

Document downloaded from:

<http://hdl.handle.net/10251/35938>

This paper must be cited as:

Hernández Orallo, J.; Dowe, DL.; España Cubillo, S.; Hernández-Lloreda, MV.; Insa Cabrera, J. (2011). On more realistic environment distributions for defining, evaluating and developing intelligence. En Artificial General Intelligence. Springer Verlag (Germany). 6830:82-91. doi:10.1007/978-3-642-22887-2\_9.



The final publication is available at

[http://link.springer.com/chapter/10.1007/978-3-642-22887-2\\_9](http://link.springer.com/chapter/10.1007/978-3-642-22887-2_9)

Copyright Springer Verlag (Germany)

# On more realistic environment distributions for defining, evaluating and developing intelligence

José Hernández-Orallo<sup>1</sup>   David L. Dowe<sup>2</sup>   Sergio España-Cubillo<sup>3</sup>  
M.Victoria Hernández-Lloreda<sup>4</sup>   Javier Insa-Cabrera<sup>1</sup>

<sup>1</sup> DSIC, Universitat Politècnica de València, Spain. {jorallo, jinsa}@dsic.upv.es

<sup>2</sup> Clayton School of Information Technology, Monash University, Australia.

david.dowe@monash.edu

<sup>3</sup> ProS Research Center, Universitat Politècnica de València, Spain.

sergio.espana@pros.upv.es

<sup>4</sup> Departamento de Metodología de las Ciencias del Comportamiento, Universidad Complutense de Madrid, Spain. vhlloreda@psi.ucm.es

**Abstract.** One insightful view of the notion of intelligence is the ability to perform well in a diverse set of tasks, problems or environments. One of the key issues is therefore the choice of this set and the probability of each individual, which can be formalised as a ‘distribution’. Formalising and properly defining this distribution is an important challenge to understand what intelligence is and to achieve artificial general intelligence (AGI). In this paper, we agree with previous criticisms that a universal distribution using a reference universal Turing machine (UTM) over tasks, environments, etc., is perhaps much too general, since, e.g., the probability of other agents appearing on the scene or having some social interaction is almost 0 for most reference UTMs. Instead, we propose the notion of Darwin-Wallace distribution for environments, which is inspired by biological evolution, artificial life and evolutionary computation. However, although enlightening about where and how intelligence should excel, this distribution has so many options and is uncomputable in so many ways that we certainly need a more practical alternative. We propose the use of intelligence tests over multi-agent systems, in such a way that agents with a certified level of intelligence at a certain degree are used to construct the tests for the next degree. This constructive methodology can then be used as a more realistic intelligence test and also as a testbed for developing and evaluating AGI systems.

**Keywords:** Intelligence, Evolutionary Computation, Artificial Life, Social Intelligence, Intelligence Test, Universal Distribution.

## 1 Introduction

Understanding what intelligence is (and is not) plays a crucial role in developing truly general intelligent machines. Apart from the many informal definitions from psychology, philosophy, biology, artificial intelligence and other disciplines (see an account in [16]), there have been some definitions which include the notion of compression, Kolmogorov Complexity or related concepts such as

Solomonoff’s universal distribution (see, e.g. [4, 11, 7, 17]). Some of these proposals have claimed that they are necessary and sufficient, while others have only claimed that the ability which is defined is necessary for intelligence but not sufficient (so suggesting further factors, see e.g. [8]).

Apart from the view of ‘intelligence as compression’ (or compression – one-part or (MML) two-part – being a necessary part of intelligence), the previous approaches are based on a distribution of tasks, exercises or environments. Then, intelligence is defined as good performance in this distribution of exercises. This clearly connects the notion of an ‘intelligence definition’ with the notion of an ‘intelligence test’, notions which are tightly intertwined. The kind of tasks (static vs. dynamic, predictive vs. explanatory, etc.) distinguishes the previous approaches. The common feature is the distribution which is used for the selection of tasks. While psychometrics does not choose tasks using a formal and independent distribution, these works take the tasks from a universal distribution [22], which gives higher probability to tasks whose shortest description is smaller.

Nonetheless, the use of a universal distribution gives a very high weight (i.e., probability) to very simple tasks. The problem is that it is not true that intelligent systems have a monotone behaviour on the complexity of a task. In many simple tasks, a very simple program can perform better than the smartest human. Thus, aggregating results using this distribution would assign higher intelligence to simple programs. One possible solution is to set a minimum complexity value [13], but this is clearly an arbitrary choice. An alternative option would be to define intelligence as the maximum complexity level where a system can score ‘significantly better’ than random. This can be shaped as an adaptive (or anytime) intelligence test [10], where the complexity of tasks (and speed) is adapted to the intelligence of the agent which is being evaluated.

However, even though we think that these ideas are in the right direction, there is a problem about the class of tasks we might generate from the universal distribution. Using environments like those typical in reinforcement learning, which is the approach taken in [17] and followed in [10], we see very clearly that even very complex environments will be very different to the environments a human (or an animal or a robot) faces during their life. We can, of course, enrich the environments to be more *physically* realistic, with simple physics, in order to allow the complex perception and mechanisms we are used to in this universe. However, apart from being anthropocentric, this 3D physical world would not ensure that other agents may appear, as it happens in, e.g., Mars.

If we generate environments randomly using a universal distribution over a Turing-complete reference machine (Universal Turing Machine, UTM), it is clear that some environments might contain some other agents. In fact, some environments might contain life, and some others might even contain intelligent life<sup>5</sup>. However, the probability of any of these environments is almost 0 for many

---

<sup>5</sup> An early exploration of the idea of computable universes where life can emerge is given in [20] and a discussion of how big a universe we need before intelligence might appear can be found in [1, sec. 0.2.7, p545, col. 1].

reference UTMs. So, we are quite far from what intelligence is supposed to be, an ability to interact with a physical world, full of plants and animals.

In fact, this is one of the drifts taken by artificial intelligence in the last decades. The appearance of multi-agent systems is a sign that the future of machine intelligence will not be found in monolithic systems solving tasks without other agents to compete or collaborate with. This is also better understood in comparative psychology and biology as well. For instance, [12] shows that cognitive abilities for dealing with the physical world (concerning things such as space, quantities, and causality) are similar for chimpanzees and human children, while it is in cognitive skills for dealing with the social world where the balance is in favour of humans. Furthermore, it is now better understood (see, e.g., [24]) that co-operation and communication are special traits of human cognition. In fact, “evolutionarily, the key difference is that humans have evolved not only social-cognitive skills geared toward competition, but also social-cognitive skills and motivations geared toward complex forms of co-operation.” [24]

So using a universal distribution to define or evaluate intelligence might be like evaluating or defining intelligent agents for Mars. A similar criticism has been raised by artificial (general) intelligence, stating that intelligence must be more embodied/natural/social [5]. In fact, there have been some proposals of simplified worlds, such as Goertzel’s AGI preschool [6], to test and develop AGI systems, whose complexity (in terms of shortest description) is very high (and hence, with very low probability for a universal distribution).

However, a non-principled way of selecting environments mimicking the environments that adult humans or small children interact with dramatically deviates from the goal of deriving a formal, mathematical and non-anthropomorphic definition of intelligence, and would be useless far beyond human intelligence.

In this paper, we take a different approach. Inspired by evolutionary computation, artificial life, multi-agent systems and social cognition, we develop a more realistic distribution of environments. The basic idea is straightforward: intelligence is the result of evolution through millions of generations interacting with other live beings. Thus we define intelligence in this context, interacting with other agents of similar intelligence.

We formalise the so-called Darwin-Wallace distribution for agents and environments. Despite the many options and the many sources of uncomputability, we claim that, conceptually, the notion of Darwin-Wallace distribution is useful to re-visit previous definitions of intelligence [4, 11, 7, 17, 8, 10].

The next step is how this notion can be used for AGI development and evaluation. We present a procedure which approximates a Darwin-Wallace distribution by using intelligence tests over environments such that ‘certified’ systems are incorporated into the environments, so making them socially more complex. This iterated process (in [19, sec. 5.1] a recursion is applied for the Turing Test) may lead to more realistic testbeds for AGI or for reinforcement learning.

The paper is organised as follows. The next section deals with the frequency of individuals in terms of evolution. This paves the way for the definition of the Darwin-Wallace distribution in section 3. Section 4 introduces a constructive

approximation of the distribution in the form of incremental intelligence tests from non-social environments to more social ones. Finally, section 5 discusses the contributions, implications and open questions which derive from this work.

## 2 Artificial Life, Biology and Intelligence

Many ideas from Artificial Life (Alife) [15] (and evolutionary computation) are useful for understanding which environments may hold intelligence. The starting point of every Alife project is a virtual (or artificial) environment where we place some individuals (a population) inside. One problem in Alife systems (and evolutionary computation) is that evolution stagnates at some point.

Of course, in Alife (and biology) we find the notion of ‘evolutionarily successful individual’. However, many species on Earth outnumber the *Homo Sapiens* (by several orders of magnitude), while their individual adaptability is poor (their adaptability as a species may be high, like cockroaches or some kinds of bacteria). The good thing about Alife is that we can force all the individuals to have the same body, while only letting their behaviour evolve.. So even though the notion of frequency in the real world is not related to the adaptability of a species, it *can* be related to the adaptability of a species or individual in a virtual world where the bodies and rules are the same for all the individuals.

Before this, we have to consider co-evolution and mind-reading. Co-evolution is the evolution of one species in a direction which is triggered or shaped by other species, as plants with insects, and mammals with intestinal bacteria. The genotype and phenotype (both its body and its behaviour) of many species can only be explained as the result of co-evolution. Intelligence, as a phenotypical trait, is not an exception. Intelligence can only be explained as the result of the co-evolution with other species and the co-evolution with individuals of the same species. In particular, when we focus on behaviour innovations (and not on physical changes), we find the concepts of mind-reading and manipulation [14], which can be applied for competition but also for co-operation, which are essential for predators and preys. Although ‘mind-reading’ has been used (especially in the past) to refer to both the adaptation of the species and the adaptation (i.e. intelligence) of the individual, in what follows we refer to the second. So, by ‘mind-reading’ and ‘manipulation’ we mean the use of information processing and the use of mental representation of other individuals, and not a genetically pre-programmed manipulation or simple pre-defined predator-prey patterns, as, e.g., an anteater can have. Human intelligence is the clearest example of such a social context. But again, evolutionarily, the key difference of humans in front of other animals (great apes included) is their ability at knowledge acquisition and transmission in order to perform better in very rich social contexts, “in terms of intention-reading, social learning, and communication” [24].

So, a more realistic definition of intelligence must give more weight to environments which contain other agents with social abilities.

### 3 Darwin-Wallace Distribution

Systems in Alife do not generally create environments where life emerges from scratch. On the contrary, they include a start-up population, from which things evolve. Environments are created in such a way that agents can do some actions and have a (generally partial) observation of the world. This setting is very similar to the way environments are created in reinforcement learning (RL), but the notion of reward is generally understood as a fitness function, which may affect the way agents die, reproduce or mutate.

So, following [10] (which in turns follows [17]), we define actions as a finite set of symbols  $\mathcal{A}$  (e.g.,  $\{left, right, up, down\}$ ); rewards are taken from any subset  $\mathcal{R}$  of rational numbers between  $-1$  and  $1$ ; and observations are also limited by a finite set  $\mathcal{O}$  of possibilities. We consider these sets to be the same for all the agents. This means that, physically, all the agents are equal.

An agent is an interactive (Turing) machine which receives a pair  $\langle r, o \rangle$  of reward and observation, known as a perception, and outputs an action  $a$ . Note that an agent is not a (Markov) function, so the action may depend on old rewards and observations. Additionally, an agent might also be stochastic. An agent can be properly formalised as a probability measure. For instance, given an agent, denoted by  $\pi^j$ , the term  $\pi^j(a_k^j | r_1^j o_1^j a_1^j r_2^j o_2^j a_2^j \dots r_k^j o_k^j)$  denotes the probability of agent  $\pi^j$  executing action  $a_k^j$  after the sequence of events  $r_1^j o_1^j a_1^j r_2^j o_2^j a_2^j \dots r_k^j o_k^j$ .

A *base* multi-agent environment is also an interactive (Turing) machine which takes as input the action of every agent and outputs rewards and observations for each and all of them. Again, it can be stochastic and it can be represented by a probability measure. For instance, for an environment accepting  $m$  agents, the term  $\mu(r_k^1 o_k^1 r_k^2 o_k^2 \dots r_k^m o_k^m | r_1^1 o_1^1 a_1^1 r_1^2 o_1^2 a_1^2 \dots r_1^m o_1^m a_1^m r_2^1 o_2^1 a_2^1 r_2^2 o_2^2 a_2^2 \dots r_2^m o_2^m a_2^m \dots r_{k-1}^1 o_{k-1}^1 a_{k-1}^1 r_{k-1}^2 o_{k-1}^2 a_{k-1}^2 \dots r_{k-1}^m o_{k-1}^m a_{k-1}^m)$  denotes the probability of the environment outputting the specified rewards and observations for all the agents given previous rewards, observations and actions. Environments should treat all agents equally, so changing the agent indices should not change the result. Other properties, such as sensitivity, minimum working (living) space, etc., could also be imposed, in order to define an appropriate multi-agent environment *class*<sup>6</sup>. Given this class, we can define a universal distribution by properly (prefix-free) coding its environments over a reference UTM  $U_e$ , using the formula  $p_E(\mu) := 2^{-K_{U_e}(\mu)}$ , where  $K()$  refers to Kolmogorov complexity. This is the probability of the base multi-agent environment not considering the agents. Another option here would be to define  $p_E$  using Schmidhuber's Speed Prior [21] or any other computable variant of  $K$ , such as Levin's  $Kt$  [18].

Using just one agent ( $m = 1$ ), where the agent is the evaluated agent, we would have an environment distribution similar to the one introduced in [17]. But we are interested in social environments where other agents appear. So we

<sup>6</sup> The notions of sensitivity and balancedness, introduced in [10], could be valid here as well if properly extended to multi-agent environments, in such a way that what an agent does may affect the others. For instance, the environment class introduced in [9] might be used, where rewards could be shared by the agents.

need to introduce more agents. In order to do this, we define an agent class and derive a universal distribution over its agents, (prefix-free) coding them over a reference UTM  $U_a$ , using the formula  $p_A(\pi) := 2^{-K_{U_a}(\pi)}$ .

When we join a base environment  $\mu$  with a set of  $m$  agents, we get an  $m$ -agent environment (a ‘social’ environment with population of size  $m$ ). From here we define the distribution of *start-up*  $m$ -agent environments (denoted by  $\sigma$ ) as:

$$p_S(\sigma) = p_S(\langle \mu, \pi_1, \pi_2, \dots, \pi_m \rangle) := p_E(\mu) \times \prod_{j=1}^m p_A(\pi^j)$$

Note that this definition considers base environments and agents to be independent. Of course, agents and environments are not independent in the real world, and they are not independent in evolutionary processes. The advantage of making them independent is that we can combine and analyse the distributions separately, and this may ease the construction of approximation without a fully (and intractable) evolutionary process.

And now systemic properties emerge when we let the social environment run (or “evolve”). We define the average reward for each agent  $j$  at step  $i$ , denoted by  $R_i^j$  as the mean of the rewards that the agent has received from step 0 to step  $i$ . We use this value  $R_i^j$  as a fitness function in a simplified ‘evolution’. We define the probability of dying  $d$  as a function of  $R_i^j$ . Since rewards range from  $-1$  to  $+1$  one possible choice is  $d(R_i) = \delta(1 - R_i^j)/2 + \epsilon$ , with  $0 < \delta < 1$  and  $0 \leq \epsilon < (1 - \delta)$  being small positive real numbers.

Finally, instead of considering gradual environment changes (as usual in evolution), we define a probability of the environment being completely replaced by a new one, while keeping the agents. The goal is to favour individual adaptability instead of pre-programmed specialisation to the environment. In particular, we define this probability as a constant (for simplicity). This small constant factor  $0 \leq c \leq 1$  means that for each step the environment might (with probability  $c$ ) be replaced by another environment, which would be chosen using  $p_E$ . Of course, if  $c = 0$  we stick to the original environment forever. The rationale for all this is that we do not want the agents to be optimised for the environments, but to be optimised to behave in the context of other agents. In general, environments will be very simple, so it is the multi-agent scenario which must greatly determine the distributions.

After each step, we apply  $d$  to all agents to determine whether they die or survive, replacing the dead agents by new agents using  $p_A$  (or  $q_{(d,c,i)}$ , see below). We do similarly for the environment, using  $c$ . Then, for each  $m$ -agent environment  $\sigma$ , function  $d$  and step  $i$ , this defines a probability of agent  $\pi^j$  after step  $i$ , denoted by  $q_{(d,c,i)}$ .

From here, we denote the Darwin-Wallace<sup>7</sup> distribution of degree  $i$  as:

$$p_{d,c,i}(\sigma) = p_i(\langle \mu, \pi_1, \pi_2, \dots, \pi_m \rangle) := p_E(\mu) \times \prod_{j=1}^m q_{(d,c,i)}(\pi^j)$$

---

<sup>7</sup> From Charles Darwin (1809-1882) and Alfred Russel Wallace (1823-1913).

Logically,  $p_0(\sigma) = p_S(\sigma)$ , since  $q_{(d,c,i)}(\pi^j) = p_A(\pi^j)$ .

Given the previous definition, what does it represent? It just assigns probabilities to multi-agent environments. For instance, an environment with sophisticated and evolutionary-adapted agents is much less likely according to  $p_{d,c,i}$  for low values of  $i$  than for large values of  $i$ , as the probability of a mammal on the Earth in the Precambrian is much lower than today. Note that this significantly differs from an evolutionary system (either artificial or natural) in many ways, since there is no concept of generation, there is no reproduction, no phenotype, no inheritance, no mutations and no cross-over. In addition, the environment may change drastically in just one step, so favouring general adaptive agents ahead of agents which specialise to (or in) a particular environment.

The purpose of all this is to isolate social adaptability instead of the adaptability to a single environment. The use of any mechanism from natural or artificial evolution that is not strictly necessary is then ruled out in this proposal.

The previous definition synthesises the notion that complex agents with complex behaviours might have a higher probability the higher the value of  $i$  is. Many of these complex behaviours are only useful when the environment is full of other complex agents, which we can only explain with a kind of evolutionary process. This is the reason why we use the name Darwin-Wallace for this distribution.

Consequently, previous definitions of intelligence [4, 11, 7, 17, 8, 10] can be re-understood by using variants of the previous distribution instead of variants of the universal distribution<sup>8</sup> for regular UTMs. Of course, this gives infinitely many definitions, since (apart from the reference UTM) we can parametrise  $i$  and  $m$ . The higher  $i$  and  $m$  are, the more ‘social’ the intelligence definition is. Nonetheless, despite the insightful and philosophical lure of the concept, it is difficult to apply in practice, because its uncomputability or intractability (if we use computable versions based on Levin’s Kt or Schmidhuber’s Speed Prior [21]) and many different options that could be taken (e.g., the probabilities  $d$  and  $c$ ).

## 4 Approximating the distribution through testing

The introduction of the previous distribution does not mean that we suggest ‘constructing’ intelligence (or deriving social environments) as the result of an inefficient, artificial evolutionary process. The alternative is an artificial selection rather than a natural selection. Starting with a multi-agent environment, we can introduce a first generation of agents, whose intelligence we still do not know. After a sufficient number of steps, we would get an average reward for each of them. From this individual assessment, we would (probabilistically) keep the agents with the best results, while removing the agents with the worst results, and introducing some new agents. With a large number of agents, a variety of

---

<sup>8</sup> There can be a machine  $U$  such that the universal distribution using this machine may match (at least approximately) a Darwin-Wallace distribution of degree  $i$  and population size  $m$ . In theory this is possible (ignoring time in  $K_U$  and excluding some environments without agents), but finding such a machine  $U$  is similar to giving the definition of the Darwin-Wallace distribution above, which is our goal.

environments and several interactions with this assisted process, we would have a testbed where only socially intelligent agents would score well.

It can be argued that this process is not very different to recreating natural environments and embedding humans inside. In fact, the Turing Test is an example of this. However, in our proposal, environments are not anthropomorphic, they do not rely on natural language, and they accept any kind of agent. Furthermore, this process can scale up far beyond human intelligence, since it is intelligence inside a social system that we measure (and use) to feed the system.

There are, of course, many issues about this process. For instance, it is possible that a highly intelligent system  $\pi_a$  could score worse than another less intelligent system  $\pi_b$  for a low degree of the test, say  $i_1$ . This is possible, but considering that this is a probabilistic process and that  $\pi_a$  might be re-introduced again for a higher degree  $i_2 > i_1$  (we have checked its intelligence for  $i_1$ , but not for  $i_2$ ), then we could eventually have  $\pi_a$  in its right place.

But of course, there might be some environments where rewards are a limited resource. In these environments, competition would be expected. On the contrary, co-operation could be encouraged if rewards could be shared (without limits) by a group of individuals. For instance, in the environment class introduced in [9], we can define rewards in such a way that they are eaten by the agents (to favour competitiveness) but we can also define rewards which remain.

## 5 Discussion

There is an increasing consensus in biology, comparative psychology and even artificial intelligence that intelligent systems must be social. Multi-agent systems appeared as a realisation of this fact, and multi-agent reinforcement learning is also in this direction. Several works in the AGI community have also advocated for a social approach to intelligence. However, we did not have a formal definition of intelligence with social abilities *playing an important role*.

Here we have proposed a novel environment distribution and a constructive method to build social environments incrementally. In fact, the previous Darwin-Wallace distribution can replace the universal distribution used in other formal definitions and tests of intelligence presented to date. Of course, the appropriateness and applicability of this would require theoretical or empirical results.

Although the distribution is related to (because it is inspired by) evolutionary computation, artificial life and natural biology, it is a novel approach to understanding the set of environments and tasks we want general intelligent systems to cope with. There are three features in the definition which distinguish it from a ‘distribution of life forms’: i) Physical things do not matter, placing the focus on behaviour, since the ‘body’ is the same for all. In other words, it is a ‘distribution of mind forms’. ii) There is no genotype, cross-over, mutation, etc., so selection does not work for genes or species, but for individuals. In fact, in biological evolution, genes compete and collaborate, and we do not consider here the ‘distribution of genes’, either. iii) Environments are replaced, so avoiding

specialisation (to or) in an environment. Instead, adaptability to a wide range of environments (i.e., intelligence) is the only fitness function for selection.

One of the problems of this approach is that now we are far away from the notion of ‘intelligence as compression’. Nonetheless, the notions of mind-reading and manipulation imply that agents need to be able to capture the models of other agents, i.e. to compress their behaviours. Consequently, intelligence is related to the concept of adversarial (reinforcement) learning [23] and, eventually, to the elusive-model paradox<sup>9</sup>. These two concepts are of course related to measuring intelligence using games as environments, as suggested in [10]. And the role of two-part compression can be vindicated again in social environments in terms of communication, since having one model, the first-part of a MML message (namely, the model/theory/hypothesis/concept), which is concise and can be (relatively quickly and concisely) transmitted between agents (in contrast to a weighted ensemble of models) must be crucial for communication.

Summing up, we think that the main contributions of this paper span over the definition, evaluation and development of intelligence. We are conscious that there are many open questions and many implications, especially on the re-understanding of previous works on defining and evaluating intelligence, and on the direct applicability of these ideas to develop more intelligent machines.

## Acknowledgments

We thank the anonymous reviewers for their helpful comments. We acknowledge the funding from the Spanish MEC and MICINN for projects TIN2009-06078-E/TIN, Consolider-Ingenio CSD2007-00022 and TIN2010-21062-C02, for MEC FPU grant AP2006-02323, and Generalitat Valenciana for Prometeo/2008/051.

## References

1. D. L. Dowe. Foreword re C. S. Wallace. *Computer Journal*, 51(5):523 – 560, Sep. 2008. Christopher Stewart WALLACE (1933-2004) memorial special issue.
2. D. L. Dowe. Minimum Message Length and statistically consistent invariant (objective?) Bayesian probabilistic inference - from (medical) “evidence”. *Social Epistemology*, 22(4):433 – 460, October - December 2008.
3. D. L. Dowe. MML, hybrid Bayesian network graphical models, statistical consistency, invariance and uniqueness. In P. S. Bandyopadhyay and M. R. Forster,

---

<sup>9</sup> The elusive-model paradox can be summarised as follows. Imagine two agents where one (perhaps the prey) is trying to anticipate where the other will be so as to avoid meeting, and the other (perhaps the predator) is trying to anticipate the other so that they do meet. In re-discovering this “elusive model paradox” [1, footnote 211][2, p 455], we note that, if the one agent is better at inference then it will be able to anticipate the behaviour of the other agent better than vice versa [3, sec. 7.5]. But, if both agents have the ability to infer Turing-complete functions and have statistically consistent inference methods, then, we will indeed encounter the halting problem and decisions (regarding actions) that can not be made in finite time.

- editor, *Handbook of the Philosophy of Science - Volume 7: Philosophy of Statistics*, pages 901–982. Elsevier, 2011.
4. D. L. Dowe and A. R. Hajek. A computational extension to the Turing Test. *4th Conf. of the Australasian Cognitive Science Society, Newcastle, Australia*, 1997.
  5. B. Goertzel. The Embodied Communication Prior: A characterization of general intelligence in the context of Embodied social interaction. In *Cognitive Informatics, 2009. ICCI'09. 8th IEEE International Conference on*, pages 38–43. IEEE, 2009.
  6. B. Goertzel and S.V. Bugaj. AGI Preschool: a framework for evaluating early-stage human-like AGIs. In *Intl. Conf. on Artificial General Intelligence (AGI-09)*, 2009.
  7. J. Hernández-Orallo. Beyond the Turing Test. *J. Logic, Language & Information*, 9(4):447–466, 2000.
  8. J. Hernández-Orallo. On the computational measurement of intelligence factors. In A. Meystel, editor, *Performance metrics for intelligent systems workshop*, pages 1–8. National Institute of Standards and Technology, Gaithersburg, MD, 2000.
  9. J. Hernández-Orallo. A (hopefully) non-biased universal environment class for measuring intelligence of biological and artificial systems. In M. Hutter et al., editor, *Artificial General Intelligence*, pages 182–183, 2010.
  10. J. Hernández-Orallo and D. L. Dowe. Measuring universal intelligence: Towards an anytime intelligence test. *Artificial Intelligence*, 174(18):1508 – 1539, 2010.
  11. J. Hernández-Orallo and N. Minaya-Collado. A formal definition of intelligence based on an intensional variant of Kolmogorov complexity. In *Proc. Intl Symposium of Engineering of Intelligent Systems (EIS'98)*, pages 146–163. ICSC Press, 1998.
  12. E. Herrmann, J. Call, M. V. Hernández-Lloreda, B. Hare, and M. Tomasello. Humans have evolved specialized skills of social cognition: The cultural intelligence hypothesis. *Science*, Vol 317(5843):1360–1366, 2007.
  13. B. Hibbard. Bias and No Free Lunch in Formal Measures of Intelligence. *Journal of Artificial General Intelligence*, 1(1):54–61, 2009.
  14. J.R. Krebs and R. Dawkins. Animal signals: mind-reading and manipulation. *Behavioural Ecology: an evolutionary approach*, 2:380–402, 1984.
  15. C.G. Langton. *Artificial life: An overview*. The MIT Press, 1997.
  16. S. Legg and M. Hutter. A collection of definitions of intelligence. In *Proc. of the 2007 Conf. on Artificial General Intelligence*, pages 17–24. IOS Press, 2007.
  17. S. Legg and M. Hutter. Universal intelligence: A definition of machine intelligence. *Minds and Machines*, 17(4):391–444, 2007.
  18. L. A. Levin. Universal sequential search problems. *Problems of Information Transmission*, 9(3):265–266, 1973.
  19. P. Sanghi and D. L. Dowe. A computer program capable of passing IQ tests. In *Proc. 4th ICCS International Conference on Cognitive Science (ICCS'03), Sydney, Australia*, pages 570–575, 2003.
  20. J. Schmidhuber. A computer scientist’s view of life, the universe, and everything. In *Foundations of Computer Science*, page 201. Springer, 1997.
  21. J. Schmidhuber. The Speed Prior: a new simplicity measure yielding near-optimal computable predictions. In *Computational Learning Theory*, pages 123–127. Springer, 2002.
  22. R. J. Solomonoff. A formal theory of inductive inference. Part I. *Information and control*, 7(1):1–22, 1964.
  23. P. Stone and M. Veloso. Towards collaborative and adversarial learning: A case study in robotic soccer. *Intl. J. of Human-Computers Studies*, 48(1):83–104, 1998.
  24. M. Tomasello and E. Herrmann. Ape and human cognition : What’s the difference? *Current Directions in Psychological Science*, 19(1):3–8, March 2010.