
UNIVERSAL ALGORITHMIC INTELLIGENCE

A Mathematical Top→Down Approach

Marcus Hutter

Canberra, ACT, 0200, Australia

<http://www.hutter1.net/>



ANU



RSISE



NICTA

Abstract

The dream of creating artificial devices that reach or outperform human intelligence is many centuries old. Nowadays most research is more modest, focussing on solving narrower, specific problems, associated with only some aspects of intelligence, like playing chess or natural language translation, either as a goal in itself or as a bottom-up approach. The dual top down approach investigates theories of general-purpose intelligent agents: the power of such theoretical agents, how to scale them down, and the involved key concepts. Necessary ingredients seem to be Occam's razor; Turing machines; Kolmogorov complexity; probability theory; Solomonoff induction; Bayesian sequence prediction; minimum description length principle; agent framework; sequential decision theory; adaptive control theory; reinforcement learning; Levin search and extensions, which are all important subjects in their own right. From a mathematical point of view these concepts also seem to be sufficient.

Contents

- What is (Artificial) Intelligence.
- Philosophical Foundations.
(Ockham, Epicurus, Induction)
- Mathematical Foundations.
(Information, Complexity, Bayesian & Algorithmic Probability, Solomonoff induction, Sequential Decision)
- Framework: Rational Agents.
(in Known and Unknown Environments)
- AIXI: Universal Artificial Intelligence.
- Approximations & Applications.
(Universal Search, $AIXI_{tl}$, AI_{ξ} , MC-AIXI-CTW, Φ MDP)
- Comparison to Other Approaches.
- Discussion.
(Summary, Criticism, Questions, Next Steps, Literature)

What is (Artificial) Intelligence?

Intelligence can have many faces \Rightarrow formal definition difficult

- reasoning
- creativity
- association
- generalization
- pattern recognition
- problem solving
- memorization
- planning
- achieving goals
- learning
- optimization
- self-preservation
- vision
- language processing
- classification
- induction
- deduction
- ...

What is AI?	Thinking	Acting
humanly	Cognitive Science	Turing test, Behaviorism
rationally	Laws Thought	Doing the Right Thing

Collection of 70+ Defs of Intelligence

<http://www.vetta.org/>

[definitions-of-intelligence/](http://www.vetta.org/definitions-of-intelligence/)

Real world is nasty: partially unobservable, uncertain, unknown, non-ergodic, reactive, vast, but luckily structured, ...

Informal Definition of (Artificial) Intelligence

Intelligence measures an agent's ability to achieve goals in a wide range of environments. [S. Legg and M. Hutter]

Emergent: Features such as the ability to learn and adapt, or to understand, are implicit in the above definition as these capacities enable an agent to succeed in a wide range of environments.

The science of **Artificial Intelligence** is concerned with the construction of intelligent systems/artifacts/agents and their analysis.

What next? Substantiate all terms above: agent, ability, utility, goal, success, learn, adapt, environment, ...

Never trust a ~~theory~~ if it is not supported by an ~~experiment~~
experiment **theory**

Induction → Prediction → Decision → Action

Having or acquiring or *learning* or *inducing* a model of the environment an agent interacts with allows the agent to make *predictions* and utilize them in its *decision* process of finding a good next *action*.

Induction infers general models from specific observations/facts/data, usually exhibiting regularities or properties or relations in the latter.

Example

Induction: Find a model of the world economy.

Prediction: Use the model for predicting the future stock market.

Decision: Decide whether to invest assets in stocks or bonds.


Action: Trading large quantities of stocks influences the market.

Science \approx Induction \approx Occam's Razor

- Grue Emerald Paradox:

Hypothesis 1: All emeralds are green.

Hypothesis 2: All emeralds found till y2020 are green,
thereafter all emeralds are blue.



- Which hypothesis is more plausible? **H1!** Justification?
- **Occam's razor**: take simplest hypothesis consistent with data.
is the most important principle in machine learning and science.

Information Theory & Kolmogorov Complexity

- Quantification/interpretation of Occam's razor:
- Shortest description of object is best explanation.
- Shortest program for a string on a Turing machine T leads to best extrapolation=prediction.



$$K_T(x) = \min_p \{l(p) : T(p) = x\}$$

- Prediction is best for a universal Turing machine U .

$$\text{Kolmogorov-complexity}(x) = K(x) = K_U(x) \leq K_T(x) + c_T$$

- Recent impressive applications: universal similarity metric \uparrow [LV08].
[Clustering by Compression, Cilibrasi & Vitanyi 2005]

Bayesian Probability Theory

Given (1): Models $P(D|H_i)$ for probability of observing data D , when H_i is true.

Given (2): Prior probability over hypotheses $P(H_i)$

Goal: Posterior probability $P(H_i|D)$ of H_i , after having seen data D .



Solution:

Bayes' rule:


$$P(H_i|D) = \frac{P(D|H_i) \cdot P(H_i)}{\sum_i P(D|H_i) \cdot P(H_i)}$$

(1) Models $P(D|H_i)$ usually easy to describe (objective probabilities)

(2) But Bayesian prob. theory does not tell us how to choose the prior $P(H_i)$ (subjective probabilities)

Algorithmic Probability Theory

- **Epicurus**: If more than one theory is consistent with the observations, keep all theories.
- \Rightarrow uniform prior over all H_i ?
- Refinement with **Occam's razor** quantified in terms of **Kolmogorov complexity**:




$$P(H_i) := 2^{-K_{T/U}(H_i)}$$

- **Fixing T** we have a complete theory for prediction.
Problem: How to choose T .
- **Choosing U** we have a universal theory for prediction.
Observation: Particular choice of U does not matter much.
Problem: Incomputable.

Inductive Inference & Universal Forecasting

- Solomonoff combined Occam, Epicurus, Bayes, and Turing into one formal theory of sequential prediction.
- $M(x)$ = probability that a universal Turing machine outputs x when provided with fair coin flips on the input tape.
- A posteriori probability of y given x is $M(y|x) = M(xy)/M(x)$.
- Given x_1, \dots, x_{t-1} , the probability of x_t is $M(x_t|x_1 \dots x_{t-1})$.
- Immediate “applications”:
 - Weather forecasting: $x_t \in \{\text{sun, rain}\}$.
 - Stock-market prediction: $x_t \in \{\text{bear, bull}\}$.
 - Continuing number sequences in an IQ test: $x_t \in \mathbb{N}$.
- Optimal universal inductive reasoning system!



Sequential Decision Theory

Setup: For $t = 1, 2, 3, 4, \dots$

Given sequence x_1, x_2, \dots, x_{t-1}

(1) predict/make decision y_t ,

(2) observe x_t ,

(3) suffer loss $\text{Loss}(x_t, y_t)$,

(4) $t \rightarrow t + 1$, goto (1)

Goal: Minimize expected Loss.

Greedy minimization of expected loss **is optimal** if:

Important: Decision y_t does not influence env. (future observations).

Loss function is known.

Problem: Expectation w.r.t. what?

Solution: W.r.t. universal distribution M if true distr. is unknown.


Example: Weather Forecasting

Observation $x_t \in \mathcal{X} = \{\text{sunny, rainy}\}$

Decision $y_t \in \mathcal{Y} = \{\text{umbrella, sunglasses}\}$

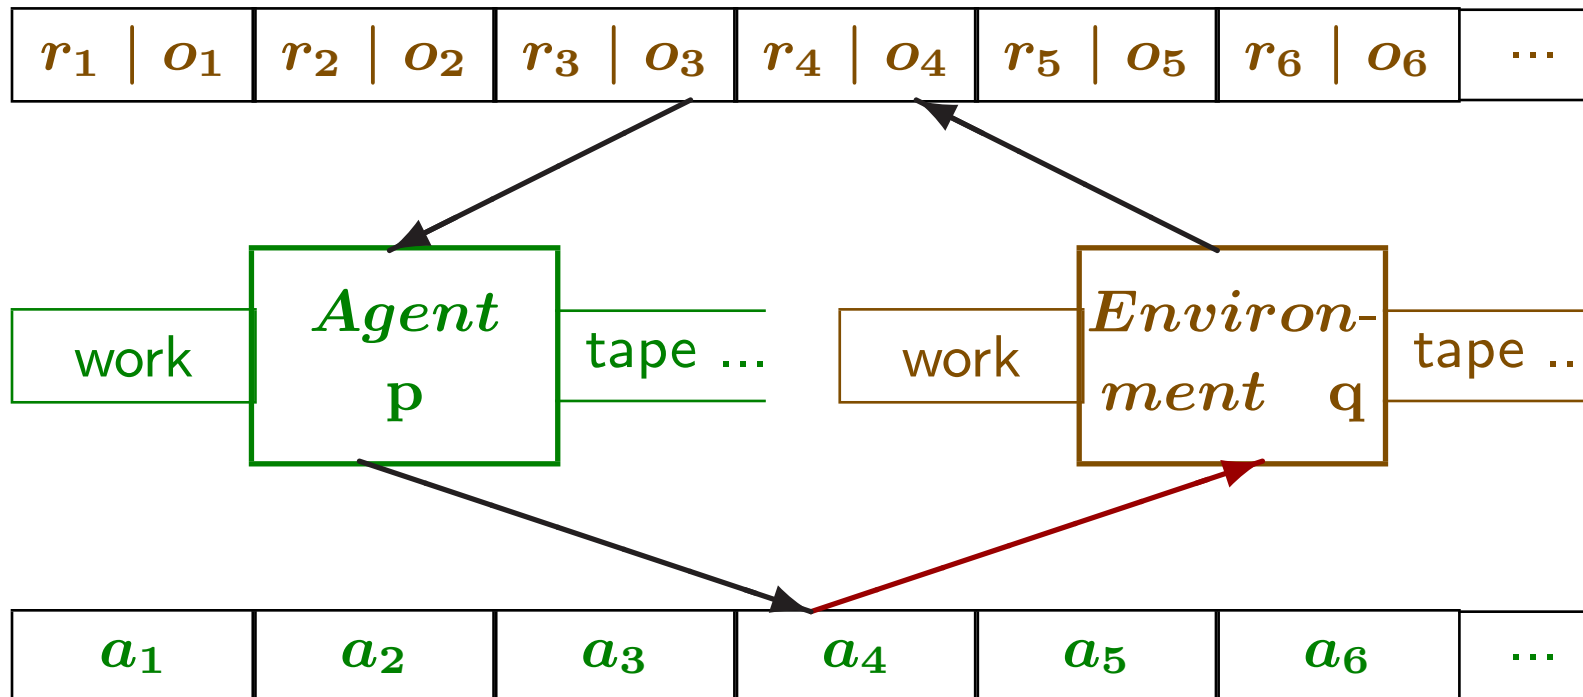
Loss	sunny	rainy
umbrella	0.1	0.3
sunglasses	0.0	1.0

Taking umbrella/sunglasses does not influence future weather
(ignoring butterfly effect)



Agent Model with Reward

if actions/decisions a
influence the environment q



Rational Agents in Known Environment

- **Setup:** Known deterministic or probabilistic environment
- **Greedy** maximization of reward r ($= -\text{Loss}$) **no longer optimal.**
Example: Chess
- **Exploration versus exploitation problem.**
 \Rightarrow **Agent has to be farsighted.**
- **Optimal solution:** Maximize future (expected) reward sum, called value.
- **Problem:** Things drastically change if environment is unknown

Rational Agents in Unknown Environment

Additional problem: (probabilistic) environment unknown.

Fields: reinforcement learning and adaptive control theory

Bayesian approach: Mixture distribution.

1. What performance does Bayes-optimal policy imply?

It does not necessarily imply self-optimization
(Heaven&Hell example).

2. Computationally very hard problem.

3. Choice of horizon? Immortal agents are lazy.

Universal Solomonoff mixture \Rightarrow universal agent AIXI.

Represents a formal (math., non-comp.) solution to the AI problem?

Most (all?) problems are easily phrased within AIXI.

The AIXI Model in one Line

complete & essentially unique & limit-computable

$$\text{AIXI: } a_k := \arg \max_{a_k} \sum_{O_k r_k} \dots \max_{a_m} \sum_{O_m r_m} [r_k + \dots + r_m] \sum_{q: U(q, a_1 \dots a_m) = O_1 r_1 \dots O_m r_m} 2^{-\ell(q)}$$

action, reward, observation, Universal TM, qrogram, k=now

AIXI is an elegant mathematical theory of AI

Claim: AIXI is the most intelligent environmental independent, i.e. universally optimal, agent possible.

Proof: For formalizations, quantifications, and proofs, see [Hut05].

Potential Applications: Agents, Games, Optimization, Active Learning, Adaptive Control, Robots.

Computational Issues: Universal Search

- **Levin search:** Fastest algorithm for inversion and optimization problems.
- **Theoretical application:**
Assume somebody found a non-constructive proof of $P=NP$, then Levin-search is a polynomial time algorithm for every NP (complete) problem.
- **Practical (OOPS) applications** (J. Schmidhuber)
Maze, towers of hanoi, robotics, ...
- **FastPrg:** The asymptotically fastest and shortest algorithm for all well-defined problems.
- **Computable Approximations of AIXI:**
 $AIXItl$ and $AI\xi$ and MC-AIXI-CTW and Φ MDP.
- **Human Knowledge Compression Prize:** (50'000€)



The Time-Bounded AIXI Model ($AIXI_{tl}$)

An algorithm p^{best} has been constructed for which the following holds:









- Let p be any (extended chronological) policy
- with length $l(p) \leq \tilde{l}$ and computation time per cycle $t(p) \leq \tilde{t}$
- for which there exists a proof of length $\leq l_P$ that p is a valid approximation of AIXI.
- Then an algorithm p^{best} can be constructed, depending on \tilde{l}, \tilde{t} and l_P but not on knowing p
- which is effectively more or equally intelligent than any such p .
- The size of p^{best} is $l(p^{best}) = O(\ln(\tilde{l} \cdot \tilde{t} \cdot l_P))$,
- the setup-time is $t_{setup}(p^{best}) = O(l_P^2 \cdot 2^{l_P})$,
- the computation time per cycle is $t_{cycle}(p^{best}) = O(2^{\tilde{l}} \cdot \tilde{t})$.


Brute-Force Approximation of AIXI

- **Truncate expectimax tree** depth to a small fixed lookahead h .
Optimal action computable in time $|\mathcal{Y} \times \mathcal{X}|^h \times \text{time to evaluate } \xi$.
- Consider mixture over **Markov Decision Processes** (MDP) only, i.e.
 $\xi(x_{1:m}|y_{1:m}) = \sum_{\nu \in \mathcal{M}} w_\nu \prod_{t=1}^m \nu(x_t|x_{t-1}y_t)$. Note: ξ is *not* MDP
- Choose **uniform prior** over w_μ .
Then $\xi(x_{1:m}|y_{1:m})$ can be computed in linear time.
- Consider (approximately) Markov problems
with very **small action and perception space**.
- **Example application:** 2×2 Matrix Games like Prisoner'S Dilemma,
Stag Hunt, Chicken, Battle of Sexes, and Matching Pennies. [PH'06]

AIXI Learns to Play 2x2 Matrix Games

- Repeated prisoners dilemma.
- Game unknown to AIXI.
Must be learned as well
- AIXI behaves appropriately.

Loss matrix			
		cooperates	defects
	cooperates	 = 0.3 years	 = 1 year
	defects	 free	 = 0.7 years



A Monte-Carlo AIXI Approximation


Consider class of **Variable-Order Markov Decision Processes**.

The **Context Tree Weighting (CTW)** algorithm can efficiently mix (exactly in essentially linear time) all prediction suffix trees.

Monte-Carlo approximation of expectimax tree:

Upper Confidence Tree (UCT) algorithm:

- **Sample** observations from CTW distribution.
- **Select** actions with highest upper confidence bound.
- **Expand** tree by one leaf node (per trajectory).
- **Simulate** from leaf node further down using (fixed) playout policy.
- **Propagate back** the value estimates for each node.




Repeat until timeout.

[VNHS'09]

Guaranteed to **converge** to exact value.

Extension: Predicate CTW not based on raw obs. but features thereof.

Monte-Carlo AIXI Applications



Properties of Learning Algorithms

Comparison of AIXI to Other Approaches

Algorithm	time efficient	data efficient	explo-ration	conver-geance	global optimum	genera-lization	POMDP	learning	active
Value/Policy iteration	yes/no	yes	-	YES	YES	NO	NO	NO	yes
TD w. func.approx.	no/yes	NO	NO	no/yes	NO	YES	NO	YES	YES
Direct Policy Search	no/yes	YES	NO	no/yes	NO	YES	no	YES	YES
Logic Planners	yes/no	YES	yes	YES	YES	no	no	YES	yes
RL with Split Trees	yes	YES	no	YES	NO	yes	YES	YES	YES
Pred.w. Expert Advice	yes/no	YES	-	YES	yes/no	yes	NO	YES	NO
OOPS	yes/no	no	-	yes	yes/no	YES	YES	YES	YES
Market/Economy RL	yes/no	no	NO	no	no/yes	yes	yes/no	YES	YES
SPXI	no	YES	-	YES	YES	YES	NO	YES	NO
AIXI	NO	YES	YES	yes	YES	YES	YES	YES	YES
AIXI _{tl}	no/yes	YES	YES	YES	yes	YES	YES	YES	YES
MC-AIXI-CTW	yes/no	yes	YES	YES	yes	NO	no/yes	YES	YES
Feature RL	yes/no	YES	yes	yes	yes	yes	yes	YES	YES
Human	yes	yes	yes	no/yes	NO	YES	YES	YES	YES

Machine Intelligence Tests & Definitions

Intelligence Test	Valid	Informative	Wide Range	General	Dynamic	Unbiased	Fundamental	Formal	Objective	Fully Defined	Universal	Practical	Test vs. Def.
Turing Test	●	·	·	·	●	·	·	·	·	●	·	●	⊥
Total Turing Test	●	·	·	·	●	·	·	·	·	●	·	·	⊥
Inverted Turing Test	●	●	·	·	●	·	·	·	·	●	·	●	⊥
Toddler Turing Test	●	·	·	·	●	·	·	·	·	·	·	●	⊥
Linguistic Complexity	●	★	●	·	·	·	·	●	●	·	●	●	⊥
Text Compression Test	●	★	★	●	·	●	●	★	★	★	●	★	⊥
Turing Ratio	●	★	★	★	?	?	?	?	?	·	?	?	T/D
Psychometric AI	★	★	●	★	?	●	·	●	●	●	·	●	T/D
Smith's Test	●	★	★	●	·	?	★	★	★	·	?	●	T/D
C-Test	●	★	★	●	·	★	★	★	★	★	★	★	T/D
AIXI	★	★	★	★	★	★	★	★	★	★	★	·	D


★ = yes, · = no,
 ● = debatable,
 ? = unknown.

Summary

- Sequential **Decision Theory** solves the problem of rational agents in uncertain worlds if the environmental probability distribution is known.
- Solomonoff's theory of **Universal Induction** solves the problem of sequence prediction for unknown prior distribution.
- Combining both ideas one arrives at

A Unified View of Artificial Intelligence

$$\begin{array}{rcl}
 & = & \\
 \text{Decision Theory} & = & \text{Probability} + \text{Utility Theory} \\
 + & & + \\
 \text{Universal Induction} & = & \text{Ockham} + \text{Bayes} + \text{Turing}
 \end{array}$$



Common Criticisms

- AIXI is obviously wrong.
(intelligence cannot be captured in a few simple equations)
- AIXI is obviously correct. (everybody already knows this)
- Assuming that the environment is computable is too strong|weak.
- All standard objections to strong AI also apply to AIXI.
(free will, lookup table, Lucas/Penrose Gödel argument)
- AIXI doesn't deal with X or cannot do X .
(X = consciousness, creativity, imagination, emotion, love, soul, etc.)
- AIXI is not intelligent because it cannot choose its goals.
- Universal AI is impossible due to the No-Free-Lunch theorem.

See [Legg:08] for refutations of these and more criticisms.

General Murky&Quirky AI Questions

- Does current mainstream AI research has anything todo with AGI?
- Are sequential decision and algorithmic probability theory all we need to well-define AI?
- What is (Universal) AI theory good for?
- What are robots|embodiment good for in AI?
- Is intelligence a fundamentally simple concept?
(compare with fractals or physics theories)
- What can we (not) expect from super-intelligent agents?
- Is maximizing the expected reward the right criterion?

Next Steps

- Address the many open theoretical questions (see Hutter:05).
- Bridge the gap between (Universal) AI theory and AI practice.
- Explore what role logical reasoning, knowledge representation, vision, language, etc. play in Universal AI.
- Determine the right discounting of future rewards.
- Develop the right nurturing environment for a learning agent.
- Consider embodied agents (e.g. internal \leftrightarrow external reward)
- Analyze AIXI in the multi-agent setting.

The Big Questions

- Is non-computational physics relevant to AI? [Penrose]
- Could something like the number of wisdom Ω prevent a simple solution to AI? [Chaitin]
- Do we need to understand consciousness before being able to understand AI or construct AI systems?
- What if we succeed?

Introductory Literature

- [HMU01] J. E. Hopcroft, R. Motwani, and J. D. Ullman. *Introduction to Automata Theory, Language, and Computation*. Addison-Wesley, 2nd edition, 2001.
- [RN03] S. J. Russell and P. Norvig. *Artificial Intelligence. A Modern Approach*. Prentice-Hall, Englewood Cliffs, NJ, 2nd edition, 2003.
- [LV08] M. Li and P. M. B. Vitányi. *An Introduction to Kolmogorov Complexity and its Applications*. Springer, Berlin, 3rd edition, 2008.
- [SB98] R. S. Sutton and A. G. Barto. *Reinforcement Learning: An Introduction*. MIT Press, Cambridge, MA, 1998.
- [Leg08] S. Legg. *Machine Super Intelligence*. PhD Thesis, Lugano, 2008.
- [Hut05] M. Hutter. *Universal Artificial Intelligence: Sequential Decisions based on Algorithmic Probability*. Springer, Berlin, 2005.