



Incremental and developmental perspectives for general-purpose learning systems

Fernando Martínez-Plumed

Universitat Politècnica de València, fmartinez@dsic.upv.es

Abstract The stupefying success of Artificial Intelligence (AI) for *specific* problems, from recommender systems to self-driving cars, has not yet been matched with a similar progress in *general* AI systems, coping with a variety of (different) problems. This dissertation deals with the long-standing problem of creating more general AI systems, through the analysis of their development and the evaluation of their cognitive abilities. It presents a declarative general-purpose learning system and a developmental and lifelong approach for knowledge acquisition, consolidation and forgetting. It also analyses the use of more ability-oriented evaluation techniques for AI evaluation and provides further insight for the understanding of the concepts of development and incremental learning in AI systems.

Keywords: artificial intelligence, general-purpose learning systems, inductive programming, reinforcement learning, forgetting, task difficulty, cognitive development, evaluation of artificial systems, intelligence tests

Extended Abstract

In the light of all the astonishing achievements in recently AI research, it is becoming increasingly clear that creating artificial intelligence is much more than “pattern matching”. However, although it would be unfair to deny that some current AI systems exhibit some intelligent behaviour (especially those that incorporate some learning potential), in general terms, most AI research is focused on designing AI systems for a particular functionality or adapted for a specific problem with no intention whatsoever of featuring intelligence. Up to date, the vast majority of the computer models are mindless rule-followers or cleverly written computer program doing statistical calculations and making predictions based on them. However, what it would mean for a computer to behave in an intelligent way? This thesis [9] states that the answer lies in the construction of systems that go beyond task specific scenarios into more general-purpose ones thus able to learn automatically, not pre-programmed or without fixed handcrafted features.

Given the above challenge, in the presented dissertation we characterise a series of human intelligence attributes (incremental, developmental and lifelong learning) and cognitive-oriented procedures (memory and forgetting) that, combined with the use of symbolic AI and symbolic learning, have helped us to develop both a general-purpose learning approach as well as a knowledge handling tool. This ambitious issue should, furthermore, pervade the evaluation procedures in AI where systems are usually evaluated in terms of task performance, not really in terms of intelligence. Hence, AI evaluation must necessarily be linked to the purpose of the discipline: general AI systems should require an ability-oriented evaluation in the same way that specialised AI systems should require a task-oriented evaluation.

Particularly, and regarding the construction of more general AI approaches, this thesis contributes with a pair of settings for learning and knowledge acquisition. Firstly we present a general-purpose

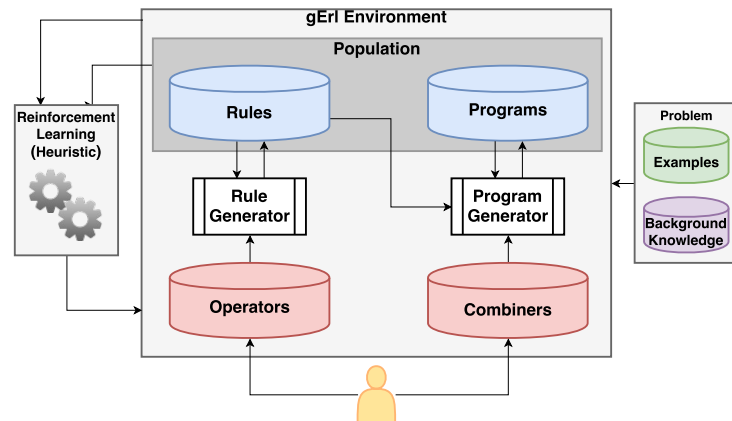


Figure 1: **gErI** takes examples and learning operators as input and returns functional programs. There are two internal repositories containing rules and programs which are updated at each learning step by the *Rule Generator* and the *Program Generator* processes. The *Reinforcement Learning Module* is in charge of defining combinations of states and actions (operator applied over a rule) which will be used by the *Rule Generator* process to select future actions to perform.

declarative learning system (**gErI**) [1, 2, 3, 4] that meets several desirable characteristics in terms of expressiveness, comprehensibility and versatility. We have shown that more general systems can be constructed by not only giving power to data and background knowledge representation but also to a flexible operator redefinition and the reuse of heuristics across problems and systems. **gErI** (Fig. 1) relies on two compatible mechanisms. The former is the definition of customised learning operators, depending on the data structures and problem at hand, done by the user, using a functional language. The latter mechanism is the use of generalised heuristics, since the use of different operators precludes the system from using specialised heuristics for each of them. The choice of the right pair of operator and rule has been reframed as a decision process (using a reinforcement learning approach). Therefore, not only is this a novel approach, but also allows us to better understand the role of operators and heuristics in machine learning. By performing a series of illustrative experiments we show where the flexibility stands out, since **gErI** is able to solve a wide range of problems (from recursive ones to several IQ tests).

Secondly, the learning process is also overhauled with a new developmental and lifelong approach for knowledge acquisition, consolidation and forgetting, which is necessary when bounded resources (memory and time) are considered. In this sense we present a parametrisable (hierarchical) approach [6, 5] for structuring knowledge which is able to check whether the new learnt knowledge can be considered redundant, irrelevant or inconsistent with the old one, and whether it may be built upon previously acquired knowledge. This approach is designed to combine any rule-based inductive engine with a deductive engine (is, therefore, parametrisable to other cognitive or intelligent systems) and integrates them into a lifelong learner through the use of a hierarchical knowledge assessment structure (based on coverage) and by introducing several information theory-based metrics. Therefore, given a lifelong learning problem, our approach is able to discover and develop knowledge incrementally by means of assessing the usefulness of the rules and gradually generating a large repository of consolidated knowledge where the knowledge is revised in order to generate a rich and reusable knowledge base. Particularly, we have analysed how appropriate these cognitive mechanisms are in order to deal with declarative knowledge bases in intelligent systems that are meant to have a non-ephemeral life. This complex knowledge organisation and assessment mechanisms allows for a straightforward and principled approach to knowledge acquisition, consolidation (promotion), revision (demotion) and forgetting.

Thirdly, and moving towards AI evaluation, this thesis analyses whether the use of more ability-oriented evaluation techniques for AI (such as intelligence tests) is a much better alternative to most task-oriented evaluation approaches in AI. Accordingly, we make a review of what has been done when AI systems have been confronted against tasks taken from intelligence tests [7, 8]. In this regard, we scrutinised what intelligence tests measure in machines, whether they are useful to evaluate AI systems,

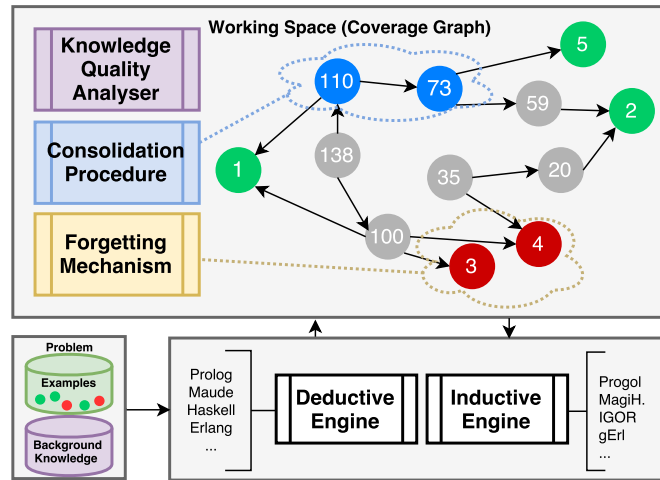


Figure 2: Organisation of hierarchical knowledge assessment structures in terms of coverage and information theory-based principles. This approach is designed to combine any rule-based inductive engine with a deductive engine and integrates them into a lifelong learner.

whether they are really challenging problems, and whether they are useful to understand (human) intelligence by analysing over 30 papers featuring AI systems addressing intelligence test problems. We have studied and characterised each system by their relationships, the range of intelligence test tasks they address, the purpose of the models, how general or specialised these models are, the AI techniques they use in each case, their comparison with human performance and their evaluation of item difficulty. Through this analysis we have realised that those systems have different purposes and applications: to advance AI by the use of challenging problems, to use intelligence tests for the evaluation of AI systems, to better understand intelligence tests and what they measure (including item difficulty), and, finally, to better understand what human intelligence is. Furthermore, we have seen that these systems systematically ignore results and ideas already present in previous related approaches specialising to the task and, therefore, losing the opportunity to understand what a computer model passing an intelligence test really means. Our aim here is both to encourage any future computer model taking intelligence tests to link with and build upon previous research, and to contribute to a more widespread realisation that more general classes of problems are needed when constructing benchmarks for AI evaluation.

By the same token, as a final contribution, we show that intelligence tests can also be useful to examine concept dependencies (mental operational constructs) in the cognitive development of artificial systems (although a superficial score comparison is misleading), therefore supporting the assumption that, even for fluid intelligence tests, the difficult items require a more advanced cognitive development than the simpler ones. In this sense, we show [10] how several fluid intelligence test problems (odd-one-out problems, Raven's Progressive Matrices and Thurstone's letter series) are addressed by our general-purpose learning system **gErl**, which, although lacks any mental epigenetic development and physical embodiment and it is not particularly designed on purpose to solve intelligence tests, is able to perform relatively well for this kind of tests. **gErl** makes it explicitly how complex each pattern is and what operators are used for each problem due its symbolic and declarative nature: rule-based representation language for examples, patterns and operators. This provides useful information about the role of the cognitive operational constructs that are needed to solve a problem or task. Therefore, the goal has not been to evaluate **gErl** but to use it as a tool to gain some insights into the characteristics and usefulness of these tests and how careful we need to be when applying human test problems to assess the abilities and cognitive development other AI systems. We do think that, in general terms, for both humans and machines, human intelligence tests are useful to evaluate cognitive development through the diversity of cognitive operational constructs required, therefore supporting the assumption that, even for fluid intelligence tests, the difficult items require a more advanced cognitive development than the simpler ones.

Summing up, this dissertation represents one step forward in the hard and long pursuit of making more

general AI systems and fostering less customary (and challenging) ability-oriented evaluation approach. Comprehensibility, expressiveness, incrementality and developmental knowledge discovery are all desirable features for this general-purpose AI development, apart from the requirement of accurate, effective and meaningful ability-oriented ways for evaluating its progress. For this purpose we have integrated different topics both within and outside AI, such as machine learning, inductive programming, reinforcement learning, cognitive science and psychometrics. From a methodological point of view, we have considered some conceptual developments with systematic empirical evidence.

Acknowledgements

I would like to thank my Ph.D. advisors, José Hernández-Orallo, María José Ramírez-Quintana and Cèsar Ferri, for all his help and guidance.

This work has been partially supported by the EU (FEDER) and Spanish MINECO grant TIN2015-69175-C4-1-R, and the REFRAME project, granted by the European Coordinated Research on Long-term Challenges in Information and Communication Sciences Technologies ERA-Net (CHIST-ERA), and funded by MINECO in Spain (PCIN-2013-037) and by Generalitat Valenciana PROMETEOII/2015/013.

References

- [1] F. Martínez-Plumed, C. Ferri, J. Hernández-Orallo, and M. J. Ramírez-Quintana. Accelerating learning by reusing abstract policies in gerl. In Concha Bielza et al., editor, *Multiconferencia CAEPIA'13*, VII Simposio de Teoría y Aplicaciones de Minería de Datos, TAMIDA'13, pages 1383–1392, 2013.
- [2] F. Martínez-Plumed, C. Ferri, J. Hernández-Orallo, and M. J. Ramírez-Quintana. Learning with configurable operators and rl-based heuristics. In A. Appice, editor, *New Frontiers in Mining Complex Patterns*, volume 7765 of *LNCS*, pages 1–16. Springer, 2013.
- [3] F. Martínez-Plumed, C. Ferri, J. Hernández-Orallo, and M. J. Ramírez-Quintana. On the definition of a general learning system with user-defined operators. *CoRR*, abs/1311.4235, 2013.
- [4] F. Martínez-Plumed, C. Ferri, J. Hernández-Orallo, and M. J. Ramírez-Quintana. Policy reuse in a general learning framework. In Concha Bielza et al., editor, *Multiconferencia CAEPIA'13*, 15th Conference of the Spanish Association for Artificial Intelligence, CAEPIA'13, pages 119–128, 2013.
- [5] F. Martínez-Plumed, C. Ferri, J. Hernández-Orallo, and M. J. Ramírez-Quintana. A knowledge growth and consolidation framework for lifelong machine learning systems. In *Machine Learning and Applications (ICMLA), 2014 13th International Conference on*, pages 111–116, Dec 2014.
- [6] F. Martínez-Plumed, C. Ferri, J. Hernández-Orallo, and M. J. Ramírez-Quintana. Knowledge acquisition with forgetting: an incremental and developmental setting. *Adaptive Behavior*, 23(5):283–299, 2015.
- [7] J. Hernández-Orallo, F. Martínez-Plumed, Ute Schmid, Michael Siebers, and David L. Dowe. Computer models solving intelligence test problems: Progress and implications. *Artificial Intelligence*, 230:74–107, 2016.
- [8] F. Martínez-Plumed, J. Hernández-Orallo, Ute Schmid, Michael Siebers, and David L. Dowe. Historical account of computer models solving iq test problems. In *Evaluating General-Purpose AI (EGPAI 2016)*, 22nd European Conference on Artificial Intelligence (ECAI 2016), pages 20–21, Sept 2016.
- [9] Fernando Martínez-Plumed. *Incremental and developmental perspectives for general-purpose learning systems*. PhD thesis, Technical University of Valencia, Spain, April 2016.
- [10] F. Martínez-Plumed, C. Ferri, J. Hernández-Orallo, and M. J. Ramírez-Quintana. A computational analysis of general intelligence tests for evaluating cognitive development. *Cognitive Systems Research*, pages –, 2017.