Diverse Consequences of Algorithmic Probability

Eray Özkural
Bilkent University Computer Engineering Dept.
Ankara, Turkey

Solomonoff 85th Memorial Conference

Axiomatization of AI: basic axioms

- These three (four) axioms are **absolutely** necessary
- **AI0**: AI possesses a universal computer U
  - e.g.: probabilistic universal comp., quantum comp.
- **AI1**: AI can learn any model expressed in U
- **AI2**: AI uses probabilistic prediction (Bayes' theorem)
- **AI3**: AI applies a principle of induction
- Physics envy? Perhaps so!

Incremental learning

- Without any memory, machine learning is infeasible
  - Solution: incremental machine learning
- For two solutions, every common bit of information can double the speed of learning the second at best
- **AI4**: AI must use incremental machine learning
- There are practical incremental ML systems
  - Teramachine: faster than human for simple seqs
  - OOPS: solved towers of hanoi problem

Incremental learning (cont.)

- Adaptive Levin Search (because LS is infeasible)
- OOPS:
  - machine model: FORTH
  - simple probability model: instruction probability
  - dynamic modification of probabilities (bump instr.)
  - results not very convincing for transfer learning

Incremental learning (cont.)

- Teramachine (2011):
  - General purpose U: full Scheme R5RS + library
  - SCFG based probability model
  - Real memory with efficient update algorithms
  - Parallel search algorithm
  - Promising experimental results at 0.5 teraflop/sec
    - *first* results with actually *long* training sequence
    - for deterministic problems: terminating proggies

Cognitive architecture

- A simple Levin Search algorithm may be insufficient
  - Brain is equipped with a lot of innate information
- **AI5**: AI must be arranged so that it can improve itself\
  - this must hold in *practice*, not just on paper
- Promising cognitive architectures:
  - Alpha: functional, higher-order, general, probabilistic
  - Gödel Machine: agent, self-reflective, logical

Philosophical consequences

- Functional synthesis of three big ideas:
  - Carnap's inductive inference
  - Turing's universal computer
  - Chomsky's generative grammars
- Same method used for both math and physics:
  - There is no epistemological difference
  - Math is empirical: thought experiment=computation

Philosophy of AIT

- Mathematics is computational and self-reflective
  - Mathematicians develop useful programs and a precise language to communicate them
    - Mathematical theories themselves are inductive:
      - They improve prediction of physical theory!
      - Multiple theories are useful: ZFC / category theory
    - Computability of prob. density function:
      - a kind of constructivism: information finitism

Philosophy of AIT (cont.)

- Conjecture: Ω_U is the golden standard of math!
  - every meaningful math. fact is reducible to a computational fact, which is present in Ω_U
- Every U corresponds to a nomologically possible world
- Every halting problem corresponds to a thought expr.
- Thus, we are just making general statements about computation, nothing mysterious is going on
- Chaitin's big idea: is evolution computing Omega?

AIT and epistemology

- Non-reductionism: there are no strict bridge laws
  - **certainly false** from an AIT viewpoint
    - much better understanding: algorithmic irreducibility
    - AIT is compatible with evolution
- Is there an objective U?
  - our proposal (2007): universe itself
  - possibilities: RUCA, universal quantum computer

AIT and epistemology (cont.)

- Formalization of knowledge:
  - analytical philosophy: justified true belief
    - our proposal (2007):
      - mutual information between world and mind
    - partial formalization: *justified* is not covered
    - conjecture: is it *predictive* information instead?
    - verification in induction seems sufficient

The vindication of Positivism

- ALP realizes two main tenets of logical positivism:
  - a finite cognitive procedure for inductive inference
  - a unified language for science (AI's private language)
- Analytic-synthetic distinction?
  - not strictly required: ALP can invent & revise logic
  - however, computation has some analytic character

AIT and evolution

- Life is reducible to molecular biology:
  - Molecular biology can’t contain infinite information
- Chemistry is reducible to quantum physics
- Lloyd 2002: every local quantum system can be simulated by a universal quantum computer
  - An AI can use this universal QC for AI0
- Non-reductionism is intelligent design:
  - Flagella of bacteria are too complex to have evolved!

Infinity Point and beyond!

- Infinity Point (Solomonoff 1985):
  - AI improves efficiency of computing at fixed energy
  - More efficient computing accelerates AI
  - Positive feedback: infinite acceleration in finite time!
  - In practice, infinity cannot be reached
- AI milestone G: an AI several times smarter than the entire CS community, 2040's?

A weaker condition

- Suppose we cannot write the right program:
  - Brain simulation is a bio-information based AI
    - Energy efficiency of computing:
      - Today’s GPU: 6 gigaflops/watt
      - Human brain: 5 teraflops/watt
      - 2023 GPU’s, 2026 multicore: 5 teraflops/watt
    - Cheap intellectual work drives brain simulation

Intellectual property and AI

- IP laws assume a fair playing ground:
  - That changes with AI: monopolies wreck economy
- AI4 implies that:
  - Independent discoveries are normal
  - Knowledge is more important than computing speed
- If IP laws are too rigid and costly:
  - Decreases global information sharing
  - Culture of secrecy decreases global intelligence

Intellectual Property and AI

- When technology starts to advance rapidly:
  - IP latency becomes a severe bottleneck
    - Scientists *already* suffer from this bottleneck
- Solutions:
  - Shorten IP protection
  - Make it much more difficult to get IP protection
  - Ignore IP law: it will be obsolete soon enough

References

- Solomonoff, R.J.: Progress in incremental machine learning. Technical Report IDSIA-16-03, IDSIA, Lugano, Switzerland (2003)
- Solomonoff, R.J.: The time scale of artificial intelligence: Reflections on social effects. Human Systems Management 5 (1985) 149–153
- Özkural,E.:Towards heuristic algorithmic memory.In Schmidhuber,J.,Thorisson, K.R., Looks, M., eds.: AGI. Volume 6830 of Lecture Notes in Computer Science., Springer (2011) 382–387
- Özkural, E.: A compromise between reductionism and non-reductionism. In: Worldviews, Science and Us: Philosophy and Complexity. World Scientific Books (2007)
- Lloyd, S. (1996). "Universal quantum simulators". *Science* **273** (5278): 1073–8.