Towards a Universal Theory of Artificial Intelligence based on Algorithmic Probability and Sequential Decision Theory

Marcus Hutter

IDSIA, Galleria 2, CH-6928 Manno-Lugano, Switzerland

marcus@idsia.ch http://www.idsia.ch

Technical Report IDSIA-14-00, 15. December 2000

Abstract

Decision theory formally solves the problem of rational agents in uncertain worlds if the true environmental probability distribution is known. Solomonoff's theory of universal induction formally solves the problem of sequence prediction for unknown distribution. We unify both theories and give strong arguments that the resulting universal AIξ model behaves optimal in any computable environment. The major drawback of the AIξ model is that it is uncomputable. To overcome this problem, we construct a modified algorithm AIξ^{tl}, which is still superior to any other time t and space l bounded agent. The computation time of AIξ^{tl} is of the order $t \cdot 2^l$.

Key Words

Rational agents, sequential decision theory, universal Solomonoff induction, algorithmic probability, reinforcement learning, computational complexity, theorem proving, probabilistic reasoning, Kolmogorov complexity, Levin search.

1 Introduction

The most general framework for Artificial Intelligence is the picture of an *agent* interacting with an environment [Russell and Norvig, 1995]. If the goal is not pre-specified, the agent has to learn by occasional reinforcement feedback [Sutton and Barto, 1998]. If the agent shall be universal, no assumption about the environment may be made, besides that there *exists* some exploitable structure at all. We may ask for the most intelligent way an agent could behave, or, about the optimal way of learning in terms of real world interaction cycles. *Decision theory* formally[1] solves this problem only if the true environmental probability distribution is known (e.g. Backgammon) [Bellman, 1957; Bertsekas and Tsitsiklis, 1996]. [Solomonoff, 1964; 1978] formally solves the problem of *induction* if the true distribution is unknown but only if the agent cannot influence the environment (e.g. weather forecasts) [Li and Vitányi, 1997]. We combine both ideas and get *a parameterless model AIξ of an acting agent which we claim to behave optimally in any computable environment* (e.g. prisoner or auction problems, poker, car driving). To get an effective solution, a modification AIξ^{tl}, superior to any other time t and space l bounded agent, is constructed. The computation time of AIξ^{tl} is of the order $t \cdot 2^l$. The main goal of this work is to derive and discuss the AIξ and the AIξ^{tl} model, and to clarify the meaning of *universal*, *optimal*, *superior*, *etc.* Details can be found in [Hutter, 2000b].

[1] With a formal solution we mean a rigorous mathematically definition, uniquely specifying the solution. For problems considered here this always implies the existence of an algorithm which asymptotically converges to the correct solution.

2 Rational Agents & Sequential Decisions

Agents in probabilistic environments: A very general framework for intelligent systems is that of rational agents [Russell and Norvig, 1995]. In cycle k, an agent performs *action* $y_k \in Y$ (output word) which results in a *perception* $x_k \in X$ (input word), followed by cycle $k+1$ and so on. If agent and environment are deterministic and computable, the entanglement of both can be modeled by two Turing machines with two common tapes (and some private tapes) containing the action stream $y_1 y_2 y_3$... and the perception stream $x_1 x_2 x_3$... (The meaning of $x_k \equiv x'_k r_k$ is explained in the next paragraph):

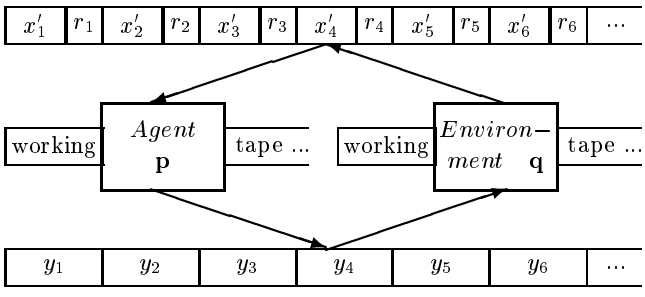

p is the *policy* of the agent interacting with environment q. We write $p(x_{<k}) = y_{1:k}$ to denote the output $y_{1:k} \equiv y_1 ... y_k$ of the agent p on input $x_{<k} \equiv x_1 ... x_{k-1}$ and similarly $q(y_{1:k}) = x_{1:k}$ for the environment q. We call Turing machines p and q behaving in this way *chronological*. In the more general case of a *probabilistic environment*, given the history $\dot{y}\dot{x}_{<k}y_k \equiv y_1 x_1 ... y_{k-1} x_{k-1} y_k$, the probability that the environment leads to perception x_k in cycle k is (by definition) $\mu(\dot{y}\dot{x}_{<k}\underline{y}\underline{x}_k)$. The underlined argument $\underline{x}_k$ in μ is a probability variable and the other non-underlined arguments $\dot{y}\dot{x}_{<k}y_k$ represent conditions. We call probability distributions like μ *chronological*.

The AIμ Model: The goal of the agent is to maximize future *rewards*, which are provided by the environment through the inputs x_k. The inputs $x_k \equiv x'_k r_k$ are divided into a regular part x'_k and some (possibly empty or delayed) reward r_k. The μ-expected reward sum of future cycles k to m with outputs $y_{k:m} = y^p_{k:m}$ generated by the agent's policy p can be written compactly as

$$V^p_\mu(\dot{y}\dot{x}_{<k}) := \sum_{x_k ... x_m} (r_k + ... + r_m)\mu(\dot{y}\dot{x}_{<k}\underline{y}\underline{x}_{k:m}), \quad (1)$$

where m is the *lifespan* of the agent, and the dots above $\dot{y}\dot{x}_{<k}$ indicate the actual action and perception history. The μ-expected reward sum of future cycles k to m with outputs y_i generated by the *ideal agent*, which maximizes the expected future rewards is

$$V^*_\mu(\dot{y}\dot{x}_{<k}) := \max_{y_k}\sum_{x_k} ... \max_{y_m}\sum_{x_m}(r_k + ... + r_m)\mu(\dot{y}\dot{x}_{<k}\underline{y}\underline{x}_{k:m}), \quad (2)$$

i.e. the best expected credit is obtained by averaging over the x_i and maximizing over the y_i. This has to be done in chronological order to correctly incorporate the dependency of x_i and y_i on the history. The output $\dot{y}_k$, which achieves the maximal value defines *the AIμ model*:

$$\dot{y}_k := \max_{y_k}\arg\sum_{x_k} ... \max_{y_m}\sum_{x_m}(r_k + ... + r_m)\mu(\dot{y}\dot{x}_{<k}\underline{y}\underline{x}_{k:m}). \quad (3)$$

The AIμ model is optimal in the sense that no other policy leads to higher μ-expected reward. A detailed derivation and other recursive and functional versions can be found in [Hutter, 2000b].

Sequential decision theory: Eq. (3) is essentially an Expectimax algorithm/sequence. One can relate (3) to the Bellman equations [Bellman, 1957] of sequential decision theory by identifying complete histories $\dot{y}\dot{x}_{<k}$ with states, $\mu(\dot{y}\dot{x}_{<k}\underline{y}\underline{x}_k)$ with the state transition matrix, $V^*_\mu(\dot{y}\dot{x}_{<k})$ with the value of history/state $\dot{y}\dot{x}_{<k}$, and y_k with the action in cycle k [Russell and Norvig, 1995; Hutter, 2000b]. Due to the use of complete histories as state space, the AIμ model neither assumes stationarity, nor the Markov property, nor complete accessibility of the environment. Every state occurs at most once in the lifetime of the system. As we have in mind a universal system with complex interactions, the action and perception spaces Y and X are huge (e.g. video images), and every action or perception itself occurs usually only once in the lifespan m of the agent. As there is no (obvious) universal similarity relation on the state space, an effective reduction of its size is impossible, but there is no principle problem in determining $\dot{y}_k$ as long as μ is known and computable and X, Y and m are finite.

Reinforcement learning: Things dramatically change if μ is unknown. Reinforcement learning algorithms [L.P. Kaelbling, 1996; Sutton and Barto, 1998; Bertsekas and Tsitsiklis, 1996] are commonly used in this case to learn the unknown μ. They succeed if the state space is either small or has effectively been made small by generalization or function approximation techniques. In any case, the solutions are either ad hoc, work in restricted domains only, have serious problems with state space exploration versus exploitation, or have non-optimal learning rate. There is no universal and optimal solution to this problem so far. In the Section 4 we present a new model and argue that it formally solves all these problems in an optimal way. The true probability distribution μ will not be learned directly, but will be replaced by a universal prior ξ, which is shown to converge to μ in a sense.

3 Algorithmic Complexity and Universal Induction

The problem of the unknown environment: We have argued that currently there is no universal and optimal solution to solving reinforcement learning problems. On the other hand, [Solomonoff, 1964] defined a universal scheme of inductive inference, based on Epicurus' principle of multiple explanations, Ockham's razor, and Bayes' rule for conditional probabilities. For an excellent introduction one should consult the book of [Li and Vitányi, 1997]. In the following we outline the theory and the basic results.

Kolmogorov complexity and universal probability: Let us choose some universal prefix Turing machine U with unidirectional binary input and output tapes and a bidirectional working tape. We can then define the (conditional) prefix Kolmogorov complexity [Chaitin, 1975; Gács, 1974; Kolmogorov, 1965; Levin, 1974] as the length l of the shortest program p, for which U outputs the binary string $x = x_{1:n}$ with $x_i \in \{0,1\}$:

$$K(x) := \min_p\{l(p) : U(p) = x\},$$

and given y

$$K(x|y) := \min_p\{l(p) : U(p,y) = x\}.$$

The *universal semimeasure* $\xi(\underline{x})$ is defined as the probability that the output of U starts with x when provided with fair coin flips on the input tape [Solomonoff, 1964;

1978]. It is easy to see that this is equivalent to the formal definition

$$\xi(\underline{x}) \;:=\; \sum_{p\,:\,\exists\omega:U(p)=x\omega} 2^{-l(p)} \qquad (4)$$

where the sum is over minimal programs p for which U outputs a string starting with x. U might be nonterminating. As the short programs dominate the sum, ξ is closely related to $K(x)$ as $\xi(\underline{x}) = 2^{-K(x)+O(K(l(x)))}$. ξ has the important universality property [Solomonoff, 1964] that it dominates every computable probability distribution ρ up to a multiplicative factor depending only on ρ but not on x:

$$\xi(\underline{x}) \;\geq\; 2^{-K(\rho)-O(1)}\cdot\rho(\underline{x}). \qquad (5)$$

The Kolmogorov complexity of a function like ρ is defined as the length of the shortest self-delimiting coding of a Turing machine computing this function. ξ itself is *not* a probability distribution[2]. We have $\xi(\underline{x}0)+\xi(\underline{x}1)<\xi(\underline{x})$ because there are programs p, which output just x, neither followed by 0 nor 1. They just stop after printing x or continue forever without any further output. We will call a function $\rho\geq 0$ with the properties $\rho(\epsilon)\leq 1$ and $\sum_{x_n}\rho(\underline{x}_{1:n})\leq\rho(\underline{x}_{<n})$ a *semimeasure*. ξ is a semimeasure and (5) actually holds for all enumerable semimeasures ρ.

Universal sequence prediction: (Binary) sequence prediction algorithms try to predict the continuation x_n of a given sequence $x_1...x_{n-1}$. In the following we will assume that the sequences are drawn from a probability distribution and that the true probability of a string starting with $x_1...x_n$ is $\mu(\underline{x}_{1:n})$. The probability of x_n given $x_{<n}$ hence is $\mu(x_{<n}\underline{x}_n)$. If we measure prediction quality as the number of correct predictions, the best possible system predicts the x_n with the highest probability. Usually μ is unknown and the system can only have some belief ρ about the true distribution μ. Now the universal probability ξ comes into play: [Solomonoff, 1978] has proved that the mean squared difference between ξ and μ is finite for computable μ:

$$\sum_{k=1}^{\infty}\sum_{x_{1:k}}\mu(\underline{x}_{<k})(\xi(x_{<k}\underline{x}_k)-\mu(x_{<k}\underline{x}_k))^2 \qquad (6)$$
$$<\; \ln 2\cdot K(\mu)+O(1).$$

A simplified proof can be found in [Hutter, 1999]. So the difference between $\xi(x_{<n}\underline{x}_n)$ and $\mu(x_{<n}\underline{x}_n)$ tends to zero with μ probability 1 for *any* computable probability distribution μ. The reason for the astonishing property of a single (universal) function to converge to *any* computable probability distribution lies in the fact that the set of μ-random sequences differ for different μ. The universality property (5) is the central ingredient for proving (6).

[2]It is possible to normalize ξ to a probability distribution as has been done in [Solomonoff, 1978; Hutter, 1999] by giving up the enumerability of ξ. Bounds (6) and (8) hold for both definitions.

Error bounds: Let SPρ be a probabilistic sequence predictor, predicting x_n with probability $\rho(x_{<n}\underline{x}_n)$. If ρ is only a semimeasure the SPρ system might refuse any output in some cycles n. Further, we define a deterministic sequence predictor SPΘ_ρ predicting the x_n with highest ρ probability. $\Theta_\rho(x_{<n}\underline{x}_n):=1$ for one x_n with $\rho(x_{<n}\underline{x}_n)\geq\rho(x_{<n}\underline{x}'_n)\,\forall x'_n$ and $\Theta_\rho(x_{<n}\underline{x}_n):=0$ otherwise. SPΘ_μ is the best prediction scheme when μ is known. If $\rho(x_{<n}\underline{x}_n)$ converges quickly to $\mu(x_{<n}\underline{x}_n)$ the number of additional prediction errors introduced by using Θ_ρ instead of Θ_μ for prediction should be small in some sense. Let us define the total number of expected erroneous predictions the SPρ system makes for the first n bits:

$$E_{n\rho} \;:=\; \sum_{k=1}^{n}\sum_{x_{1:k}}\mu(\underline{x}_{1:k})(1-\rho(x_{<k}\underline{x}_k)). \qquad (7)$$

The SPΘ_μ system is best in the sense that $E_{n\Theta_\mu}\leq E_{n\rho}$ for any ρ. In [Hutter, 1999] it has been shown that SPΘ_ξ is not much worse

$$E_{n\Theta_\xi}-E_{n\rho} \;\leq\; H+\sqrt{4E_{n\rho}H+H^2} \;=\; O(\sqrt{E_{n\rho}}) \qquad (8)$$
$$\text{with}\quad H \;<\; \ln 2\cdot K(\mu)+O(1)$$

and the tightest bound for $\rho=\Theta_\mu$. For finite $E_{\infty\Theta_\mu}$, $E_{\infty\Theta_\xi}$ is finite too. For infinite $E_{\infty\Theta_\mu}$, $E_{n\Theta_\xi}/E_{n\Theta_\mu}\stackrel{n\to\infty}{\longrightarrow}1$ with rapid convergence. One can hardly imagine any better prediction algorithm as SPΘ_ξ without extra knowledge about the environment. In [Hutter, 2000a], (6) and (8) have been generalized from binary to arbitrary alphabet and to general loss functions. Apart from computational aspects, which are of course very important, the problem of sequence prediction could be viewed as essentially solved.

4 The Universal AIξ Model

Definition of the AIξ Model: We have developed enough formalism to suggest our universal AIξ model. All we have to do is to suitably generalize the universal semimeasure ξ from the last section and to replace the true but unknown probability μ in the AIμ model by this generalized ξ. In what sense this AIξ model is universal and optimal will be discussed thereafter.

We define the generalized universal probability ξ^{AI} as the $2^{-l(q)}$ weighted sum over all chronological programs (environments) q which output $x_{1:k}$, similar to (4) but with $y_{1:k}$ provided on the "input" tape:

$$\xi(\underline{yx}_{1:k}) \;:=\; \sum_{q:q(y_{1:k})=x_{1:k}} 2^{-l(q)}. \qquad (9)$$

Replacing μ by ξ in (3) the iterative AIξ system outputs

$$\dot{y}_k := \operatorname*{maxarg}_{y_k}\sum_{x_k}...\max_{y_m}\sum_{x_m}(c_k+...+c_m)\xi(\dot{y}\dot{x}_{<k}\underline{yx}_{k:m}). \qquad (10)$$

in cycle k given the history $\dot{y}\dot{x}_{<k}$.

(Non)parameters of AIξ: The AIξ model and its behaviour is completely defined by (9) and (10). It (slightly) depends on the choice of the universal Turing machine. The AIξ model also depends on the choice of X and Y, but we do not expect any bias when the spaces are chosen sufficiently large and simple, e.g. all strings of length 2^{16}. Choosing $\mathbb{N}$ as word space would be ideal, but whether the maxima (or suprema) exist in this case, has to be shown beforehand. The only non-trivial dependence is on the horizon m. Ideally we would like to chose $m=\infty$, but there are several subtleties to be discussed later, which prevent at least a naive limit $m\to\infty$. So apart from m and unimportant details, the AIξ system is uniquely defined by (10) and (9) without adjustable parameters. It does not depend on any assumption about the environment apart from being generated by some computable (but unknown!) probability distribution as we will see.

ξ is only a semimeasure: One subtlety should be mentioned. Like in the SP case, ξ is not a probability distribution but still satisfies the weaker inequalities

$$\sum_{x_n}\xi(\dot{y}\underline{x}_{1:n}) \;\leq\; \xi(\dot{y}\underline{x}_{<n}) \quad , \quad \xi(\epsilon) \;\leq\; 1 \qquad (11)$$

Note, that the sum on the l.h.s. is *not* independent of y_n unlike for the chronological probability distribution μ. Nevertheless, it is bounded by something (the r.h.s) which is independent of y_n. The reason is that the sum in (9) runs over (partial recursive) chronological functions only and the functions q which satisfy $q(y_{1:n})=x_{<n}x'_n$ for some $x'_n\in X$ are a subset of the functions satisfying $q(y_{<n})=x_{<n}$. We will in general call functions satisfying (11) *chronological semimeasures*. The important point is that the conditional probabilities (9) are ≤ 1 like for true probability distributions.

Universality of ξ^{AI}: It can be shown that ξ^{AI} defined in (9) is universal and converges to μ^{AI} analogously to the SP case (5) and (6). The proofs are generalizations from the SP case. The actions y are pure spectators and cause no difficulties in the generalization. This will change when we analyze error/value bounds analogously to (8). The major difference when incorporating y is that in (5), $U(p)=x\omega$ produces strings starting with x, whereas in (9) we can demand q to output exactly n words $x_{1:n}$ as q knows n from the number of input words $y_1...y_n$. ξ^{AI} dominates all *chronological enumerable semimeasures*

$$\xi(\dot{y}\underline{x}_{1:n}) \;\geq\; 2^{-K(\rho)-O(1)}\rho(\dot{y}\underline{x}_{1:n}). \qquad (12)$$

ξ is a universal element in the sense of (12) in the set of all enumerable chronological semimeasures. This can be proved even for infinite (countable) alphabet [Hutter, 2000b].

Convergence of ξ^{AI} to μ^{AI}: From (12) one can show

$$\sum_{k=1}^{n}\sum_{x_{1:k}}\mu(\dot{y}\underline{x}_{<k})\Big(\mu(\dot{y}x_{<k}\dot{y}\underline{x}_k)-\xi(\dot{y}x_{<k}\dot{y}\underline{x}_k)\Big)^2 \;<\; \ln 2\cdot K(\mu)+O(1) \qquad (13)$$

for computable chronological measures μ. The main complication in generalizing (6) to (13) is the generalization to non-binary alphabet [Hutter, 2000a]. The y are, again, pure spectators. (13) shows that the μ-expected squared difference of μ and ξ is finite for computable μ. This, in turn, shows that $\xi(\dot{y}x_{<k}\dot{y}\underline{x}_k)$ converges to $\mu(\dot{y}x_{<k}\dot{y}\underline{x}_k)$ for $k\to\infty$ with μ probability 1. If we take a finite product of ξ's and use Bayes' rule, we see that also $\xi(\dot{y}x_{<k}\dot{y}\underline{x}_{k:k+r})$ converges to $\mu(\dot{y}x_{<k}\dot{y}\underline{x}_{k:k+r})$. More generally, in case of a bounded horizon $h_k\equiv m_k-k+1\leq h_{max}<\infty$, it follows that

$$\xi(\dot{y}x_{<k}\dot{y}\underline{x}_{k:m_k}) \stackrel{k\to\infty}{\longrightarrow} \mu(\dot{y}x_{<k}\dot{y}\underline{x}_{k:m_k}) \qquad (14)$$

Convergence is only guaranteed for one (e.g. the true) i/o sequence $\dot{y}\dot{x}_{<k}\dot{y}\dot{x}_{k:m_k}$ but not for alternate sequences $\dot{y}\dot{x}_{<k}\dot{y}x_{k:m_k}$. Since (10) takes an average over all possible future actions and perceptions $\dot{y}x_{k:m_k}$; not only the one which will finally occur, (14) does not guarantee $\dot{y}^{\xi}_k\to\dot{y}^{\mu}_k$. This gap is already present in the SPΘ_ρ models, but nevertheless good error bounds could be proved. This gives confidence that the outputs $\dot{y}_k$ of the AIξ model (10) could converge to the outputs $\dot{y}_k$ of the AIμ model (3), at least for a bounded horizon h_k. The problems with a fixed horizon $m_k=m$ and especially $m\to\infty$ will be discussed later.

Universally optimal AI systems: We want to call an AI model *universal*, if it is μ-independent (unbiased, model-free) and is able to solve any solvable problem and learn any learnable task. Further, we call a universal model, *universally optimal*, if there is no program, which can solve or learn significantly faster (in terms of interaction cycles). As the AIξ model is parameterless, ξ converges to μ in the sense of (13,14), the AIμ model is itself optimal, and we expect no other model to converge faster to AIμ by analogy to SP (8),

we expect AIξ to be universally optimal.

This is our main claim. Further support is given in [Hutter, 2000b] by a detailed analysis of the behaviour of AIξ for various problem classes, including prediction, optimization, games, and supervised learning.

The choice of the horizon: The only significant arbitrariness in the AIξ model lies in the choice of the lifespan m or the $h_k\equiv m_k-k+1$ if we allow a cycle dependent m. We will not discuss ad hoc choices of h_k for specific problems. We are interested in universal choices. The book of [Bertsekas, 1995] thoroughly discusses the mathematical problems regarding infinite horizon systems.

In many cases the time we are willing to run a system depends on the quality of its actions. Hence, the lifetime,

if finite at all, is not known in advance. Exponential discounting $r_k \to r_k \cdot \gamma^k$ solves the mathematical problem of $m \to \infty$ but is no real solution, since an effective horizon $h \sim \ln \frac{1}{\gamma}$ has been introduced. The scale invariant discounting $r_k \to r_k \cdot k^{-\alpha}$ has a dynamic horizon $h \sim k$. This choice has some appeal, as it seems that humans of age k years usually do not plan their lives for more than the next $\sim k$ years. From a practical point of view this model might serve all needs, but from a theoretical point we feel uncomfortable with such a limitation in the horizon from the very beginning. A possible way of taking the limit $m \to \infty$ without discounting and its problems can be found in [Hutter, 2000b].

Another objection against too large choices of m_k is that $\xi(yx_{<k}\underline{yx}_{k:m_k})$ has been proved to be a good approximation of $\mu(yx_{<k}\underline{yx}_{k:m_k})$ only for $k \gg h_k$, which is never satisfied for $m_k = m \to \infty$. On the other hand it may turn out that the rewards $r_{k'}$ for $k' \gg k$, where ξ may no longer be trusted as a good approximation of μ, are in a sense randomly disturbed with decreasing influence on the choice of $\dot{y}_k$. This claim is supported by the forgetfulness property of ξ (see next section)and can be proved when restricting to factorizable environments [Hutter, 2000b].

We are not sure whether the choice of m_k is of marginal importance, as long as m_k is chosen sufficiently large and of low complexity, $m_k = 2^{2^{16}}$ for instance, or whether the choice of m_k will turn out to be a central topic for the AIξ model or for the planning aspect of any universal AI system in general. Most if not all problems in agent design of balancing exploration and exploitation vanish by a sufficiently large choice of the (effective) horizon and/or a sufficiently general prior. We suppose that the limit $m_k \to \infty$ for the AIξ model results in correct behaviour for weakly separable (defined in the next section) μ, and that even the naive limit $m \to \infty$ may exist.

Value bounds and separability concepts: The values V_ρ^* associated with the AIρ systems correspond roughly to the negative error measure $-E_{n\rho}$ of the SPρ systems. In the SP case we were interested in small bounds for the error excess $E_{n\Theta_\xi} - E_{n\rho}$. Unfortunately, simple value bounds for AIξ or any other AI system in terms of V^* analogously to the error bound (8) can not hold [Hutter, 2000b]. We even have difficulties in specifying what we can expect to hold for AIξ or any AI system which claims to be universally optimal. In SP, the only important property of μ for proving error bounds was its complexity $K(\mu)$. In the AI case, there are no useful bounds in terms of $K(\mu)$ only. We either have to study restricted problem classes or consider bounds depending on other properties of μ, rather than on its complexity only. In [Hutter, 2000b] the difficulties are exhibited by two examples. Several concepts, which might be useful for proving value bounds are introduced and discussed. They include forgetful, relevant, asymptotically learnable, farsighted, uniform, (generalized) Markovian, factorizable and (pseudo) passive μ. They are approximately sorted in the order of decreasing generality and are called *separability concepts*. A first weak bound for passive μ is proved.

5 Time Bounds and Effectiveness

Non-effectiveness of AIξ: ξ is not a computable but only an enumerable semimeasure. Hence, the output $\dot{y}_k$ of the AIξ model is only asymptotically computable. AIξ yields an algorithm that produces a sequence of trial outputs eventually converging to the correct output $\dot{y}_k$, but one can never be sure whether one has already reached it. Besides this, convergence is extremely slow, so this type of asymptotic computability is of no direct (practical) use. Furthermore, the replacement of ξ by time-limited versions [Li and Vitányi, 1991; 1997], which is suitable for sequence prediction, has been shown to fail for the AIξ model [Hutter, 2000b]. This leads to the issues addressed next.

Time bounds and effectiveness: Let $\tilde{p}$ be a policy which calculates an acceptable output within a reasonable time $\tilde{t}$ per cycle. This sort of computability assumption, namely, that a general purpose computer of sufficient power and appropriate program is able to behave in an intelligent way, is the very basis of AI research. Here it is not necessary to discuss what exactly is meant by 'reasonable time/intelligence' and 'sufficient power'. What we are interested in is whether there is a computable version AI$\xi^{\bar{t}}$ of the AIξ system which is superior or equal to any program p with computation time per cycle of at most $\tilde{t}$.

What one can realistically hope to construct is an AI$\xi^{\bar{t}\bar{l}}$ system of computation time $c \cdot \tilde{t}$ per cycle for some constant c. The idea is to run all programs p of length $\leq \tilde{l} := l(\tilde{p})$ and time $\leq \tilde{t}$ per cycle and pick the best output in the sense of maximizing the *universal value* V_ξ^*. The total computation time is $c \cdot \tilde{t}$ with $c \approx 2^{\bar{l}}$. Unfortunately V_ξ^* can not be used directly since this measure is also only semi-computable and the approximation quality by using computable versions of ξ given a time of order $c \cdot \tilde{t}$ is crude [Li and Vitányi, 1997; Hutter, 2000b]. On the other hand, we *have* to use a measure which converges V_ξ^* for $\tilde{t}, \tilde{l} \to \infty$, since the AI$\xi^{\bar{t}\bar{l}}$ model should converge to the AIξ model in that case.

Valid approximations: A solution satisfying the above conditions is suggested in [Hutter, 2000b]. The main idea is to consider *extended chronological incremental policies* p, which in addition to the regular output y_k^p *rate* their own output with w_k^p. The AI$\xi^{\bar{t}\bar{l}}$ model selects the output $\dot{y}_k = y_k^p$ of the policy p with highest rating w_k^p. p might suggest any output y_k^p but it is not allowed to rate itself with an arbitrarily high w_k^p if one wants w_k^p to be a reliable criterion for selecting the best p. One must demand that no policy p is allowed to claim that it is better than it actually is. In [Hutter, 2000b]

a (logical) predicate VA(p), called *valid approximation*, is defined, which is true if, and only if, p *always* satisfies $w_k^p \leq V_\xi^p(\dot{y}\dot{x}_{<k})$, i.e. never overrates itself. $V_\xi^p(\dot{y}\dot{x}_{<k})$ is the ξ expected future reward under policy p. Valid policies p can then be (partially) ordered w.r.t. their rating w_k^p.

The universal time bounded AI$\xi^{\tilde{t}\tilde{l}}$ system: In the following, we describe the algorithm p^* underlying the universal time bounded AI$\xi^{\tilde{t}\tilde{l}}$ system. It is essentially based on the selection of the best algorithms p_k^* out of the time $\tilde{t}$ and length $\tilde{l}$ bounded policies p, for which there exists a proof P of VA(p) with length $\leq l_P$.

1. Create all binary strings of length l_P and interpret each as a coding of a mathematical proof in the same formal logic system in which VA($\cdot$) has been formulated. Take those strings which are proofs of VA(p) for some p and keep the corresponding programs p.
2. Eliminate all p of length $>\tilde{l}$.
3. Modify all p in the following way: all output $w_k^p y_k^p$ is temporarily written on an auxiliary tape. If p stops in $\tilde{t}$ steps the internal 'output' is copied to the output tape. If p does not stop after $\tilde{t}$ steps a stop is forced and $w_k^p = 0$ and some arbitrary y_k^p is written on the output tape. Let $\mathcal{P}$ be the set of all those modified programs.
4. Start first cycle: $k := 1$.
5. Run every $p \in \mathcal{P}$ on extended input $\dot{y}\dot{x}_{<k}$, where all outputs are redirected to some auxiliary tape: $p(\dot{y}\dot{x}_{<k}) = w_1^p y_1^p ... w_k^p y_k^p$. This step is performed incrementally by adding $\dot{y}\dot{x}_{k-1}$ for $k > 1$ to the input tape and continuing the computation of the previous cycle.
6. Select the program p with highest rating w_k^p: $p_k^* := \text{maxarg}_p\, w_k^p$.
7. Write $\dot{y}_k := y_k^{p_k^*}$ to the output tape.
8. Receive input $\dot{x}_k$ from the environment.
9. Begin next cycle: $k := k+1$, goto step 5.

Properties of the p^* algorithm: Let p be any extended chronological (incremental) policy of length $l(p) \leq \tilde{l}$ and computation time per cycle $t(p) \leq \tilde{t}$, for which there exists a proof of VA(p) of length $\leq l_P$. The algorithm p^*, depending on $\tilde{l}$, $\tilde{t}$ and l_P but not on p, has always higher rating than any such p. The setup time of p^* is $t_{setup}(p^*) = O(l_P^2 \cdot 2^{l_P})$ and the computation time per cycle is $t_{cycle}(p^*) = O(2^{\tilde{l}} \cdot \tilde{t})$. Furthermore, for $\tilde{t}, \tilde{l} \to \infty$, p^* converges to the behavior of the AIξ model.

Roughly speaking, this means that if there exists a computable solution to some AI problem at all, then the explicitly constructed algorithm p^* is such a solution. Although this claim is quite general, there are some limitations and open questions, regarding the setup time regarding the necessity that the policies must rate their own output, regarding true but not efficiently provable VA(p), and regarding "inconsistent" policies [Hutter, 2000b].

6 Outlook & Discussion

This section contains some discussion and remarks on otherwise unmentioned topics.

Value bounds: Rigorous proofs of value bounds for the AIξ theory are the major theoretical challenge – general ones as well as tighter bounds for special environments μ. Of special importance are suitable (and acceptable) conditions to μ, under which $\dot{y}_k$ and finite value bounds exist for infinite Y, X and m.

Scaling AIξ down: [Hutter, 2000b] shows for several examples how to integrate problem classes into the AIξ model. Conversely, one can downscale the AIξ model by using more restricted forms of ξ. This could be done in a similar way as the theory of universal induction has been downscaled with many insights to the Minimum Description Length principle [Li and Vitányi, 1992; Rissanen, 1989] or to the domain of finite automata [Feder *et al.*, 1992]. The AIξ model might similarly serve as a super model or as the very definition of (universal unbiased) intelligence, from which specialized models could be derived.

Applications: [Hutter, 2000b] shows how a number of AI problem classes, including *sequence prediction*, *strategic games*, *function minimization* and *supervised learning* fit into the general AIξ model. All problems are claimed to be formally solved by the AIξ model. The solution is, however, only formal, because the AIξ model is uncomputable or, at best, approximable. First, each problem class is formulated in its natural way (when μ^{problem} is known) and then a formulation within the AIμ model is constructed and their equivalence is proven. Then, the consequences of replacing μ by ξ are considered. The main goal is to understand how the problems are solved by AIξ. For more details see [Hutter, 2000b].

Implementation and approximation: The AI$\xi^{\tilde{t}\tilde{l}}$ model suffers from the same large factor $2^{\tilde{l}}$ in computation time as Levin search for inversion problems [Levin, 1973; 1984]. Nevertheless, Levin search has been implemented and successfully applied to a variety of problems [Schmidhuber, 1997; Schmidhuber *et al.*, 1997]. Hence, a direct implementation of the AI$\xi^{\tilde{t}\tilde{l}}$ model may also be successful, at least in toy environments, e.g. prisoner problems. The AI$\xi^{\tilde{t}\tilde{l}}$ algorithm should be regarded only as the first step toward a *computable universal AI model*. Elimination of the factor $2^{\tilde{l}}$ without giving up universality will probably be a very difficult task. One could try to select programs p and prove VA(p) in a more clever way than by mere enumeration. All kinds of ideas like, heuristic search, genetic algorithms, advanced theorem

provers, and many more could be incorporated. But now we have a problem.

Computability: We seem to have transferred the AI problem just to a different level. This shift has some advantages (and also some disadvantages) but presents, in no way, a solution. Nevertheless, we want to stress that we have reduced the AI problem to (mere) computational questions. Even the most general other systems the author is aware of, depend on some (more than complexity) assumptions about the environment, or it is far from clear whether they are, indeed, universally optimal. Although computational questions are themselves highly complicated, this reduction is a non-trivial result. A formal theory of something, even if not computable, is often a great step toward solving a problem and has also merits of its own (see previous paragraphs).

Elegance: Many researchers in AI believe that intelligence is something complicated and cannot be condensed into a few formulas. They believe it is more a combining of enough *methods* and much explicit *knowledge* in the right way. From a theoretical point of view, we disagree as the AIξ model is simple and seems to serve all needs. From a practical point of view we agree to the following extent. To reduce the computational burden one should provide special purpose algorithms (*methods*) from the very beginning, probably many of them related to reduce the complexity of the input and output spaces X and Y by appropriate pre/post-processing methods.

Extra knowledge: There is no need to incorporate extra *knowledge* from the very beginning. It can be presented in the first few cycles in *any* format. As long as the algorithm that interprets the data is of size $O(1)$, the AIξ system will 'understand' the data after a few cycles (see [Hutter, 2000b]). If the environment μ is complicated but extra knowledge z makes $K(\mu|z)$ small, one can show that the bound (13) reduces to $\ln 2 \cdot K(\mu|z)$ when $x_1 \equiv z$, i.e. when z is presented in the first cycle. Special purpose algorithms could also be presented in x_1, but it would be cheating to say that no special purpose algorithms have been implemented in AIξ. The boundary between implementation and training is blurred in the AIξ model.

Training: We have not said much about the training process itself, as it is not specific to the AIξ model and has been discussed in literature in various forms and disciplines. A serious discussion would be out of place. To repeat a truism, it is, of course, important to present enough knowledge x'_k and evaluate the system output y_k with r_k in a reasonable way. To maximize the information content in the reward, one should start with simple tasks and give positive reward to approximately the better half of the outputs y_k, for instance.

The big questions: [Hutter, 2000b] contains a discussion of the "big" questions concerning the mere existence of any computable, fast, and elegant universal theory of intelligence, related to non-computable μ [Penrose, 1994] and the 'number of wisdom' Ω [Chaitin, 1975; 1991].

References

[Bellman, 1957] R. Bellman. *Dynamic Programming.* Princeton University Press, New Jersey, 1957.

[Bertsekas and Tsitsiklis, 1996] D. P. Bertsekas and J. N. Tsitsiklis. *Neuro-Dynamic Programming.* Athena Scientific, Belmont, MA, 1996.

[Bertsekas, 1995] D. P. Bertsekas. *Dynamic Programming and Optimal Control, Vol. (II).* Athena Scientific, Belmont, Massachusetts, 1995.

[Chaitin, 1975] G. J. Chaitin. A theory of program size formally identical to information theory. *Journal of the ACM*, 22(3):329–340, 1975.

[Chaitin, 1991] G. J. Chaitin. Algorithmic information and evolution. *in O.T. Solbrig and G. Nicolis, Perspectives on Biological Complexity, IUBS Press*, pages 51–60, 1991.

[Feder *et al.*, 1992] M. Feder, N. Merhav, and M. Gutman. Universal prediction of individual sequences. *IEEE Transactions on Information Theory*, 38:1258–1270, 1992.

[Gács, 1974] P. Gács. On the symmetry of algorithmic information. *Russian Academy of Sciences Doklady. Mathematics (formerly Soviet Mathematics–Doklady)*, 15:1477–1480, 1974.

[Hutter, 1999] M. Hutter. New error bounds for Solomonoff prediction. *Journal of Computer and System Science, in press*, (IDSIA-11-00):1–13, 1999. ftp://ftp.idsia.ch/pub/techrep/IDSIA-11-00.ps.gz.

[Hutter, 2000a] M. Hutter. Optimality of universal prediction for general loss and alphabet. Technical Report IDSIA-15-00, Istituto Dalle Molle di Studi sull'Intelligenza Artificiale, Manno(Lugano), Switzerland, 2000. In progress.

[Hutter, 2000b] M. Hutter. A theory of universal artificial intelligence based on algorithmic complexity. Technical report, 2000. 62 pages, http://xxx.lanl.gov/abs/cs.AI/0004001.

[Kolmogorov, 1965] A. N. Kolmogorov. Three approaches to the quantitative definition of information. *Problems of Information and Transmission*, 1(1):1–7, 1965.

[Levin, 1973] L. A. Levin. Universal sequential search problems. *Problems of Information Transmission*, 9:265–266, 1973.

[Levin, 1974] L. A. Levin. Laws of information conservation (non-growth) and aspects of the foundation of probability theory. *Problems of Information Transmission*, 10:206–210, 1974.

[Levin, 1984] L. A. Levin. Randomness conservation inequalities: Information and independence in mathematical theories. *Information and Control*, 61:15–37, 1984.

[Li and Vitányi, 1991] M. Li and P. M. B. Vitányi. Learning simple concepts under simple distributions. *SIAM Journal on Computing*, 20(5):911–935, 1991.

[Li and Vitányi, 1992] M. Li and P. M. B. Vitányi. Inductive reasoning and Kolmogorov complexity. *Journal of Computer and System Sciences*, 44:343–384, 1992.

[Li and Vitányi, 1997] M. Li and P. M. B. Vitányi. *An introduction to Kolmogorov complexity and its applications*. Springer, 2nd edition, 1997.

[L.P. Kaelbling, 1996] A.W. Moore L.P. Kaelbling, M.L. Littman. Reinforcement learning: a survey. *Journal of AI research*, 4:237–285, 1996.

[Penrose, 1994] R. Penrose. *Shadows of the mind, A search for the missing science of consciousness*. Oxford Univ. Press, 1994.

[Rissanen, 1989] J. Rissanen. *Stochastic Complexity in Statistical Inquiry*. World Scientific Publ. Co., 1989.

[Russell and Norvig, 1995] S. J. Russell and P. Norvig. *Artificial Intelligence. A Modern Approach*. Prentice-Hall, Englewood Cliffs, 1995.

[Schmidhuber *et al.*, 1997] J. Schmidhuber, J. Zhao, and M. Wiering. Shifting inductive bias with success-story algorithm, adaptive Levin search, and incremental self-improvement. *Machine Learning*, 28:105–130, 1997.

[Schmidhuber, 1997] J. Schmidhuber. Discovering neural nets with low Kolmogorov complexity and high generalization capability. *Neural Networks*, 10(5):857–873, 1997.

[Solomonoff, 1964] R. J. Solomonoff. A formal theory of inductive inference: Part 1 and 2. *Inform. Control*, 7:1–22, 224–254, 1964.

[Solomonoff, 1978] R. J. Solomonoff. Complexity-based induction systems: comparisons and convergence theorems. *IEEE Trans. Inform. Theory*, IT-24:422–432, 1978.

[Sutton and Barto, 1998] R. Sutton and A. Barto. *Reinforcement learning: An introduction*. Cambridge, MA, MIT Press, 1998.