Computing:

The Science of Nearly Everything.

Some possibilities for computing education that relate to deep thinking and have not been widely acknowledged.

Why ideas from computing are important
for many other academic disciplines
besides computing, and engineering.

We should be doing much more to stretch young minds using these ideas.

Aaron Sloman, School of Computer Science, UoB. `http://www.cs.bham.ac.uk/~axs`

These slides are available here:
`http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#everything`

What's the content of computing education?

(1) A large collection of skills to be learnt.

   This is how many people think of a computing education, whether it is IT, ITC or CS.

(2) A body of knowledge: as in the CAS BOK (Body of Knowledge)

   Knowledge about computers, computing, programming, programs, networks, applications, theory, ...

(3) Ways of having a lot of fun while learning to program, or learning to use a variety of powerful tools.

IS THAT ALL?

Something else to teach

There are aspects of computing that are important for students interested in other academic disciplines, including:

philosophy,

linguistics,

psychology,

biology,

social science,

mathematics,

physics,

earth sciences,

and many more.

I.e. those who are not going to use programming techniques to build something useful or entertaining, but who could use computing ideas and techniques

to help them understand the world they live in, including aspects of their own minds.

Example: What is the universe made of?

It is clear that a great deal of what we find in the universe is composed of, and involves transformations of,

- matter
- energy,   and
- information.

Example: a table, the energy involved in moving it, the information it provides about how you can move.

It's not just man-made machines that process information.

Not all information processing systems use bit patterns implemented in transistors.

All organisms are information-processors but the information to be processed has changed and so have the means
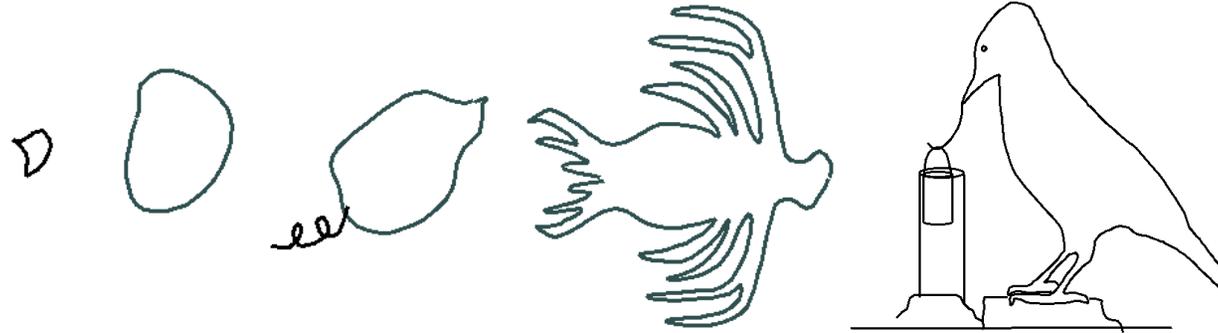
Types of environment with different information-processing requirements

- Chemical soup
- Soup with detectable gradients
- Soup plus some stable structures (places with good stuff, bad stuff, obstacles, supports, shelters)
- Things that have to be manipulated to be eaten (e.g. disassembled)
- Controllable manipulators
- Things that try to eat you
- Food that tries to escape
- Mates with preferences
- Competitors for food and mates
- Collaborators that need, or can supply, information.
- and so on .....

Discussion topic: How do the information-processing requirements change across these cases?

Betty the hook-making crow

Show demo.

This is very far beyond anything current robots can do?

Why?

Compare BigDog

Information and life

All biological organisms use and deploy energy in what they do.

The processes involve selection from options and control of details.

Informed control: both selections among alternatives and online modifications of behaviours are all based on **information**, about: needs, opportunities, constraints, achievements, discrepancies, resources available, structures, processes, opponents, helpers ....

The information can come from various sources, e.g. from external and internal sensors, from things learnt previously in the individual's life, by reasoning, and from the genome.

We need to educate far more people so that they can think about the workings of information-processing systems, on the basis of first-hand practical experience of designing, building, testing, debugging and explaining (initially simple) examples:

This should be a pre-requisite for studying and teaching biology, psychology, neuroscience, philosophy ... and several other subjects. It can also help mathematicians explore mathematical spaces.

Living things are informed control systems

Multiple designs are found in living things.

Not only varied designs for physical structures, from microbes to elephants –
but also a wide variety of ways of processing information at the level of

- cells (e.g. detecting intruders or repairing damage)
- whole organisms (e.g. selecting actions, controlling movements)
- groups of organisms – e.g. foraging insects

Living things are information processors: informed control systems

Organisms are control systems:
They acquire and use information (external and internal),
in order to control how they use matter and energy

> (in order to acquire more matter, energy and information,
> and also reproduce, repair, defend against intruders, dispose of waste products...).

Over many millions of years, evolution produced

- more and more sophisticated information processors,
- driven partly by changes in physical environments,
- partly by results of prior evolutionary history
- in their ancestors and in other organisms (enemies, food, symbiants, ...).

Much of that process is unexplained: people studying evolution have not learnt to think
about possible requirements for information processing systems, and possible designs.

Demos depending on time

Designing "bumpy" agents

vs

designing "thinky" agents.

Not just designing, testing, running, but also analysing, explaining, comparing, criticising, documenting...

Topics for possible discussion within CAS

1. Do we want to broaden the scope of CAS to include:

   teaching about the ways in which computing ideas and programming experience can illuminate other disciplines,
   especially understanding natural intelligence, in humans and other species?

   [Including the nature of mind and consciousness.]

2. How do those goals affect the choice of computing/programming concepts, techniques and principles that are relevant?

3. What are good ways to do that?

   E.g. what sorts of languages and tools help and what sorts of learning/teaching activities? How much freedom should teachers have to choose what suits them and their pupils?

4. Which children should learn about this?

   Contrast
   - Offering specialised versions for learners interested in biology, psychology, economics, linguistics, philosophy, mathematics, etc...
   - Offering a study of computation as part of a general science syllabus.
     (Better to start with this then move to specialised versions.)

5. Is there any scope for that within current syllabus structures, and if not, what can be done about making space?

Using Computation to Illuminate

Instead of always teaching programming because of its practical applications, we can also teach some learners to use computing ideas and programming experiments to illuminate other disciplines.

Especially helping us think about natural intelligence, in humans and other species

[Including re-shaping questions about the nature of mind and consciousness.]

Example topics for learners:

Linguistics	Visual perception	Analogical reasoning
Motivation	Emotions	Planning
Plan execution	Learning	Evolution of consciousness
Curiosity	Play	Developing numerical competence

and many more...

For example:

- Many learners like creating programs that manipulate fragments of natural language, initially in textual form, and initially in very simple ways (as in "toy" chatbots), then gradually adding deeper models of uses of language to talk about something.

- We can show how Darwin's thinking (and thinking of some of his opponents) about evolution of consciousness might have been different if he had known about virtual machines in computers:
  `http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#darwin`

What else? Biology, psychology, economics, linguistics, philosophy, mathematics.

The remaining slides illustrate only a small subset of the possibilities, largely from an AI standpoint.

Remaining Slides

- The remaining slides present examples of topics that illustrate the preceding points

- about teaching young learners how to think about natural as well as artificial information processing systems

- using ideas derived from attempts to build working models as well as problems encountered in attempts to build human-like or animal-like machines.

- The list is merely illustrative: different teachers will prefer to focus on different topics.
    (There is great scope for teacher creativity.

- Some of the examples are fairly sophisticated and would only be suitable for advanced learners.

- Others are very simple and could even be used in primary school.

- This is a personal list, based on my experience of teaching AI in various contexts: a lot more work is required to produce a broad list of examples to help teachers choose, or design their own topics.

- I have used examples based on a programming language (Pop-11) and development environment (Poplog) that I know provides powerful tools for teachers and learners. There is no implication that this combination is the only option: it merely illustrates what is possible.

Programming "bumpy" agents vs "thinky" agents

There are now powerful tools enabling learners to program simulated agents moving around in a simulated world where they sense other things in the world and react accordingly, e.g. by producing simulated physical changes in their behaviour, such as

- shooting
- chasing
- bumping into
- pushing
- following

A quite different sort of task is to develop simulated agents that acquire store, use, share, and later modify information obtained from the environment.

An example is a simulated agent that explores some terrain and builds up a map of where things are which it can later use to make plans of various sorts involving taking, fetching, visiting, etc.

This could be extended to include cooperative or competitive behaviours, e.g. sharing information about where food or dangers or shelter are to be found, possibly entering into trading arrangements.

This will require some thinking not only about what information is acquired and used, but also what sorts of motives the agents can have, and how their motives are generated, how conflicts are detected and resolved (e.g. between low short term gain and high long term gain).

Exercises of this sort could be lead to discussions about how societies evolve, the function of money, the need for law-makers and law enforcers, etc.

Use of language involves information processing

Learners can enjoy getting computers to manipulate fragments of their own language.

- Initially a child can play with several "chatbot" programs, such as this (rather simple)
  eliza : `http://www.cs.bham.ac.uk/research/projects/cogaff/eliza`

- And then later learn how to build such a thing
  - starting with a very simple version,
  - Then later adding more sophistication,
    e.g. use of a simple grammar.

- For this, it is a great help to have a programming language that
  - includes facilities for list processing (including **lists** of strings, instead of only strings)
 and which includes
    - a pattern matcher for recognising, decomposing and re-creating patterns in list structures.

    One example is the Pop-11 pattern matcher:
    `http://www.cs.bham.ac.uk/research/projects/poplog/teach/matches`

The enjoyment and learning that comes from getting a machine to manipulate language does not depend on producing a perfect conversationalist.

Analysing and reducing the imperfections of a flawed program can provide deeper learning than simply producing a bug-free program.

A good teacher will set traps that produce buggy programs:
experience of fixing one's own buggy thoughts and designs is one of the most important forms of learning.

Beyond "surface" manipulation

Here's a sample interaction with the aforementioned Eliza

(Human types in lower case, and machine responds in upper case.):

    ? I would like to teach you to swim.
    WOULD YOU REALLY LIKE TO TEACH ME TO SWIM
    ? Yes you would enjoy it a lot
    WOULD THIS ENVIRONMENT MAKE YOU HAPPY?
    ? We can't swim here
    DID SOMEONE TELL YOU YOU CAN'T SWIM HERE
    ? There isn't any water here
    COULDN'T THERE REALLY ANY WATER HERE ?
    ? Can you see any water here?
    I WISH I COULD SEE ANY WATER HERE
    ? if there were water here we could learn to swim
    HOW CAN IF THERE WERE WATER HERE YOU LEARN TO SWIM ?
    ? Your grammar is pathetic
    PATHETIC ? – NO NOT MY GRAMMAR .
    ? Yes – your grammar
    DIDN'T YOU PREVIOUSLY SAY YOU WOULD LIKE TO TEACH ME TO SWIM

Could a child playing with such a system come up with theories about how it works, and perhaps test those theories, without looking at the code?

Beyond "surface" manipulation

Eliza-like systems merely look at and manipulate patterns of words, whereas natural language processing techniques go beyond that.

including using generative grammars to parse or produce novel sentences.

It is possible to introduce learners to ways of getting a computer to use a grammar and a lexicon

- to analyse simple sentences in accordance with the grammar
- to generate sentences (e.g. randomly) using a grammar and a lexicon
- and later on to relate grammatical structure to semantic content so as to be able to answer questions about some simple world, or obey commands to act in the world

    As Terry Winograd demonstrated around 1971 at MIT:
    A simplified version of his model (SHRDLU) is shown in this online video (GBLOCKS):
    http://www.cs.bham.ac.uk/research/projects/poplog/figs/simagent/#gblocks

    Students can be given the package and invited to modify the grammar, and/or the lexicon, to allow more varied forms of interaction, or to enrich the initial environment. Some snapshots are below.

    Less experienced learners could be given a partially built (possibly buggy) simpler version of this program and invited to extend it (or debug it – looking only at a small subset of the code).
        An example of that approach is the move from this exercise
            http://www.cs.bham.ac.uk/research/projects/poplog/examples/#river
        to this one:
            http://www.cs.bham.ac.uk/research/projects/poplog/teach/riverchat
        (among many other possibilities).

Snapshots of a toy version of SHRDLU running

This video shows a "conversational" simulated robot inspired by Winograd's SHRDLU:

`http://www.cs.bham.ac.uk/research/projects/poplog/figs/simagent/#gblocks`

Here is a snapshot from a run of the program:

```
Move mouse pointer to this window
Type a sentence - 'help' for help - type 'bye' to exit
? help


You can:
  1. give a command to MOVE or PUT a block somewhere.
  2. ask where a block is, ask whether a block is on
         another block, ask what blocks are on a block, or ask
         which block is on another block
Examples of things you can type:

    put a green block on a blue block
    move the little red block onto a big green block
    put a block on the table onto a blue block
    put the block on a block on a block on a red block
    is the big red block on the small green one
    where is the small blue block
    put it on the table
    put the big green block on it
    what is on the big green block
    what is on the block on the big green one
```

```
/-\
```

QUIT

The picture shows a simulated robot hand above a table on which there are six blocks, of two sizes and three colours.

The 'help' command gives instructions, as shown.

Several more examples follow.

The Toy SHRDLU - snapshot 2

The parse tree for "is the small red block on a blue one" is shown both as a textual representation of a list structure (list of lists of lists...), on the left, and as a picture, below.

The Toy SHRDLU - snapshot 3

It goes from the parse tree to a "semantic interpretation", working out which entities in the known world (represented as a simple database of facts) might be referred to in the sentence, then answers the question if it is unambiguous:
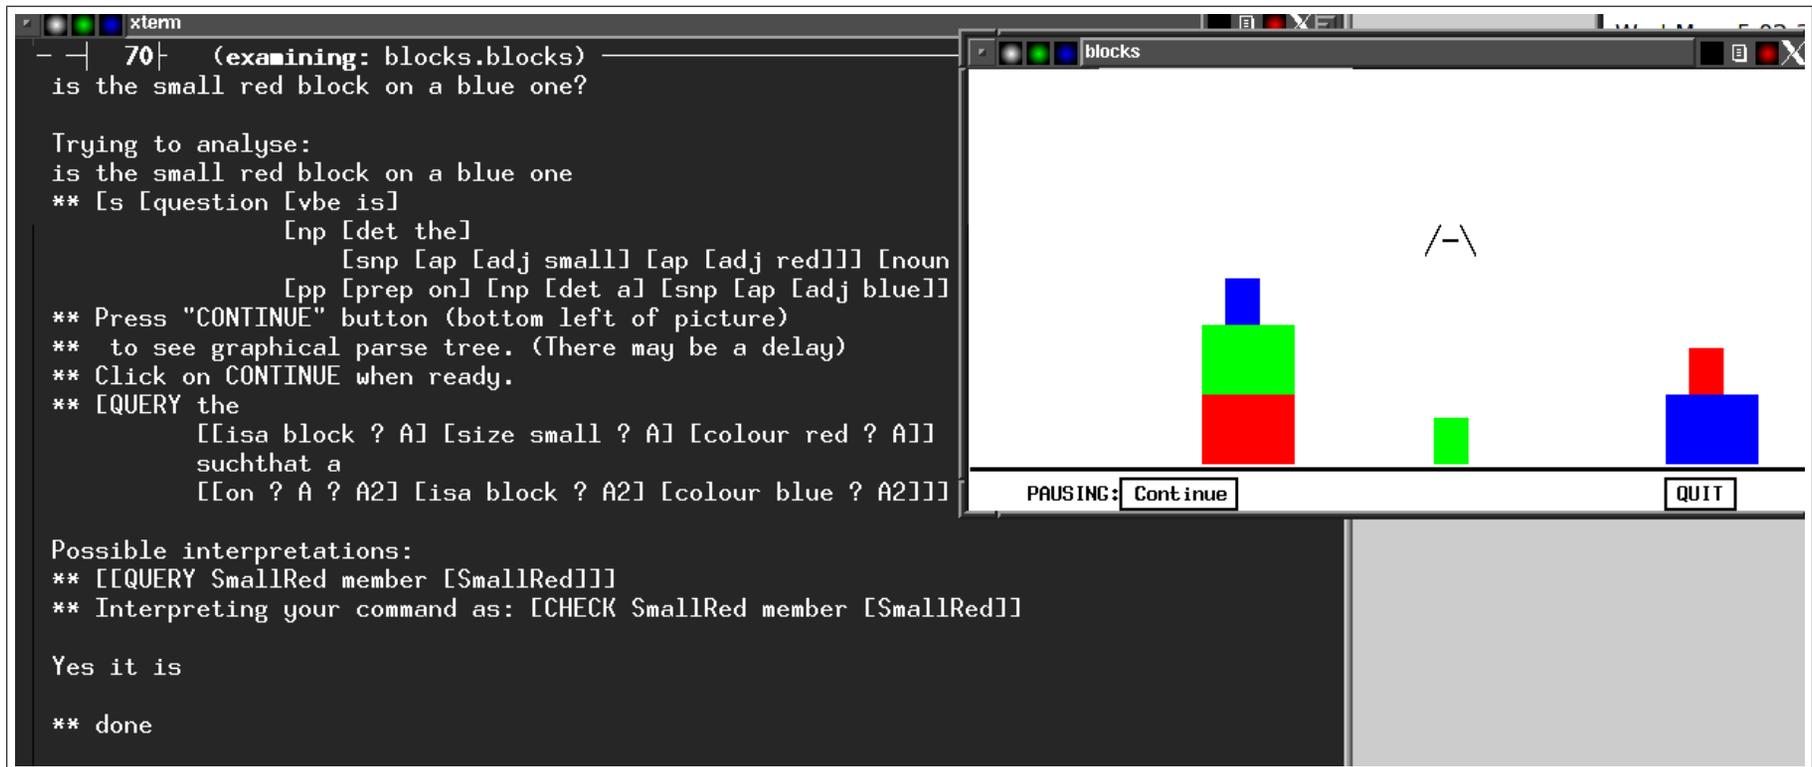
It interprets the question "is the small red block on a blue one" as being about what is in the set of blocks that are on blue blocks, and decides that the block known as SmallRed is a member of the set of blocks that are on blue blocks.

The Toy SHRDLU - snapshot 4

Now showing the result of parsing "Put the big green block on the little red one."

From the parse tree it derives an interpretation of the sentence as a command to achieve a new state of the world: the goal state, then finds a plan to achieve the state, by moving blocks around (see the video).

The result of carrying out the plan to achieve the goal.


```
─ ─┤  94├─  (examining: blocks.blocks) ─────────
? put the big green block on the little red one

Trying to analyse:
put the big green block on the little red one
** [s [vp [v put]
        [np [det the] [snp [ap [adj big] [ap [adj green]]]
        [pp [prep on]
            [np [det the]
                [snp [ap [adj little] [ap [adj red]]] [noun
** Press "CONTINUE" button (bottom left of picture)
**  to see graphical parse tree. (There may be a delay)

        [[isa block ? A] [size large ? A] [colour green ? A]]
        on the
        [[isa block ? A] [size small ? A] [colour red ? A]]]

Possible interpretations:
** [[PUT BigGreen on SmallRed]]
** Interpreting your command as: [PUT BigGreen on SmallRed]

** Plan:
** [[move SmallBlue SmallRed]
    [move SmallBlue SmallGreen]
    [move BigGreen SmallRed]]
** Plan analysed and improved. New Plan:
** [[move SmallBlue SmallGreen] [move BigGreen SmallRed]]
WATCH THE GRAPHIC WINDOW
** [Doing move SmallBlue SmallGreen]
** Moving above SmallBlue
** Going left
** Moving Down
** Moving above SmallGreen
** Going right
** Moving Down
** [SmallBlue now on SmallGreen]
** [Doing move BigGreen SmallRed]
** Moving above BigGreen
** Going left
** Moving Down
** Moving above SmallRed
** Going right
** Moving Down
** [BigGreen now on SmallRed]

** done
```

(Actions of the program are shown in the SimAgent video mentioned above.)

More examples of language processing

If learners can be provided with a library package that can accept a user-provided grammar and lexicon and on that basis construct

- a sentence generator

- a sentence parser

then they can do "programming" by specifying a grammar and a lexicon.

Example exercises could include:

Create a grammar and a lexicon for
- railway station announcements
- insults
- the opening sentence of a child's story book
- simple stories
- haikus

Typically people who try this find that some of the sentences generated by the program were not supposed to be generated, e.g.

Platform five will depart from the London train in five minutes.

Trying to work out how to avoid the unwanted sentences involves thinking about adding more details to the grammar, e.g. more syntactic categories.

Using the student's grammar to parse the unwanted output can be revealing.

All that can help students understand better how their own language works.

Playing with, modifying, assembling program fragments

It would be unreasonable to expect beginners to be able, even if working in teams, to build up anything like the SHRDLU program starting from just a programming language,

even a powerful AI language, like Lisp, Prolog, or Pop-11.

- An alternative is to demonstrate the whole system, but to give them carefully designed tasks involving parts of the system.

- So a library for doing things with simple grammars can be provided, and students can play with it and try to develop an interesting grammar for communication with a robot – testing the grammar in terms of what it can generate randomly and how it parses test sentences.

- A few, who progress rapidly, may ask to see the parsing and generating code and try to extend the libraries, or produce their own versions, while others stick with the simpler tasks.

- Likewise a planning program could be provided for students to play with.

- They may learn by testing it with various problems and see where it fails.

- A few may try to extend, or modify, the program, produce a new version from scratch.

The opportunities for creative teachers are vast – if they have powerful tools and exemplars to start with. (More are becoming available.)

Visible and invisible behaviours

Many of the toolkits involving simulated agents, designed for use by beginners support development of visible behaviours, which can include moving around on a 2-D surface, rotating in 3-D, changing colour, changing direction or speed of motion, changing size or shape, etc.

However, many organisms do information-processing that is not externally visible.

In the case of humans that can involve processes like

- passive perception

- thinking about something

- making plans

- solving problems

- choosing between alternatives

- learning (e.g. analysing what was done wrong in a recent action).

I'll give some examples of apparently passive perception, where there is no external behaviour but complex internal processes occur.

Examples of "invisible" information processing

Being **conscious of something** involves processing information

A familiar example: The Necker Cube

Although you are looking at a 2-D configuration of black pixels on a white background, you see lines not pixels and you probably see a 3-D structure not just a 2-D pattern. How???

> If you stare at it long enough the experience you have will change:
>
> The 2-D lines remain where they were, but you will see them as edges of a 3-D wire-frame cube.
>
> The image is ambiguous and your percept of it can flip between two different views, with edges changing their relative distance from you, and their 3-D orientation, even though nothing changes in the image.
>
> What forms of processing can give a machine the ability to look at a 2-D pattern of marks and interpret the marks as representing a 3-D structure?
>
> What processes in the machine could make the interpretation flip while the input remains unchanged?

If you look at this web site

`http://www.math.ubc.ca/~morey/java/rotator/rot.html`
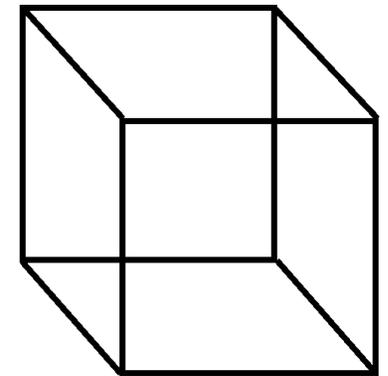
You will see several more things that can be interpreted as moving 3-D structures, even though only a flat screen is in front of you.

How can static and moving 2-D patterns be interpreted
as 3-D structures and processes in the environment?

Some things are easier to model than visual experiences....

A challenge for the class-room

Organise a discussion on what would need to go on inside a machine to enable it to take in a 2-D visual structure like the picture of the necker cube and interpret it as a 3-D wire-frame object that could not be accommodated on a flat sheet of paper, as the picture can.

- What would have to change inside the machine when the picture "flips" ?

- Are there other kinds of ambiguous pictures? Would the same mechanisms work for all of them?

- Teacher: construct some simple programming exercises involving taking in pictures made of lines and describing them, including pictures with grouping ambiguities, or figure-ground ambiguities.

- A very simple case: a rectangular array of dots could be seen as a collection of rows of dots, or a collection of columns of dots.

  Compare verbal ambiguities: She hit the man with a fish?
  Which man did she hit? What did she hit the man with?

- How should such a program describe an ambiguous picture?

- Some simple examples of getting programs to see structures in pictures made of dots are here:
  `http://www.cs.bham.ac.uk/research/projects/poplog/packages/current/teaching/teach/seepics`
  (I need to produce examples of those programs running.)

More examples of "Seeing as"

The Duck-Rabbit (Compare the Necker cube mentioned previously.)

There's a 2-D pattern as before, but you probably see this one as an animal with particular physiological features.

If you stare long enough it should flip between two animals with different features, even though there is no change in either the 2-D or the 3-D geometry of what you see.
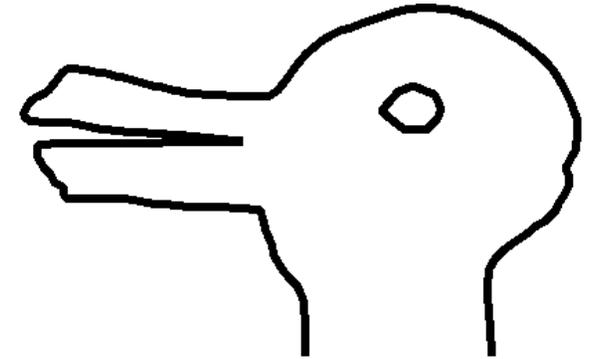
E.g. a part of the image can be seen as ears, or a bill: but both alternatives occupy the same 3-D space.

An even more abstract change probably occurs: you see the animal facing to your left or facing to your right depending whether you see it as a duck or as a rabbit.

How can a machine be made to see something as facing left, or as facing right?

The machine would have to have an ontology that includes things having states that involve facing, moving, seeing, etc.

It might relate the direction in which something is facing to what it can see, and which way it is likely to move.

This requires having the ability to think about what **could** happen even if it is not **actually** happening now.

How would such possibilities, and constraints on possibilities be represented in the machine (or in an animal's mind).

How can we find out how animals (including humans) represent such information?
Will trying to build a working model help? (Often? Rarely? Never?)

More visual information processing – for planning actions

Human vision includes many capabilities that are proving very difficult to replicate in current machines.

This may be because of our limited understanding of what the problems are.

(a)   (b)   (c)

You can look at the scene in (a) and think about how your fingers might have to move to produce the arrangement (b) (using your right hand, or your left hand, or both).

You can also think about a sequence of finger movements that would be able to perform the reverse rearrangement, from (b) to (a).

Current robots cannot do that: What would have to go on inside them to make it possible?

It's not too hard to get a machine to derive image (c) from image (a): could that be a first step???

Notice that examples like this show that seeing is FAR more than attaching labels to portions of images – which is all that some current "vision" systems can do.

Seeing should not be confused with recognising: you can see things you don't recognise.

A simple program to explore geometric process patterns

This example is a program that can be used by learners to explore ways of generating a variety of geometrical structures by varying parameters in a very simple drawing program (originally inspired by the Logo Turtle).

The key idea is this program, expressed in Pop-11, but equally simply expressible in several other languages:

```
define polyspi(side, inc, ang, num);
    ;;; Draw a polygonal spiral. Initial arm length is side.
    ;;; inc (positive or negative) is added to side at each turn.
    ;;; The angle turned (to left) is ang (in degrees).
    ;;; The total number of sides is num.
    repeat num times
        draw(side);
        turn(ang);
        side + inc ->side;
    endrepeat;
enddefine;
```

It should be obvious that this command will draw a square of side 400:
```
  polyspi(400, 0, 90, 4);
```
What will these do?
```
  polyspi(400, 0, 60, 4);
  polyspi(400, 0, 120, 4);
  polyspi(0, 5, 60, 100);
```
A program is available to encourage exploration of the space of possibilities.

Using the "polypanel" program
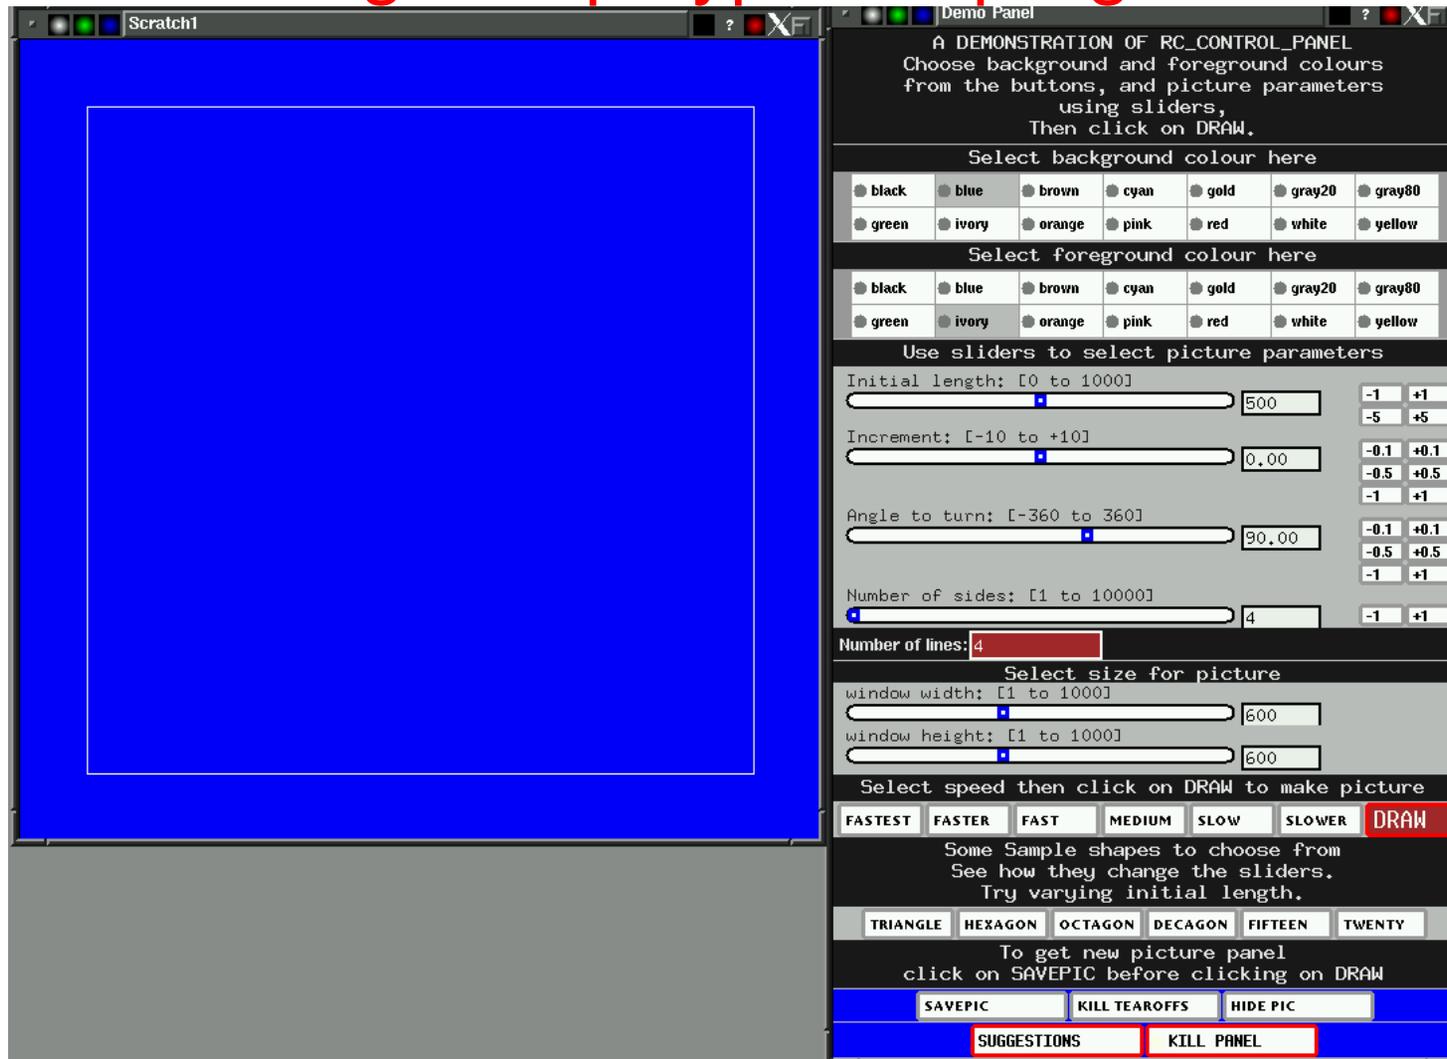
The user can choose background and foreground colours, an initial length, an increment (positive or negative), an angle to turn, a total number of sides, a drawing speed, or a pre-configured example shape, or ask for suggestions. Choices here: blue, ivory, 500, 0, 90, 4

Classroom use of the polypanel program 1

Can you tell by looking at the figure on the left, how the numbers selected in the panel menus were changed? (Depending on your display you may or may not see the picture as intended.)
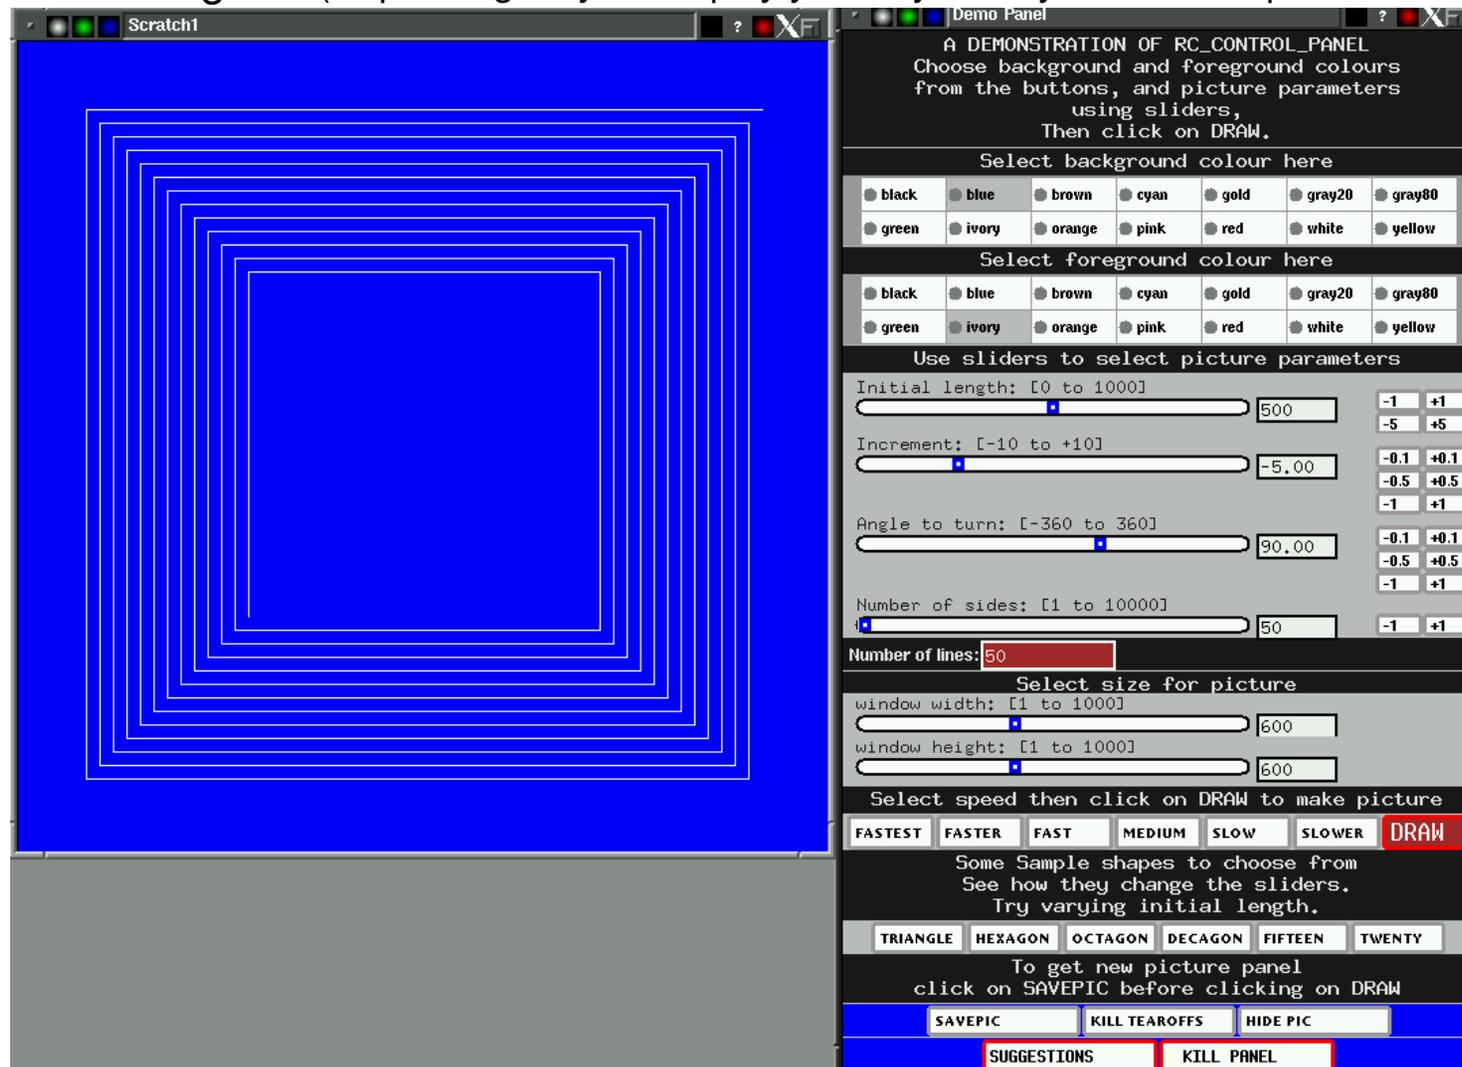
(If the lines in the picture vary in brightness that is a scaling effect.)

Classroom use of the polypanel program 2

A good thing to do in the classroom is to discuss relationships between changes in the selected numbers and the pictures. E.g. what changes from the previous input would produce the picture on the right?
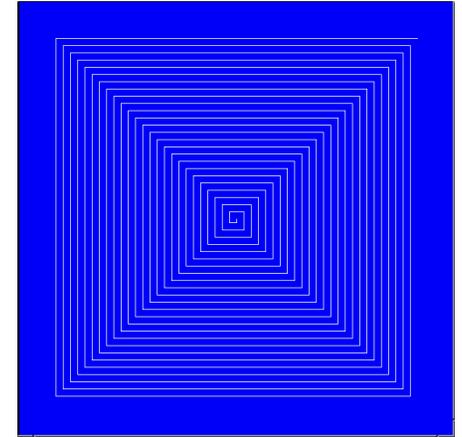
A teacher can invite suggestions as to how to produce effects like these:
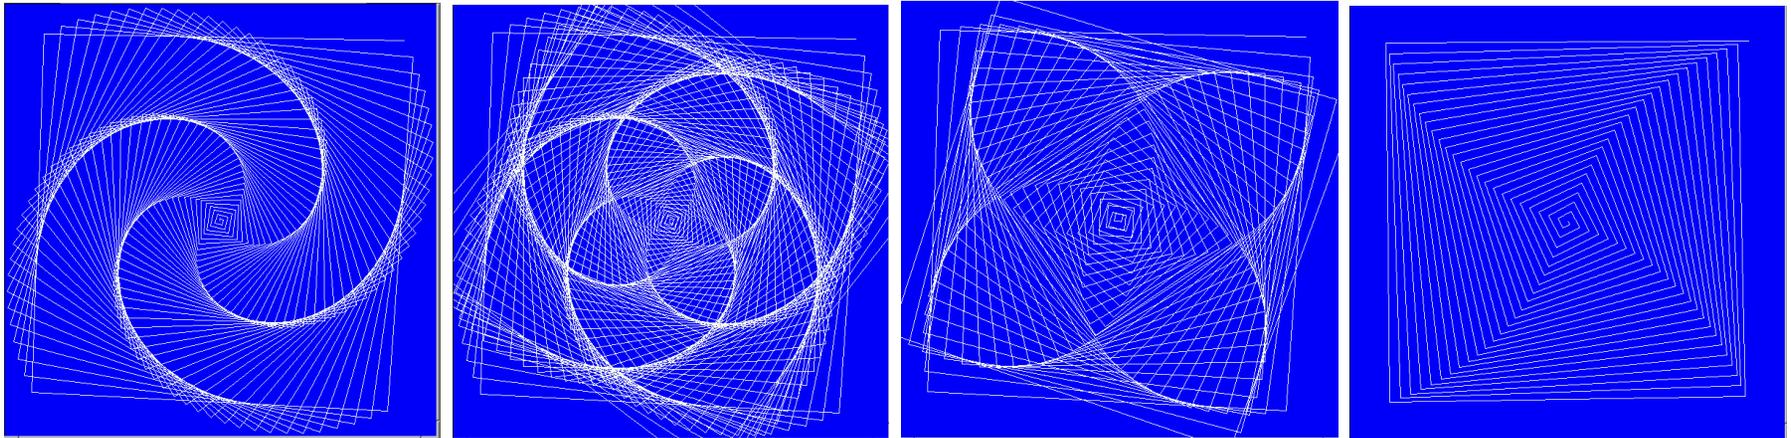
In particular, what happens when the lines drawn have negative lengths?

How could you produce a smooth looking spiral curve?

The point is not just making pretty pictures, but understanding the relations between structure, process, and generic representations of classes of processes.

How we can get machines to do things versus how machines can get themselves to do things.

The previous examples might, for some students, provide a starting point for asking questions like:

- If a robot or animal encounters different shapes in the environment how should it think about those shapes?

- One way of thinking about them is thinking about the processes that can create them
  (the philosophy behind the LOGO turtle, and our previous examples)?

- What other ways are there of thinking about shapes?
  e.g. using logical descriptions of parts and their relationships?
  or using pictures, sketches, diagrams, models?

- How do the different ways of thinking about spatial structure relate to different purposes for which we think about spatial structures?

- For what sorts of different purposes do we need to think about them?

- How could such processes be replicated in future machines?

Computation, Intelligence and the Study of Mind

A major area of impact of the development of computers has been to accelerate new thinking about the nature of minds – in humans, other animals and machines.

Although the label "Artificial Intelligence" (AI) suggests a focus on applications using machines, in fact most of the leading thinkers in AI, including John McCarthy, who invented them name, understand AI as much broader than the science and art of making smart machines.

See his online discussion:

`http://www-formal.stanford.edu/jmc/whatisai/node1.html`
WHAT IS ARTIFICIAL INTELLIGENCE?

There are also relevant papers and discussions on the web site of another of the founders of AI, Marvin Minsky,

`http://web.media.mit.edu/~minsky/`

including the nearly final draft of his book, *The Emotion Machine*.

These slides attempt to show that problems Darwin encountered, regarding evolution of consciousness might have addressed by using new ideas from computing that were not available until recently:

`http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#darwin`

Varieties of learning

There is a lot of work being done on getting machines to learn many different sorts of things, e.g. learning to recognise objects, learning to use language, learning to work, learning to grasp things, and many more.

Most of that work is concerned with achieving successful learning, as part of the process of making a machine more useful.

However, as scientists rather than engineers, we may also want to model some of the incorrect, confused forms of learning that humans go through.

The next few slides illustrate an example, though there are many more examples in the developmental psychology literature.

Probably teachers have observed many examples of mis-learning that could be amenable to an analysis in terms of information-processing bugs.

Compare the BUGGY program developed by John Seely-Brown and colleagues, many years ago.

Reasoning quirks during development:
reasoning about potentially colliding cars

Here is an example of a child who is nearly, but not quite, ready to reorganise and systematise some of the things he has learnt.

A child can learn to use a mixture of simulation and symbolic/verbal reasoning to draw conclusions and to justify conclusions.

That learning takes time and the process can go through buggy stages as illustrated by this example.

A racing car is on the left, and a van on the right.

The two vehicles start moving towards each other at the same time.

The racing car on the left moves much faster than the truck on the right.

Whereabouts will they meet – more to the left or to the right, or in the middle?

Where do you think a five year old will say they meet?

One five year old's spatial reasoning

The two vehicles start moving towards each other at the same time.

The racing car on the left moves much faster than the truck on the right.

Whereabouts will they meet – more to the left or to the right, or in the middle?

One five year old was asked this question with two toy vehicles placed on opposite ends of a window-sill, facing each other.
He answered by pointing to a location near 'b' !!!

One five year old's spatial reasoning

a b c d e f g h i j

The two vehicles start moving towards each other at the same time.

The racing car on the left moves much faster than the truck on the right.

Whereabouts will they meet – more to the left or to the right, or in the middle?

One five year old answered by pointing to a location near 'b'

Me: Why?

Child: It's going faster so it will get there sooner.

What produces this answer:
- Missing knowledge?
- Inappropriate representations?
- Missing information-processing procedures?
- An inadequate information-processing architecture?
- Inappropriate control mechanisms in the architecture?
- A buggy neural mechanism for simulating objects moving at different speeds?

Would designing simulations for parts of this scenario, or the whole thing, help a child understand how to think about relative motion of this kind?

Partly integrated competences in a five year old

Perhaps the child had not yet learnt that there can be constraints on ways in which items of information can be combined, like these:

- If two objects in a race start moving at the same time to the same target, the faster one will get there first
- Arriving earlier implies travelling for a shorter time.
- The shorter the time of travel, the shorter the distance traversed.
- So the racing car will travel a shorter distance!

The first premiss is a buggy generalisation: it does not allow for different kinds of 'race'.

The others have conditions of applicability that need to be checked.

Perhaps the child had not taken in the fact that the problem required the racing car and the truck to be travelling for the same length of time, or had not remembered to make use of that information.

Perhaps the child had the information (as could be tested by probing), but lacked the information-processing architecture required to make full and consistent use of it, and to control the derivation of consequences properly.

This is a simple example of a type of confusion I have often seen in students trying to think about processes: many features of a process are understood, but the learner has not developed a way of thinking about them simultaneously so as to ensure consistency. (Inadequate short-term memory mechanisms?)

(Compare being unable to see the impossibility of a Penrose triangle – or more complex Escher picture.)

Discussing hypotheses before producing working models

Producing a full working model of the processes that produce the confusion about the car race would be very complex and difficult.

However, even for learners who have not yet got the experience and knowledge required to propose and implement a working model, it could nevertheless be useful to think about what the problems would be, on the basis of having done some much simpler programming tasks.

There's another type of development that can occur in learning about numbers or learning about shapes and shape interactions.

- The child learns some new concepts, or new types of behaviour
- The child constructs and plays with instances of the concepts or the behaviour type.
- As a result the child makes an empirical discovery – e.g. counting a row of objects left to right and counting the same set right to left produces the same object.
- Later the child acquires a deeper understanding and realises that those are not merely empirical discoveries: the truth of the generalisations can be derived by reasoning.

There are examples related to counting, measuring, ordering, and rearranging entities in,
J. Sauvy and S. Suavy, *The Child's Discovery of Space: From hopscotch to mazes – an introduction to intuitive topology*, Penguin Education, Harmondsworth, 1974,

A group with experience of simple AI programming could start discussing what forms of representation, what algorithms, what information-processing architecture could explain the observed phenomena, led by the teacher.

How could biological evolution produce human or animal visual abilities?

How could the ability to process information as humans and other animals do have been produced by evolution – apparently starting with molecules in a chemical soup?

Some forms of human-like information processing are easier to model than others:

Some of the earliest AI successes were computer programs that learnt to play board games (e.g. checkers/draughts), worked out solutions to puzzles, made plans, and solved various problems.

There have also been various AI expert systems – some of which perform better than humans, e.g. at medical diagnosis or other tasks that depend on making good uses of much evidence.

But it has proved very much harder to give AI systems the ability to see and understanding processes in which moving 3-D objects interact, e.g. hands and cups, or an elephant's trunk and the things it picks up.

Could mechanisms known to AI researchers, computer scientists, software engineers, etc., explain what human consciousness is?

There has been partial progress insofar as we have learnt over several decades to create processes in machines that have many of the features of mental processes, e.g.

- they cannot be observed and measured using physical devices

- they have semantic content (meaning) insofar as the use information and create and send information

- their interactions can produce physical effects (e.g. flight control systems, or spelling checkers that report things on a screen)

- in some cases the machine may be able to observe what it is doing internally and keep records of what it has done and improve its skills and knowledge as a result.

We have now learnt how to create running virtual machines in computers, with those competences, i.e. virtual machines that are not mere mathematical abstractions but can actually do things, e.g. landing an aeroplane in the dark.

Those virtual machines are not physical entities that you can see or measure by opening up a computer and looking inside or using voltmeters and other tools of physics.

For a taster of how to use ideas about the development of virtual machines as a product of biological evolution, see these slides:

```
http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#darwin
http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#inside
```

Which children?

I have given a number of examples (and there are many more that could have been given.)

Which learners should be exposed to these problems and given the chance to start learning to build working models of individuals that perceive, think, deliberate and reason in deciding what to do?

Clearly the set of potential beneficiaries includes more than just the learners who wish to go into computing research or development, e.g. as software engineers.

The beneficiaries could include learners who wish to go off and do a wide variety of different further courses of study, e.g. in philosophy, biology, psychology, sociology, economics, etc.

Contrast

- Offering specialised versions for learners interested in biology, psychology, economics, linguistics, philosophy, mathematics.

- Offering a study of computation for modelling complex systems, including animal minds, as part of a general science syllabus.

How to do the teaching?

How do those goals affect the choice of computing/programming concepts, techniques and principles that are relevant?

What are good ways to do that?

E.g. what sorts of languages and tools help?

What sorts of learning/teaching activities?

To be expanded.

Some examples are here
http://www.cs.bham.ac.uk/research/projects/poplog/examples/

Practicalities – in the current climate?

Is there any scope for that within current syllabus structures, and if not, what can be done about making space?

Towards a liberal education for the 21st Century

Some things should form part of a general educational system because they are worth knowing about in their own right for many different reasons, including all the main academic disciplines:

- physical sciences (physics, chemistry, astronomy, earth sciences, ...)
- biological sciences, including neuroscience, psychology,
- mathematics
- linguistics
- philosophy
- music and other art forms
- history, geography, literature, social sciences
- application fields: engineering, medicine, ...

A subset of worthwhile competences are specially important because they provide the tools for learning, thinking, reasoning and communicating about some, or all, of the others – it is normally assumed that these include the "three-Rs": Reading, wRiting, aRithmetic.

(Actually arithmetic is a small subset of mathematics, the real third competence.)

We need Five Rs, i.e. including programming broadly construed as learning to think about, analyse, explain and use tools for thinking about:

structures, processes and how structures can produce and be produced by processes.

Computers have been able to support such education for nearly 40 years – but alas the opportunity has been mostly ignored, and it is now very difficult to use it: but perhaps not impossible. **That's the assumption of this presentation.**