
Algorithmic Randomness as Foundation of Inductive Reasoning and Artificial Intelligence

Marcus Hutter

RSISE @ ANU and SML @ NICTA
Canberra, ACT, 0200, Australia
marcus@hutter1.net www.hutter1.net

5 May 2010

Abstract

This article is a personal account of the past, present, and future of algorithmic randomness, emphasizing its role in inductive inference and artificial intelligence. It is written for a general audience interested in science and philosophy. Intuitively, randomness is a lack of order or predictability. If randomness is the opposite of determinism, then algorithmic randomness is the opposite of computability. Besides many other things, these concepts have been used to quantify Ockham's razor, solve the induction problem, and define intelligence.

Contents

1	Why were you initially drawn to the study of computation and randomness?	2
2	What have we learned?	3
3	What don't we know (yet)?	4
4	What are the most important open problems in the field?	5
5	What are the prospects for progress?	7

Keywords

algorithmic information theory; individual randomness; Ockham's razor; inductive reasoning; artificial intelligence.

1 Why were you initially drawn to the study of computation and randomness?

Some sequences of events follow a long causal “computable” path, while others are so “random” that the coherent causal path is quite short. I am able to trace back quite far my personal causal chain of events that eventually led me to computability and randomness (C&R), although the path looks warped and random.

At about 8 years of age, I got increasingly annoyed at always having to tidy up my room. It took me more than 20 years to see that computation and randomness was the solution to my problem (well – sort of). Here’s a summary of the relevant events:

First, my science fiction education came in handy. I was well-aware that robots were perfectly suitable for all kinds of boring jobs, so they should be good for tidying up my room too. Within a couple of years I had built a series of five increasingly sophisticated robots. The “5th generation” one was humanoid-shaped, about 40cm high, had two arms and hands, and one broad roller leg. The frame was metal and the guts cannibalized from my remote controlled car and other toys.

With enough patience I was able to maneuver Robbie5 with the remote control to the messy regions of my room, have it pick up some Lego pieces and place them in the box they belonged to. It worked! And it was lots of fun. But it didn’t really solve my problem. Picking up a block with the robot took at least 10 times longer than doing it by hand, and even if the throughput was the same, I felt I hadn’t gained much.

Robbie5 was born about a year before my father brought home one of the first programmable pocket calculators in 1978, a HP-19C. With its 200 bytes of RAM or so it was not quite on par with Colossus (a super computer which develops a mind of its own in the homonymous movie), but HP-19C was enough for me to realize that a computer allows programming of a robot to perform a sequence of steps autonomously. Over the following 15 years, I went through a sequence of calculators and computers, wrote increasingly sophisticated software, and studied computer science with a Masters degree in Artificial Intelligence (AI). My motivation in AI of course changed many times over the years, from the dream of a robot tidying up my room to more intellectual, philosophical, economic, and social motivations.

Around 1992 I lost confidence in any of the existing approaches towards AI, and despite considerable effort for years, didn’t have a ground-breaking idea myself.

While working in a start-up company on a difficult image interpolation problem, I realized one day in 1997 that simplicity and compression are key, not only for solving my problem at hand, but also for the grand AI problem.

It took me quite a number of weekends to work out the details. Relatively soon I concluded that the theory I had developed was too beautiful to be novel. I had rediscovered aspects of Kolmogorov complexity and Solomonoff induction. Indeed, I had done more. My system generalized Solomonoff induction to a universally optimal general-purpose reinforcement learning agent.

In order to prove some of my claims it was necessary to become more deeply and broadly acquainted with Algorithmic Information Theory (AIT).

AIT combines information theory and computation theory to an objective and absolute notion of information in an individual object, and gives rise to an objective and robust notion of randomness of individual objects. Its major sub-disciplines are Algorithmic “Kolmogorov” Complexity (AC), Algorithmic “Solomonoff” Probability (AP), Algorithmic “Martin-Löf” Randomness (AR), and Universal “Levin” Search (UL) [Hut07a].

This concludes my 25(!) year journey to C&R. In the last 10 years, I have contributed to all the 4 subfields. My primary driving force when doing research in C&R is still AI, so I’ve most to say about AP, and my answers to the following questions are biased towards my own personal interests.

2 What have we learned?

Let me begin with what *I* have learned: The most important scientific insight I have had is the following: Many scientists have a bias towards elegant or beautiful theories, which usually aligns with some abstract notion of simplicity. Others have a bias towards simple theories in the concrete sense of being analytically or computationally tractable. By ‘theories’ I essentially mean mathematical models of some aspect of the real world, e.g. of a physical system like an engine or the weather or stock market.

Way too late in my life, at age 30 or so, I realized that the most important reason for preferring simple theories is a quite different one: Simple theories tend to be better for what they are developed for in the first place, namely predicting in related but different situations and using these predictions to improve decisions and actions.

Indeed, the principle to prefer simpler theories has been popularized by William of Ockham (1285-1349) (“Entities should not be multiplied unnecessarily”) but dates back at least to Aristotle [Fra02].

Kolmogorov complexity [Kol65] is a universal objective measure of complexity and allows simplicity and hence Ockham’s “razor” principle to be quantified. Solomonoff [Sol64] developed a formal theory of universal prediction along this line, actually a few years before Kolmogorov introduced his closely related complexity measure. My contribution in the 200X [Hut00, Hut07c] was to generalize Solomonoff induction to a universally intelligent learning agent [OC06].

This shows that Ockham’s razor, inductive reasoning, intelligence, and the scientific inquiry itself are intimately related. I would even go so far as to say that science *is* the application of Ockham’s razor: *Simple* explanations of observed real-world phenomena have a higher chance of leading to correct predictions [Hut09b].

What does all this have to do with C&R? We cannot be certain about *anything* in our world. It might even end or be completely different tomorrow. Even if some

proclaimed omniscient entity told us the future, there is no scientific way to verify the premise of its omniscience. So induction has to deal with uncertainty. Making worst-case assumptions is not a generic solution; the generic worst-case is “anything can happen”. Considering restricted model classes begs the question about the validity of the model class itself, so is also not a solution. More powerful is to model uncertainty by probabilities and the latter is obviously related to randomness.

There have been many attempts to formalize probability and randomness: Kolmogorov’s axioms of probability theory [Kol33] are the default characterization. Problems with this notion are discussed in item (d) of Question 4. Early attempts to define the notion of randomness of *individual* objects/sequences by von Mises [Mis19], Wald [Wal37], and Church [Chu40] failed, but finally Martin-Löf [ML66] succeeded. A sequence is Martin-Löf random if and only if it passes all effective randomness tests or, as it turns out, if and only if it is incompressible.

3 What don’t we know (yet)?

Lots of things, so I will restrict myself to open problems in the intersection of AIT and AI. See [Hut09a] for details.

(i) Universal Induction: The induction problem is a fundamental problem in philosophy [Hum39, Ear93] and statistics [Jay03] and science in general. The most important fundamental philosophical and statistical problems around induction are discussed in [Hut07b]: Among others, they include the problems of old evidence, ad-hoc hypotheses, updating, zero prior, and invariance. The arguments in [Hut07b] that Solomonoff’s universal theory M overcomes these problems are forceful but need to be elaborated on further to convince the (scientific) world that the induction problem is essentially solved. Besides these general induction problems, universal induction raises many additional questions: for instance, it is unclear whether M can predict all computable *subsequences* of a sequence that is itself not computable, how to formally identify “natural” Turing machines [Mül10], Martin-Löf convergence of M , and whether AIXI (see below) reduces to M for prediction.

(ii) Universal Artificial Intelligence (UAI): The AIXI model integrates Solomonoff induction with sequential decision theory. As a unification of two optimal theories in their own domains, it is plausible that AIXI is optimal in the “union” of their domains. This has been affirmed by positive pareto-optimality and self-optimizingness results [Hut05]. These results support the claim that AIXI is a universally optimal generic reinforcement learning agent, but unlike the induction case, the results so far are not yet strong enough to allay all doubts. Indeed, the major problem is not to *prove* optimality but to *come up* with sufficiently strong but still satisfiable optimality notions in the reinforcement learning case. A more modest goal than proving optimality of AIXI is to ask for additional reasonable convergence properties, like posterior convergence for unbounded horizon. The probably most important fundamental and hardest problem in game theory is the grain-of-truth problem [KL93].

In our context, the question is what happens if AIXI is used in a multi-agent setup [SLB08] interacting with other instantiations of AIXI.

(iii) Defining Intelligence: A fundamental and long standing difficulty in the field of AI is that intelligence itself is not well defined. Usually, formalizing and rigorously defining a previously vague concept constitutes a quantum leap forward in the corresponding field, and AI should be no exception. AIT again suggested an extremely general, objective, fundamental, and formal measure of machine intelligence [Hut00, Leg08, GR05, Fié05], but the theory surrounding it has yet to be adequately explored. A comprehensive collection, discussion and comparison of verbal and formal intelligence tests, definitions, and measures can be found in [LH07].

4 What are the most important open problems in the field?

There are many important open technical problems in AIT. I have discussed some of those that are related to AI in [Hut09a] and in the previous answer. Here I concentrate on the most important open problems in C&R which I am able to describe in non-technical terms.

(a) The development of notions of complexity and individual randomness didn't end with Kolmogorov and Martin-Löf. Many variants of “plain” Kolmogorov complexity C [Kol65] have been developed: prefix complexity K [Lev74, Gác74, Cha75], process complexity [Sch73], monotone complexity K_m [Lev73], uniform complexity [Lov69b, Lov69a], Chaitin complexity K_c [Cha75], Solomonoff's universal prior $M = 2^{-KM}$ [Sol64, Sol78], extension semimeasure M_c [Cov74], and some others [LV08]. They often differ only by $O(\log K)$, but this can lead to important differences. Variants of Martin-Löf randomness are: Schnorr randomness [Sch71], Kurtz randomness [Kur81], Kolmogorov-Loveland randomness [Lov66], and others [Wan96, Cal02, DH07]. All these complexity and randomness classes can further be relativized to some oracle, e.g. the halting oracle, leading to an arithmetic hierarchy of classes. Invoking resource-bounds moves in the other direction and leads to the well-known complexity zoo [AKG05] and pseudo-randomness [Lub96]. Which definition is the “right” or “best” one, and in which sense? Current research on algorithmic randomness is more concerned about abstract properties and convenience, rather than practical usefulness. This is in marked contrast to complexity theory, in which the classes also sprout like mushrooms [AKG05], but the majority of classes delineate important practical problems.

(b) The historically oldest, non-flawed, most popular, and default notion of individual randomness is that of Martin-Löf. Let us assume that it is or turns out to be the “best” or single “right” definition of randomness. This would uniquely determine which individual infinite sequences are random and which are not. This unfortunately does not hold for finite sequences. This non-uniqueness problem is

equivalent to the problem that Kolmogorov complexity depends on the choice of universal Turing machine. While the choice is asymptotically, and hence for large-scale practical applications, irrelevant, it seriously hinders applications to “small” problems. One can argue the problem away [Hut05], but finding a unique “best” universal Martin-Löf test or universal Turing machine would be more satisfactory and convincing. Besides other things, it would make inductive inference absolutely objective.

(c) Maybe randomness can, in principle, only be relative: What looks random to me might be order to you. So randomness depends on the power of the “observer”. In this case, it is important to study problem-specific randomness notions, and clearly describe and separate the application domains of the different randomness notions, like classical sufficient statistics depends on the model class. Algorithmic randomness usually includes all computable tests and goes up the arithmetic hierarchy. For practical applications, limited classes, like all efficiently computable tests, are more relevant. This is the important domain of pseudo-random number generation. Could every practically useful complexity class correspond to a randomness class with practically relevant properties?

(d) It is also unclear whether algorithmic randomness or classical probability theory has a more fundamental status. While measure theory is mathematically, and statistics is practically very successful, Kolmogorov’s probability axioms are philosophically crippled and, strictly speaking, induce a purely formal but meaningless measure theory exercise. The easygoing frequentist interpretation is circular: The probability of head is p , if the long-run relative frequency tends to p almost surely (with probability one). But what does ‘almost surely’ mean? Applied statistics implicitly invokes Cournot’s somewhat forgotten principle: An event with very small probability, singled out in advance, will not happen. That is, a probability 1 event will happen for sure in the real world. Another problem is that it is not even possible to ask the question of whether a particular single sequence of observations is random (w.r.t. some measure). Algorithmic randomness makes this question meaningful and answerable. A downside of algorithmic randomness is that not every set of measure 1 will do, but only constructive ones, which can be much harder to find and sometimes do not exist [HM07].

(e) Finally, to complete the circle, let’s return to my original motivation for entering this field: Ockham’s razor (1) is the key philosophical ingredient for solving the induction problem and crucial in defining science and intelligence, and (2) can be quantified in terms of algorithmic complexity which itself is closely related to algorithmic randomness. The formal theory of universal induction [Sol78, Hut07b] is already well-developed and the foundations of universal AI have been laid [Hut05]. Besides solving specific problems like (i)-(iii) and (a)-(d) above, it is also important to “translate” the results and make them accessible to researchers in other disciplines: present the philosophical insights in a less-mathematical way; stress that sound mathematical foundations are crucial for advances in most field, and induction and AI should be no exception; etc.

5 What are the prospects for progress?

The prospects for the open problems (a)-(e) of Question 4 I believe are as follows:

(a) I am not sure about the fate of the multitude of different randomness notions. I can't see any practical relevance for those in the arithmetic hierarchy. Possibly the acquired scientific knowledge from studying the different classes and their relationship can be used in a different field in an unexpected way. For instance, the ones in the arithmetic hierarchy may be useful in the endeavor of unifying probability and logic [GS82]. Possibly the whole idea of *objectively* identifying individually which strings shall be regarded as random will be given up.

(b) All scientists, except some logicians studying logic, essentially use the same classical logic and axioms, namely ZFC, to do *deductive* reasoning. Why do not all scientists use the same definition of probability to do *inductive* reasoning? Bayesian statistics and Martin-Löf randomness are the most promising candidates for becoming universally accepted for inductive reasoning.

Maybe they will become universally accepted some time in the future, for pragmatic reasons, or simply as a generally agreed upon convention, since no one is interested in arguing over it anymore. While Martin-Löf uniquely determines infinite random sequences, the randomness for finite sequences depends on the choice of universal Turing machine. Finding a unique “best” one (if possible) is, in my opinion, the most important open problem in algorithmic randomness. A conceptual breakthrough would be needed to make progress on this hard front. See [Mül10] for a remarkable but failed recent attempt.

(c) Maybe pursuing a single definition of randomness is illusory. Noise might simply be that aspect of the data that is not useful for the particular task or method at hand. For instance, sufficient statistics and pseudo-random numbers have this task-dependence. Even with a single fundamental notion of randomness (see b) there will be many different practical approximations. I expect steady progress on this front.

(d) Bayesian statistics based on classical probability theory is incomplete, since it does not tell you how to choose the prior. Solomonoff fixes the prior to a negative exponential in the model complexity. Time and further research will convince classical statisticians to accept this (for them now) exotic choice as a kind of “gold standard” (as Solomonoff put it). All this is still within the classical measure theoretic framework, which may be combined with Cournot *or* with Martin-Löf.

(e) Finally, convincing AI researchers and philosophers about the importance of Ockham's razor, that algorithmic complexity is a suitable quantification, and that this led to a formal (albeit non-computable) conceptual solution to the induction and the AI problem should be a matter of a decade or so.

References

- [AKG05] S. Aaronson, G. Kuperberg, and C. Granade. Complexity zoo, 2005. <http://www.complexityzoo.com/>.
- [Cal02] C. S. Calude. *Information and Randomness: An Algorithmic Perspective*. Springer, Berlin, 2nd edition, 2002.
- [Cha75] G. J. Chaitin. A theory of program size formally identical to information theory. *Journal of the ACM*, 22(3):329–340, 1975.
- [Chu40] A. Church. On the concept of a random sequence. *Bulletin of the American Mathematical Society*, 46:130–135, 1940.
- [Cov74] T. M. Cover. Universal gambling schemes and the complexity measures of Kolmogorov and Chaitin. Technical Report 12, Statistics Department, Stanford University, Stanford, CA, 1974.
- [DH07] R. Downey and D. R. Hirschfeldt. *Algorithmic Randomness and Complexity*. Springer, Berlin, 2007.
- [Ear93] J. Earman. *Bayes or Bust? A Critical Examination of Bayesian Confirmation Theory*. MIT Press, Cambridge, MA, 1993.
- [Fié05] C. Fiévet. Mesurer l’intelligence d’une machine. In *Le Monde de l’intelligence*, volume 1, pages 42–45, Paris, November 2005. Mondeo publishing.
- [Fra02] J. Franklin. *The Science of Conjecture: Evidence and Probability before Pascal*. Johns Hopkins University Press, 2002.
- [Gác74] P. Gács. On the symmetry of algorithmic information. *Soviet Mathematics Doklady*, 15:1477–1480, 1974.
- [GR05] D. Graham-Rowe. Spotting the bots with brains. In *New Scientist magazine*, volume 2512, page 27, 13 August 2005.
- [GS82] H. Gaifman and M. Snir. Probabilities over rich languages, testing and randomness. *Journal of Symbolic Logic*, 47:495–548, 1982.
- [HM07] M. Hutter and An. A. Muchnik. On semimeasures predicting Martin-Löf random sequences. *Theoretical Computer Science*, 382(3):247–261, 2007.
- [Hum39] D. Hume. *A Treatise of Human Nature, Book I*. [Edited version by L. A. Selby-Bigge and P. H. Nidditch, Oxford University Press, 1978], 1739.
- [Hut00] M. Hutter. A theory of universal artificial intelligence based on algorithmic complexity. Technical Report arXiv:cs.AI/0004001, München, 62 pages, 2000.
- [Hut05] M. Hutter. *Universal Artificial Intelligence: Sequential Decisions based on Algorithmic Probability*. Springer, Berlin, 2005.
- [Hut07a] M. Hutter. Algorithmic information theory: a brief non-technical guide to the field. *Scholarpedia*, 2(3):2519, 2007.
- [Hut07b] M. Hutter. On universal prediction and Bayesian confirmation. *Theoretical Computer Science*, 384(1):33–48, 2007.
- [Hut07c] M. Hutter. Universal algorithmic intelligence: A mathematical top→down approach. In *Artificial General Intelligence*, pages 227–290. Springer, Berlin, 2007.
- [Hut09a] M. Hutter. Open problems in universal induction & intelligence. *Algorithms*, 3(2):879–906, 2009.
- [Hut09b] M. Hutter. A complete theory of everything (will be subjective). 2009. arXiv:0912.5434.
- [Jay03] E. T. Jaynes. *Probability Theory: The Logic of Science*. Cambridge University Press, Cambridge, MA, 2003.
- [KL93] E. Kalai and E. Lehrer. Rational learning leads to Nash equilibrium. *Econometrica*, 61(5):1019–1045, 1993.
- [Kol33] A. N. Kolmogorov. *Grundlagen der Wahrscheinlichkeitsrechnung*. Springer, Berlin, 1933. [English translation: *Foundations of the Theory of Probability*. Chelsea, New York, 2nd edition, 1956].

- [Kol65] A. N. Kolmogorov. Three approaches to the quantitative definition of information. *Problems of Information and Transmission*, 1(1):1–7, 1965.
- [Kur81] S. A. Kurtz. *Randomness and Genericity in the Degrees of Unsolvability*. PhD thesis, University of Illinois, 1981.
- [Leg08] S. Legg. *Machine Super Intelligence*. PhD thesis, IDSIA, Lugano, 2008.
- [Lev73] L. A. Levin. On the notion of a random sequence. *Soviet Mathematics Doklady*, 14(5):1413–1416, 1973.
- [Lev74] L. A. Levin. Laws of information conservation (non-growth) and aspects of the foundation of probability theory. *Problems of Information Transmission*, 10(3):206–210, 1974.
- [LH07] S. Legg and M. Hutter. Universal intelligence: A definition of machine intelligence. *Minds & Machines*, 17(4):391–444, 2007.
- [Lov66] D. E. Loveland. A new interpretation of von Mises’ concept of a random sequence. *Zeitschrift für Mathematische Logik und Grundlagen der Mathematik*, 12:279–294, 1966.
- [Lov69a] D. W. Loveland. On minimal-program complexity measures. In *Proc. 1st ACM Symposium on Theory of Computing*, pages 61–78. ACM Press, New York, 1969.
- [Lov69b] D. W. Loveland. A variant of the Kolmogorov concept of complexity. *Information and Control*, 15(6):510–526, 1969.
- [Lub96] M. Luby. *Pseudorandomness and Cryptographic Applications*. Princeton University Press, 1996.
- [LV08] M. Li and P. M. B. Vitányi. *An Introduction to Kolmogorov Complexity and its Applications*. Springer, Berlin, 3rd edition, 2008.
- [Mis19] R. von Mises. Grundlagen der Wahrscheinlichkeitsrechnung. *Mathematische Zeitschrift*, 5:52–99, 1919. Correction, *Ibid.*, volume 6, 1920, [English translation in: *Probability, Statistics, and Truth*, Macmillan, 1939].
- [ML66] P. Martin-Löf. The definition of random sequences. *Information and Control*, 9(6):602–619, 1966.
- [Mül10] M. Müller. Stationary algorithmic probability. *Theoretical Computer Science*, 411(1):113–130, 2010.
- [OC06] T. Oates and W. Chong. Book review: Marcus Hutter, universal artificial intelligence, Springer (2004). *Artificial Intelligence*, 170(18):1222–1226, 2006.
- [Sch71] C. P. Schnorr. *Zufälligkeit und Wahrscheinlichkeit*, volume 218 of *Lecture Notes in Mathematics*. Springer, Berlin, 1971.
- [Sch73] C. P. Schnorr. Process complexity and effective random tests. *Journal of Computer and System Sciences*, 7(4):376–388, 1973.
- [SLB08] Y. Shoham and K. Leyton-Brown. *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*. Cambridge University Press, 2008.
- [Sol64] R. J. Solomonoff. A formal theory of inductive inference: Parts 1 and 2. *Information and Control*, 7:1–22 and 224–254, 1964.
- [Sol78] R. J. Solomonoff. Complexity-based induction systems: Comparisons and convergence theorems. *IEEE Transactions on Information Theory*, IT-24:422–432, 1978.
- [Wal37] A. Wald. Die Widerspruchsfreiheit des Kollektivbegriffs in der Wahrscheinlichkeitsrechnung. In *Ergebnisse eines Mathematischen Kolloquiums*, volume 8, pages 38–72, 1937.
- [Wan96] Y. Wang. *Randomness and Complexity*. PhD thesis, Universität Heidelberg, 1996.