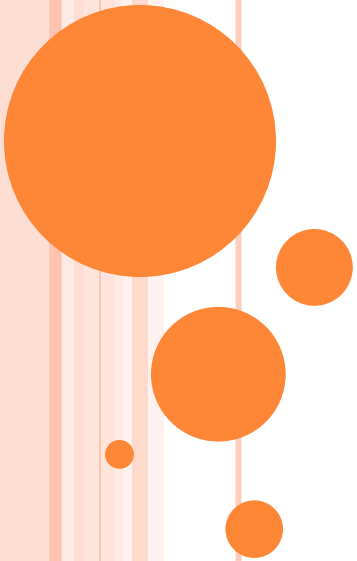

Modelling and Designing Cognitive Behaviour

Nello Cristianini

Intelligent Systems Laboratory

University of Bristol

A journalist asking for a quote …

- "How far are computers from overtaking human intelligence?"

"This is not what we do…"

- But: what do we do, really?
- [I have been running a workshop for past 3 years in Bristol to try and frame these issues differently]
- "Modelling Cognitive Behaviour" series.

Making Choices

- There are things in the world that make choices, in order to pursue their goals.
- These choices are driven by information they gather, as well as by their goals.
- Some of these things can improve with time, or adapt to changing environments.
- Some may even make plans, to achieve complex goals in adversarial situations.

- These are what I call today "**cognitive agents**".

- Some Examples of
  Adaptive Goal Driven Behaviour

A Simple Niche

- A very simple cell in a oceanic vent might be able to rely on a constant supply of energy, much like a parasite cell does.
- This would not need much information processing capability, but it would be unable to cope with change, or to exploit new environments, or variable environments.
- Competition (and evolution) would happen in the direction of energetic efficiency, etc.

Cognitive Niche

- A cell capable of responding to different situations with the appropriate actions, can colonise niches that are simply not accessible to others.
- Consider one that can exploit just 2 simple environment.
- This would require the capability to sense the state of the environment and choose the appropriate responses (e.g., fermentation or respiration? Sporulation or normal life cycle? Fructose or lactose?)
- Competition (and evolution) can then happen in the direction of better decision making.
- This is what I call the cognitive niche.

The Value of Information

- Parasitic versions of bacteria have smaller genomes that their free-living counterparts
- Entire metabolic pathways are lost, those producing resources that are either obtained from the host, or not needed in that stable environment. (see various versions of Chlamydia)

The Lac Operon

- E. coli can produce specific enzymes to digest different kinds of sugar… (decision making)
- Lactose is digested by ß-galactosidase protein
- Intracellular *regulatory protein* called the *lactose repressor* to hinder production of ß-galactosidase (in the absence of lactose it is not produces).
- In the presence of lactose, a lactose metabolite called allolactose, binds to the repressor, causing a change in its shape → unable to bind and repress → ß-galactosidase is produced.
- Internal representation of a state of the environment: encoded in the shape of repressor proteins !

Cognitive Niche

- There are various levels of environmental complexity, and increasingly sophisticated sensing / information-processing may enable cells (or higher organisms) to colonise them.
- There is always room at the top...

- A completely random environment would not be exploitable. But this is a matter of degree: for too simple a cell, a complex environment is the same as a random one.

Cognitive Niche

- Complex organisms have been evolving for billions of years.
- Some have evolved eyes, brains, nerves, in order to process information faster and better (in animals, probably driven by predation).
- But complex decision making is found also in cells, plants, and other 'brain-less' agents.

- Certainly intelligence did not arise with humans, or with language, or with logical reasoning, or with consciousness.

Evolution of Information Processing

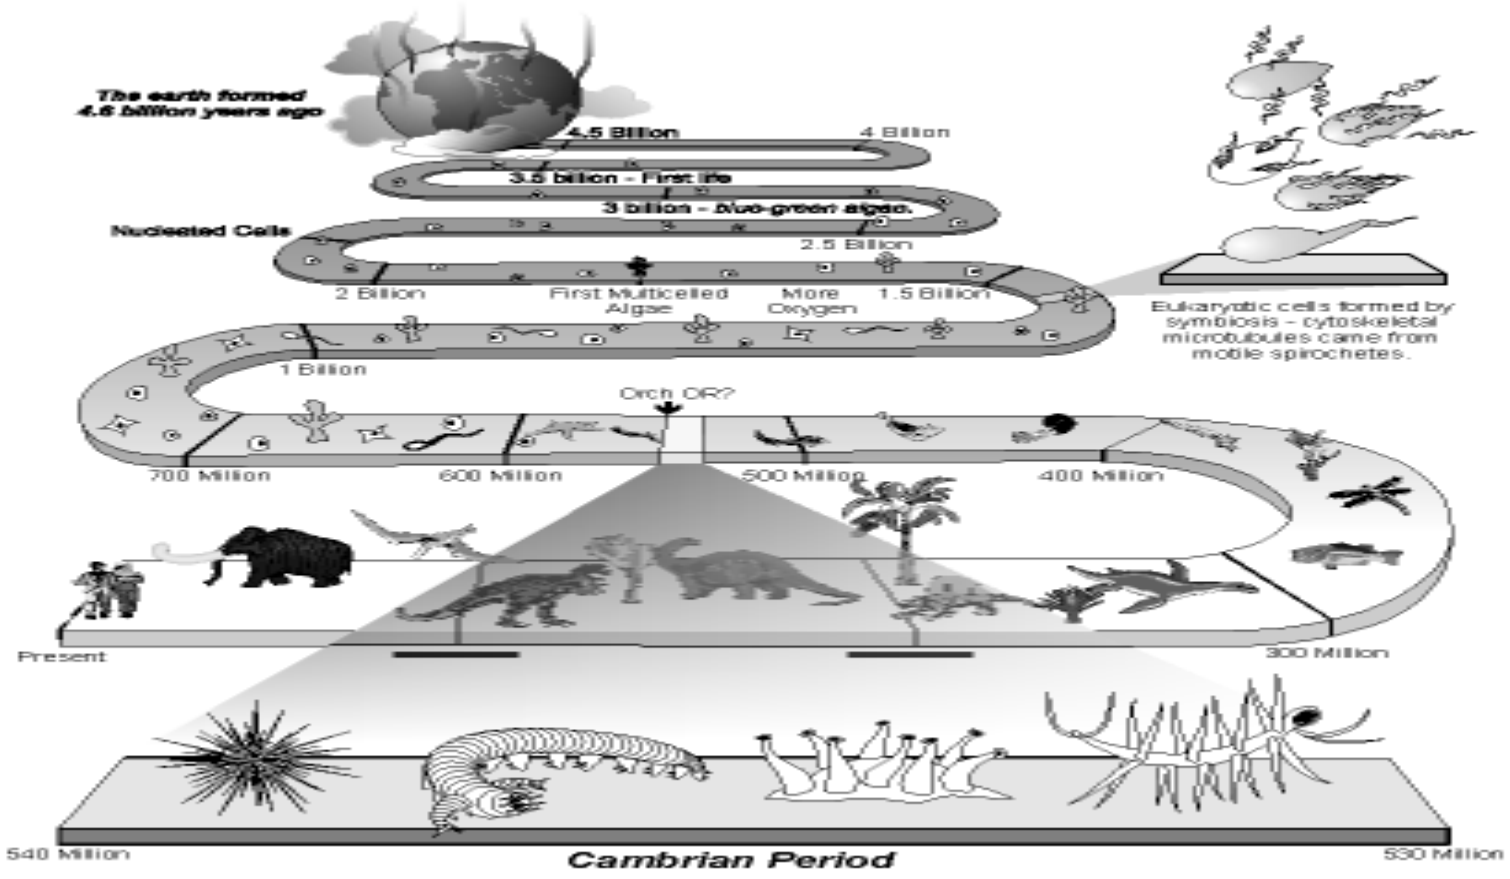


BBC


BBC

Cognitive Behaviour

- It can be very misleading to define intelligence in terms of human behaviour (and hence typically using language or logic as a criterion).

- This is something that emerged in just one species, very recently in evolution, and – while interesting – does not constitute the only instance or type of intelligence.

A working definition

- In order to avoid philosophical minefields, we will rely on a practical working definition for this talk.

- We are interested in the behaviour of an agent pursuing a goal in an environment that cannot be fully controlled by it. It should be able to make choices, respond to the state of the environment, learn the best actions, possibly make plans. It should be robust to change.

- Autonomous goal-seeking adaptive robust ....

Machine Intelligence

- Many artificial systems meet this definition.

Machine Intelligence

- The best examples are seen on the web:
  - Google
  - Amazon
  - ...
- Also some robotic systems exist with these properties.

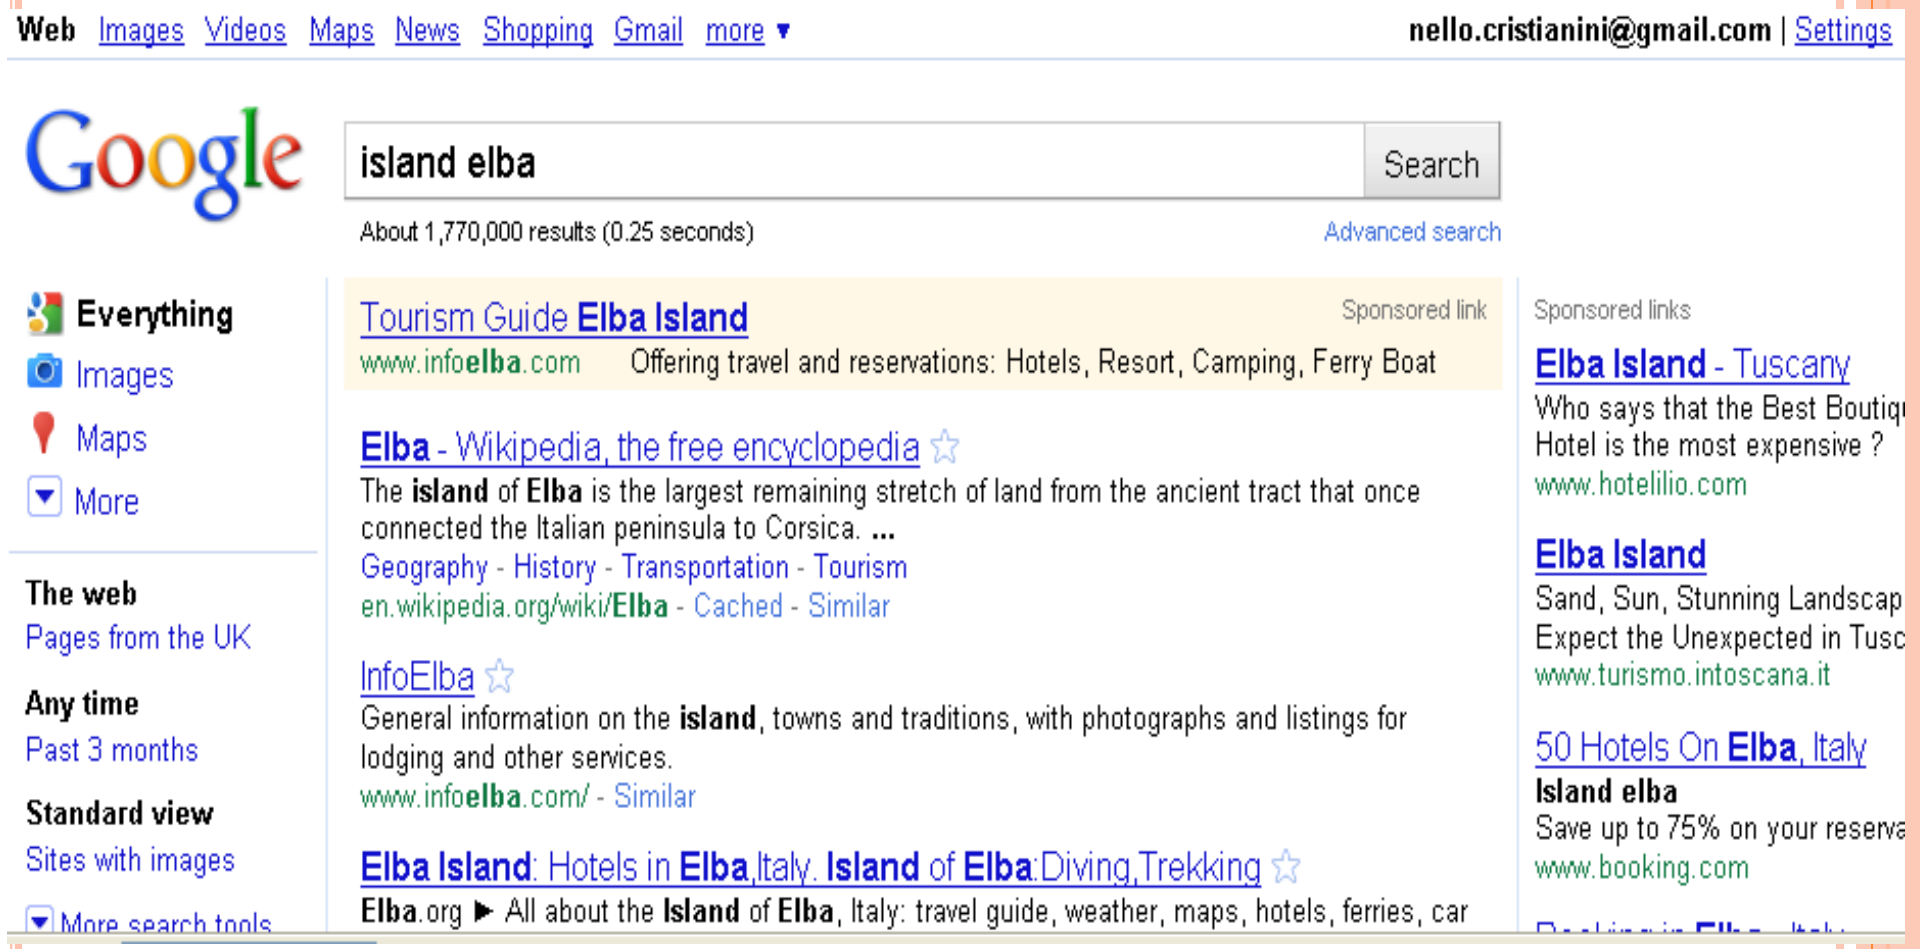
Web Images Videos Maps News Shopping Gmail more ▾
nello.cristianini@gmail.com | Settings
Google
island elba
Search
About 1,770,000 results (0.25 seconds)
Advanced search
Everything
Images
Maps
More
The web
Pages from the UK
Any time
Past 3 months
Standard view
Sites with images
More search tools
Tourism Guide Elba Island
Sponsored link
www.infoelba.com Offering travel and reservations: Hotels, Resort, Camping, Ferry Boat
Elba - Wikipedia, the free encyclopedia
The island of Elba is the largest remaining stretch of land from the ancient tract that once connected the Italian peninsula to Corsica. ...
Geography - History - Transportation - Tourism
en.wikipedia.org/wiki/Elba - Cached - Similar
InfoElba
General information on the island, towns and traditions, with photographs and listings for lodging and other services.
www.infoelba.com/ - Similar
Elba Island: Hotels in Elba,Italy. Island of Elba:Diving,Trekking
Elba.org ▶ All about the Island of Elba, Italy: travel guide, weather, maps, hotels, ferries, car
Sponsored links
Elba Island - Tuscany
Who says that the Best Boutiq
Hotel is the most expensive ?
www.hotelilio.com
Elba Island
Sand, Sun, Stunning Landscap
Expect the Unexpected in Tusc
www.turismo.intoscana.it
50 Hotels On Elba, Italy
Island elba
Save up to 75% on your reserva
www.booking.com

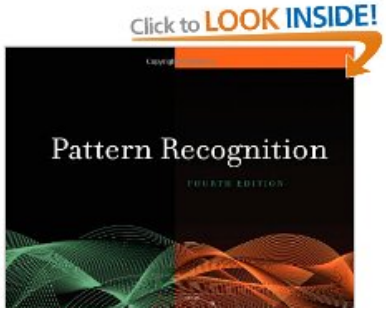

Pattern Recognition [Hardcover]

Sergios Theodoridis Dr. (Author), Konstantinos Koutroumbas (Author)

(1 customer review)

RRP: £60.99

Price: **£45.12** & this item **Delivered FREE in the UK** with Super Saver Delivery. See details and conditions

You Save: £15.87 (26%)

In stock.

Dispatched from and sold by **Amazon.co.uk**. Gift-wrap available.

Frequently Bought Together

Customers buy this book with Pattern Classification, Second Edition: 1 (A Wiley-Interscience publication) by Richard O. Duda

Price For Both: £112.21

Add both to Basket

Show availability and delivery details

Customers Who Bought This Item Also Bought

Pattern Classification, Second Edition: 1 (A W... by Richard O. Duda

(1)

£67.09

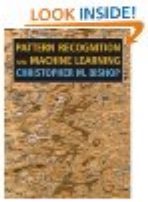

Pattern Recognition and Machine Learning... by Christopher M. Bishop

(7)

£47.65

Introduction to Pattern Recognition: A... by Sergios Theodoridis Dr.

£18.24

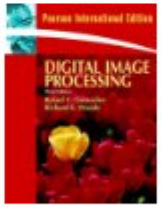

Digital Image Processing by Rafael C. Gonzalez

(3)

£43.04

Modern Multivariate Statistical Techniques: Re... by Alan J. Izenman

(1)

£51.29

Machine Intelligence

- To the journalist on the phone we should say that we are already surrounded by intelligent machines. They exist and we use them every day.

- Where do we stand today, in the quest for machine intelligence? How we got here? Where are we going next?

Machine Intelligence – 1947 AD

- We can appreciate how far we have come if we take a look from the position of Alan M. Turing who first proposed machines could display intelligent behaviour….

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

Machine Intelligence – 1947 AD

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 inteligent and another would not; the second man would have found out the rules of its behaviour.

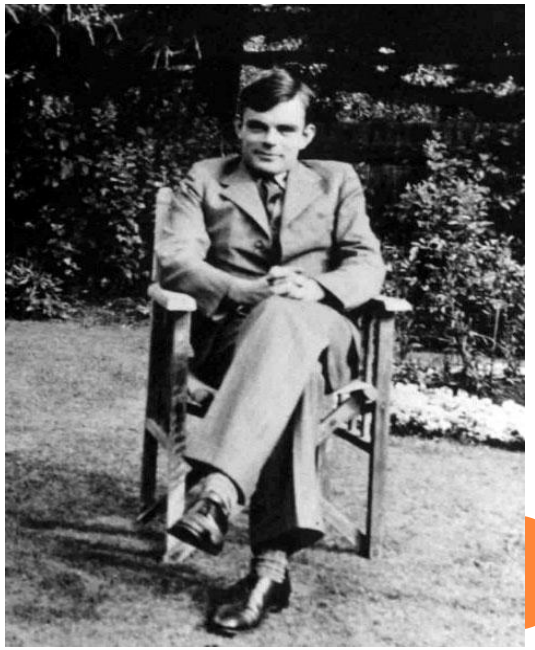

A first, not human-centric,
version of his famous test
(originally done on chess algorithms
– without language)

A bit like hypothesis testing:
can I explain the behaviour
in simple mechanical terms?

Machine Intelligence – 1958 AD

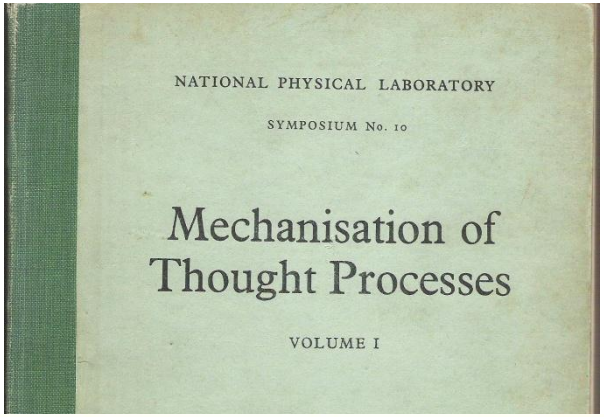

This Symposium was held to bring together scientists studying artificial thinking, character and pattern recognition, learning, mechanical language translation, biology, automatic programming, industrial planning and clerical mechanization. It was felt that a common theme in all these fields was "The Mechanization of Thought Processes" and that an interchange of ideas between these specialists would be very valuable.

Machine Intelligence – 2010 AD

- Today: we trust machines
  to make choices for us in terms of
  - Filtering our mail
  - Recommending products, books, news
  - Planning Journeys
  - …
  - (shooting down aircraft?)
  - …

Machine Intelligence 2010 AD

- From the more "practical" side, machines can perform *human* activities (not that this is a *definition* of intelligence, but sure it is useful)

- Machines today can do
  - Machine translation
  - Speech recognition
  - Information retrieval
  - Handwriting recognition
  - Gene finding
  - Face detection
  - …
- Surely this would have satisfied the criteria of 1947 or 1958

The State of the Art

- A lot of this progress has happened very recently
- A lot has happened after researchers abandoned the attempt to replicate human mental processes

- This is an interesting story

Building Intelligent Machines

- For a long time the dominant perspective has been that *logical reasoning* should be the defining feature of intelligence
- Symbolic reasoning based on logical approaches, performed by a search in a space of possible solutions (solutions as mathematical proofs) …
- Chess match / theorem proving were the paradigms

- This was stated openly by Newell and Simon and eventually led to ‘expert systems”…

REPORT ON A GENERAL PROBLEM-SOLVING PROGRAM

A. Newell
J. C. Shaw
H. A. Simon*

P-1584

30 December 1958

Revised 9 February 1959

SUMMARY

This paper reports on a computer program, called GPS-I for General Problem Solving Program I. Construction and investigation of this program is part of a research effort by the authors to understand the information processes that underlie human intellectual, adaptive, and creative abilities. The approach is synthetic — to construct computer programs that can solve problems requiring intelligence and adaptation, and to discover which varieties of these programs can be matched to data on human problem solving.

GPS-I grew out of an earlier program, the Logic Theorist, which discovers proofs to theorems in the sentential calculus. GPS-I is an attempt to fit the recorded behavior of college students trying to discover proofs. The purpose of this

Building Intelligent Machines

- That perspective was not applied just to chess and mathematical proofs…
- It was assumed that every task can be reduced to this: even navigation, robot-arm control, translation, retrieval …

- Expert-systems needed explicit listing of all relevant facts and assumptions (followed by search in solution space), in order for the reasoning to be carried out …

Building Intelligent Machines

- It turned out that machines designed in this way need listing enormous amounts of background information.
- The only logical consequence was to start an effort to list all facts to that expert systems could become viable (cyc).
- This approach had decades of time to work, and it did not.
- It is all about asking the right question, using the right metaphor… The idea of replicating human logical reasoning in this way led the entire field on this path for decades…

Machine Intelligence, AD 1973

- Debate at the Royal Institution, Sir James Lighthill

Machine Intelligence, AD 2008

- Peter Norvig, Google Director of Research

A Paradigm Shift

- Maybe life is not a chess match, and intelligence is not about proving theorems, after all (they could just have asked any biologist).

- Starting from the late 1980s many big problems were conquered, s soon as this approach was abandoned.

- What carried the day?

Statistical Artificial Intelligence

- Gathering vast amounts of data
- Using the data as examples of the desired behaviour
- Creating learning algorithms that can reproduce the desired behaviour by discovering statistical relations
  - → recognize handwriting
  - → translate documents
  - → transcribe speech
- Machine learning takes centre stage in AI

An Example: The Noisy Channel

TASK: reconstruct messages received through a noisy channel (an inverse problem).

The problem of the noisy channel is "ill posed"
in the sense that it is underspecified,

Just like inferring hidden causes from their effects...

An Example: Viterbi Algorithm

- Originally developed to reconstruct messages received through a noisy channel (an inverse problem).
- Under simple statistical assumptions, the algorithm efficiently finds the most likely message sent (optimisation + statistics).

- The problem of the noisy channel is “ill posed” in the sense that it is underspecified, unless statistical assumptions are made.
- Just like inferring causes from their effects…

A Statistical Approach to Machine Translation

Peter F. Brown, John Cocke, Stephen A. Della Pietra, Vincent J. Della Pietra, Fredrick Jelinek, John D. Lafferty, Robert L. Mercer, and Paul S. Roossin

IBM
Thomas J. Watson Research Center
Yorktown Heights, NY

4768–4778 Nucleic Acids Research, 1994, Vol. 22, No. 22

A hidden Markov model that finds genes in *E.coli* DNA

Anders Krogh, I.Saira Mian[1] and David Haussler[2]*
Nordita, Blegdamsvej 17, DK-2100 Copenhagen, Denmark, [1]Sinsheimer Laboratories, University of California, Santa Cruz, CA 95064 and [2]Computer and Information Sciences, University of California, Santa Cruz, CA 95064, USA

Received June 21, 1994; Revised and Accepted Sept[illegible]er 28, 1994

Error Bounds for Convolutional Codes and an Asymptotically Optimum Decoding Algorithm

ANDREW J. VITERBI, SENIOR MEMBER, IEEE

A Tutorial on Hidden Markov Models and Selected Applications in Speech Recognition

LAWRENCE R. RABINER, FELLOW, IEEE

PROCEEDINGS OF THE IEEE, VOL. 77, NO. 2, FEBRUARY 1989

Belief Networks and PGMs (junction tree)

Impact of Noisy Channel Model

- Humble Viterbi Algorithm more important that theorem proving by resolution,....

- Change in perspective:
  INTELLIGENCE: from search in a space, to solving an inverse problem
- The paradigm:
  from the chess match to the noisy channel

Machine Intelligence, 2010 AD

- A series of simple statistical hacks, together with optimisation technology, and vast amounts of data, delivered much of the intelligent behaviour we use today

- DDAI = data-driven artificial intelligence

Machine Intelligence, 2010 AD

- **Optimisation and Statistics are the language** behind algorithms such as Viterbi's
- Similar ideas are also behind methods like
  - Support Vector Machines (used today for handwriting recognition, text categorization, etc),
  - Probabilistic Graphical Models…
  - Association Mining (used for recommendation systems)
- Because of the centrality of Machine Learning in DDAI, these tools/concepts have become a central part of modern AI
- NOTE: series of theoretical convrgences in 1990s created a unified framework for SLT

Machine Intelligence, 2008 AD

Peter Norvig, Google

Machine Intelligence, 2010 AD

- We have come a long way as a field
- Simple DDAI methods (Statistics + Optimisation + Data) have delivered
  - automatic spelling correction;
  - handwriting recognition;
  - face recognition;
  - speech recognition
  - machine translation
  - Etc
- BUT…

BUT

- Will statistical AI be sufficient to take us “all the way”?
  - At which point are DDAI tools no longer adequate?
  - What do we miss if we cast all as statistical pattern analysis tasks, or statistical inverse problems ?
  - Planning? Reasoning? Explaining? Dynamics? Control? Understanding?
- This depends on the criterion of success, as some things clearly cannot be done by statistical means

BUT

- Has this been a case of "success by redefinition"?
- No: the AI tasks that we have recently conquered were indeed on the wish list of early AI researchers
- But there are many other tasks that we have not conquered (e.g, holding a full conversation with a human, or understanding complex meaning in text, ...)
- Has the emphasis simply shifted towards those tasks that we can handle with statistical AI?

A Change in Metaphor

- The leading metaphor for intelligent task has drifted from “chess match” to “noisy channel”.
- This is not the first big cultural shift in the field.

- Let us keep things in perspective: intelligence may not be all not about search is in a game tree, but it is also probably also not just about solving inverse problems.

Some other metaphors… ☺

We tend to use whatever technology we have at hand, to describe how the mind works…

- Mind as hydraulic system (greeks, romans)
- Mind as clockwork (enlightenment)
- Neurons and telegraph messages (Helmoltz)
- Telephone Switch Network (Hebb)
- Mind as a Computer (Newell and Simon)
- …
- Mind as: echology (Bateson), society (Minsky), market, organisation….

Intelligent Agents seen as…

- Gamblers
  (economics, theory of utility, …)
- Problem solvers
  (classical AI: decompose a complex task into smaller sub-tasks, to be solved or further decomposed…)
- Dynamical systems
- Machines/ automata
- Animals
- A bundle of reflexes, reactions, …

Each different perspective comes with different emphasis on what is important to model, what cognitive tasks are more indicative of intelligence, etc

Some more metaphors … ☺

- Mind / Brain as a web?
- Few years ago on New Scientist: the mind / brain is a Bayesian inference machine (driven by free-energy minimisation)?
- Let us keep in mind the difference between reality and our models, please.

Some more metaphors

- As we succeed in using statistical AI methods, mostly based on “noisy channel” metaphor, let us resist the temptation to draw too strong conclusions
- The “chess match”, the “dynamical system”, the “gambling game” are all valid metaphors to describe some aspects of intelligent behaviour

Cognitive Tasks

- Planning / scheduling / problem solving
- Recognition of patterns
- Solving inverse / ill-posed problems
- Controlling actuators
- …

- It seems that each task can be important for survival, and it is not obvious that they should all be corollaries of a single higher capability (general intelligence).

Swiss Army Knife

- One image that I like: mind as a Swiss Knife
  (made of separate modules, each meant for a task)

- Modules may include:
  - Face detection
  - Language understanding
  - Motor control
  - …

- Modules [definition-encapsulated, etc]

Cognitive Architectures

- Whether biologically plausible or not, modularity is a great design principle for a cognitive agent.
- We have become extremely good at solving many of the key tasks, and a key question for researchers now is:
  HOW DO WE PUT THEM TOGETHER?

- Hopefully we will come up with a system-level understanding of cognitive behaviour, one that can help us analyse natural systems and design artificial ones.

Conclusions

- For the journalist: AI is not trying to overtake humans, yet it is changing the world every day…
- Statistical AI has delivered many of the recent successes in this field.
- There is probably more to intelligent behaviour than what we can achieve by leveraging data, but data-driven AI can still do a lot

- We need to focus more on integrating together the various pieces that we have created in the past few years, to achieve a more systemic understanding.