web as a “Cognitive Ecology”?

- No point in dreaming a general purpose cognitive system: the definition itself requires the choice of a specific set of goals.
- Biology never produced a general purpose organism that is fit in every niche.
- Why should we try?

- The general endeavour to create cognitive systems is well under way. Specific systems are already operating in their own niches, and creating new niches for more abstract systems.

- An ecology of cognitive systems can arise, with systems that process low level information, and provide the output to higher level systems (think of “Google juice”, or of all mash-ups that we use every day in web 2.0).

# AI is on the move…

- In recent years, AI has been made by Google, Amazon, Expedia, …
- Cognitive systems have been studied by molecular (system?) biologists…

- Time to join the party!

"All we need" is to focus on (making / modelling) real systems
and on the following aspects….

# Check List

- Information Acquisition and Processing
- Internal States
- Information Integration
- Adaptation
- Planning

- Combination of multiple components in cognitive pathways to generate behaviour

- Not on list:
  human-like creativity;
  human-level language;
  human-like emotions;
  free will;
  self consciousness;

- On list:
  awareness
  (of environment)
  decision making
  Learning
  Seeking goals

# Summary

- *Behaviour* generation and analysis should be a focus for us
- We should be able to explain natural behaviour, create artificial systems
- Robots, web systems, cell behaviour, should be on our radar
- What can our style of work bring to the *general* field of cognitive systems?

- See thematic programme talk on monday

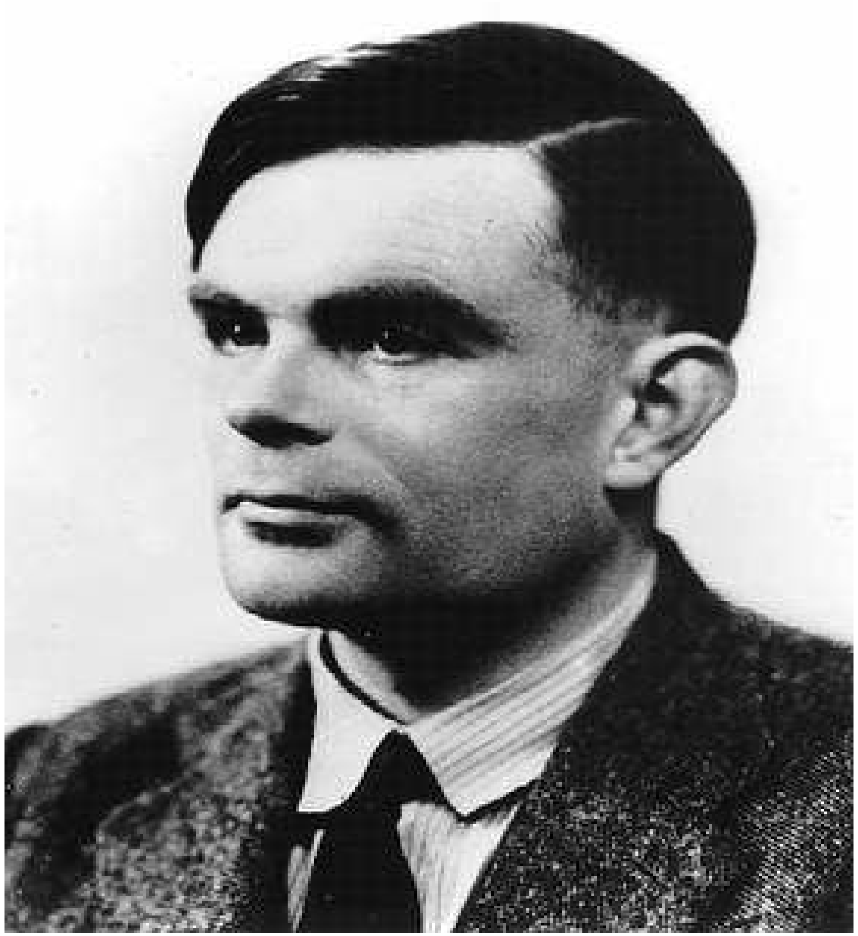

the idea of 'intelligence' is itself emotional rather than mathematical.