

Special Issue on Autonomous Learning

Barbara Hammer¹ · Marc Toussaint²

Published online: 18 September 2015
© Springer-Verlag Berlin Heidelberg 2015

1 Introduction

The ability to learn is a central aspect of human and animal intelligence. We adapt and improve our behaviour continuously when interacting with our environment—it is impossible for a human not to ‘learn’ from ubiquitous environmental clues and to adapt his/her behaviour accordingly. Learning from an explicit teacher signal (such as happens at school or university lectures) constitutes a minor part of human learning; the majority takes place in an autonomous way without explicit feedback or structuring of the learning task by a teacher. Interestingly, there are also attempts in pedagogic to not only individualise class teaching but to also put more emphasis on learner autonomy for a better learning effect and improved life-long learning abilities [1]. However, to date, existing paradigms of technical learning systems are still remote from the autonomy of human learning.

Machine learning refers to the technology to enable machines to improve their behaviour based on experience [15]. It has made great strides in the last decades and constitutes a central part in diverse technical tasks such as image classification, spam filtering, internet search, speech processing, biomedical data analysis, fraud detection, etc. [5]. Albeit the vision of machine learning is to mimic the

flexibility and efficiency of humans to learn from data [14], nowadays, classical machine learning is mostly concerned with highly restricted settings: the majority of machine learning scenarios addresses supervised tasks, batch scenarios, and learning from a fixed (often feature-based) data representation; most mathematical learning theoretical results rely on the fundamental assumption of data being independent and identically distributed [6].

A great part of machine learning emerged from an abstraction of learning processes as observed in neurobiology; quite a few early learning models are based on biological paradigms such as Hebbian learning, i.e. the corresponding algorithms are based on heuristics—the perceptron learning rule for classification or Oja’s rule for dimensionality reduction constitute prominent examples [11]. Albeit these learning rules often reveal a striking simplicity, their mathematical foundation can be quite complicated, even partially unsolved [4, 9]. With time, more and more mathematical formalisms emerged within machine learning which enable researchers to derive a learning rule based on an explicit mathematical modelling of the underlying goal: most modern techniques are accompanied by mathematical cost functions such as the least squares error for regression, constraint margin maximisation for robust classification in support vector machines, or the data log likelihood for model inference [5]. While this mathematical treatment leads to highly efficient learning techniques, it often narrows the applicability of the techniques to specific areas which are covered by the considered cost function and data representation. This way, the machine learning algorithm is uncoupled from the question what is a good cost function and parameterisation of it, what is a good data representation, and what is an informative data set based on which learning becomes possible. These crucial ingredients of machine

✉ Barbara Hammer
bhammer@techfak.uni-bielefeld.de

Marc Toussaint
marc.toussaint@informatik.uni-stuttgart.de

¹ CITEC Centre of Excellence, Bielefeld University,
33594 Bielefeld, Germany

² Machine Learning and Robotics Lab, University of Stuttgart,
70569 Stuttgart, Germany

learning are typically hand-coded and optimised by a human expert.

Hence there exists a gap between autonomous human learning and machine learning, since the latter often requires extensive human intervention to shape the learning tasks into exact mathematical formalisms which enable the use of powerful techniques from statistical modelling and mathematical optimisation. In most technical scenarios, learning is decoupled from decision making and from the question how to gather the relevant information based on which to learn. Even the area of reinforcement learning, which integrates decision making and learning into one underlying mathematical objective, crucially relies on the possibility to partially decouple these two aspects based on the notion of the value function or related quantities, and it crucially relies on a suitable parameterisation of these costs and data representation [3]. This observation has initiated a discussion about efficient ways to make machine learning more autonomous, see e.g. [7, 8]. One reason for this discussion lies in the fact that quite a few success stories offer novel stimuli towards learner autonomy such as autonomous representation learning from raw image data based on deep architectures, efficient ways to directly handle non-vectorial and symbolic data structures in machine learning, and autonomous optimisation of meta-parameters in learning systems [2, 10, 12, 13].

Autonomous learning research aims at understanding how adaptive systems can efficiently learn from the interaction with the environment, especially by having an integrated approach to decision making and learning, allowing systems to decide by themselves on actions, representations, hyper-parameters and model structures for the purpose of efficient learning. It addresses challenges such as how to autonomously learn representations for efficient model use, how to arrive at suitable cost functions from meta-objectives, how to autonomously choose model structures and hyper-parameters in possibly non-stationary environments, or how to design efficient actor-reward strategies which generalise across tasks. Application scenarios which require these types of complex models span high-impact domains such as robotics, life-long learning, intelligent tutoring, or big data analytics [16]. The German Research Foundation initiated the Priority Programme on *Autonomous Learning (SPP 1527)*¹ to support basic research in these areas, including many of the contributions in this special issue.

The collection of articles in this issue addresses central aspects of autonomous learning and includes overviews, discussions and project descriptions about novel research in this exciting area. Altogether, seven research articles, three project reports, one interview, and one Ph.D. presentation

highlight fundamental approaches to enable a greater learner autonomy in learning systems. The article *Beyond Manual Tuning of Hyperparameters* by Frank Hutter, Jörg Lücke, and Lard Schmidt-Thieme provides an overview about approaches how to autonomously choose crucial meta-parameters of a learning system, which can considerably ease the burden of model design by experts. Three prominent strategies are discussed in more detail: an automation of exhaustive search over the range of possible meta-parameters by means of an autonomous probabilistic model, an autonomous inference of regularities for an optimum parameter choice in terms of characteristics across data sets, and a reduction of the number of hyper-parameters by an autonomous unsupervised inference of data regularities from the training data itself.

The two contributions *Autonomous Learning of Representations* by Oliver Walter, Reinhold Haeb-Umbach, Bassam Mokbel, Benjamin Paassen, and Barbara Hammer and *Autonomous Learning of State Representations for Control* by Wendelin Böhmer, Jost Tobias Springenberg, Joschka Boedecker, Martin Riedmiller, and Klaus Obermayer both center around the question how to best represent training data for a specific learning task. The first article focuses on supervised and unsupervised scenarios, while the second contribution addresses simultaneous learning and decision making based on an efficient state space representation in reinforcement scenarios. All learning scenarios have in common that inadequate representations easily suffer from the curse of dimensionality and a limited suitability of the standard euclidean norm for its further processing. Thus the question is how to autonomously shape representations such that meaningful aspects of the data are emphasised. Promising technologies are offered by deep learning approaches which can reduce the data dimensionality of raw sensor data in a very effective way, efficient information extraction based on natural priors such as sparse coding or slow feature analysis, or metric learning which enables to focus the attention on relevant parts of the full data space only. The first contribution elucidates novel technology how deep learning and metric learning, respectively, can efficiently be realised in the context of sequential data structures such as language. The second addresses robotics scenarios and highlights the suitability of deep learning and slow feature analysis, respectively, in these settings. Interestingly, the specific demands of reinforcement learning allow to explicit mathematical requirements of a data representation such as Markovianity.

Symbolic data representations together with planning or logical reasoning offer efficient representations which naturally incorporate compositionally and avoid the curse of dimensionality which arises from the necessity of an enumeration within feature-based representations. *Statistical*

¹ <http://www.autonomous-learning.org>.

Relational Artificial Intelligence (StaRAI) extends these powerful representations towards probabilistic modelling, such incorporating robustness to noise and efficiency of probabilistic computations in a natural way. In their contribution, Kristian Kersting and Sriraam Natarajan give an intriguing overview about the techniques in StaRAI, in particular highlighting the merits obtained by relational reinforcement learning.

A central problem in autonomous learning is the question how to efficiently realise life-long adaptation across tasks. The problem of multi-task active and transfer learning in robotics constitutes the focus of the contribution *Accounting for Task-Difficulty in Active Multi-Task Robot Control Learning* by Alexander Fabisch, Jan Hendrik Metzen, Mario Michael Krell, and Frank Kirchner. Besides an overview of transfer learning approaches, the authors present an intriguing framework how to deal with incommensurable value functions across tasks to enable an integrated optimality modelling.

The article *The Optimization Route to Robotics—and Alternatives* by Marc Toussaint, Helge Ritter, and Oliver Brock takes a general view on robotics in discussing the question to what extent optimality principles constitute the best way to model robotics systems. While this seems unanimously true for single components, suitable design principles for complex, integrated robotics systems are less obvious. In their contribution, the authors discuss interesting alternatives and variants which could form the ground for a formalisation of complex robotics systems.

Robotics control is crucially biased by the physical constraints of the robot and its environment, and it is possible to generate complex behaviour by complex constraints rather than a complex control strategy. Information geometry offers a powerful mathematical framework in which to formalise these constraints, enabling alternative concepts for important characteristics such as approximation universality. The contribution *Geometric Design Principles for Brains of Embodied Agents* by Nihat Ay provides a unifying view on mathematical results and modelling techniques in this area.

These overview articles are accompanied by presentations of a number of projects which are conducted in the area of autonomous learning. Robotics constitutes one prime application area in which autonomous learning offers particular promises. In the project *Online Learning of Bipedal Walking Stabilization*, Marcell Missura and Sven Behnke address the challenging task of humanoid walking stabilisation, where bootstrapping of basic walking skills offers intriguing ways to achieve robust results in real hardware. Data gathered in robotics experiments need to be stored, preferably in such a way that episodic memory can be combined with prior experience and common knowledge, and higher representations enable inference and

planning on a global level. The project *OPEN-EASE—A Cloud-Based Knowledge Service for Autonomous Learning* presented by Moritz Tenorth, Jan Winkler, Daniel Beßler, and Michael Beetz addresses these demanding tasks by means of an intelligent cloud-based infrastructure. The project *Learning Feedback in Intelligent Tutoring Systems* by Sebastian Gross, Bassam Mokbel, Barbara Hammer and Niels Pinkwart tackles the question how to provide feedback in ill-posed domains for intelligent tutoring of scholars where a formalisation of correct solutions of a task are not possible. In such cases, prototype based models from machine learning offer promising approaches which enable to generate feedback from historic data for humans while they try to solve the given task.

One crucial ingredient of many deep learning models for the autonomous inference of a low dimensional data representation is a restricted Boltzmann machine which is trained to maximise the data likelihood. *Training restricted Boltzmann machines* constitutes a challenging problem. In her Ph.D. thesis, Asja Fischer mathematically and empirically investigates popular learning algorithms and proposes efficient improvements.

The special issue concludes with an interview about autonomous learning with *Werner von Seelen*, conducted by Benjamin Paassen. In his research, Werner von Seelen has pioneered connections of learning in biological systems and technical applications, and he discusses in detail demands and merits which arise from these different points of view.

This special issue is the result of the efforts of many persons. Special thanks go to the authors for their interesting contributions surveying hot topics and projects in the area, our interview partner for his very inspiring opinion about autonomous learning, our reviewers of the papers, who did a great job in providing short-notice reviews, and to the KI team, in particular the editor in charge of this issue, Christian Igel (University of Copenhagen), for excellent support.

Barbara Hammer, Marc Toussaint

2 Content

2.1 Technical Contributions

- *Beyond Manual Tuning of Hyperparameters* Frank Hutter, Jörg Lücke, Lars Schmidt-Thieme.
- *Autonomous Learning of Representations* Oliver Walter, Reinhold Haeb-Umbach, Bassam Mokbel, Benjamin Paassen, Barbara Hammer.
- *Autonomous Learning of State Representations for Control—An Emerging Field Aims to Autonomously Learn State Representations for Reinforcement*

Learning Agents from their Real-World Sensor Observations Wendelin Böhmer, Jost Tobias Springenberg, Joschka Boedecker, Martin Riedmiller, Klaus Obermayer.

- *Statistical Relational Artificial Intelligence—From Distributions through Actions to Optimization* Kristian Kersting, Sriraam Natarajan.
- *Accounting for Task-Difficulty in Active Multi-Task Robot Control Learning* Alexander Fabisch, Jan Hendrik Metzen, Mario Michael Krell, Frank Kirchner.
- *The Optimization Route to Robotics—and Alternatives* Marc Toussaint, Helge Ritter, Oliver Brock.
- *Geometric Design Principles for Brains of Embodied Agents* Nihat Ay.

2.2 Research Projects

- *Online Learning of Bipedal Walking Stabilization* Marcell Missura, Sven Behnke.
- *OPEN-EASE—A Cloud-Based Knowledge Service for Autonomous Learning.* Moritz Tenorth, Jan Winkler, Daniel Beßler, Michael Beetz.
- *Learning Feedback in Intelligent Tutoring Systems* Sebastian Gross, Bassam Mokbel, Barbara Hammer, Niels Pinkwart.

2.3 Doctoral Dissertation

- *Training restricted Boltzmann machines.* Asja Fischer.

2.4 Interview

- *Interview with Werner von Seelen* (Founding member of the Institute for Neuroinformatics at the Ruhr University Bochum, former president of the German Society for Cybernetics).
Benjamin Paassen

3 Service

Autonomous learning lies at the intersection of machine learning, artificial intelligence, and prominent application scenarios such as robotics, hence major contributions are spread over different sources. In the following, we include an incomplete list of links to interesting journals, conferences, and software in this domain. Further material is also summarised at the web site of the DFG priority program *autonomous learning*: <http://ipvs.informatik.uni-stuttgart.de/mlr/spp-wordpress/>.

3.1 Journals

Some journals with a main focus on machine learning and autonomous robotics include the following:

- Journal of Machine Learning Research <http://www.jmlr.org/>,
- Journal of Artificial Intelligence Research <http://www.jair.org/>,
- IEEE Transactions on Pattern Analysis and Machine Intelligence <http://www.computer.org/web/tpami>,
- IEEE Transaction on Neural Networks and Machine Learning <http://cis.ieee.org/ieee-transactions-on-neural-networks-and-learning-systems.html>,
- Machine Learning <http://link.springer.com/journal/10994>,
- Neurocomputing <http://www.journals.elsevier.com/neurocomputing/>,
- Artificial Intelligence <http://www.journals.elsevier.com/artificial-intelligence/>,
- International Journal of Robotics Research <http://www.ijrr.org/>,
- Autonomous robots <http://www.springer.com/engineering/robotics/journal/10514>,
- Robotics and Autonomous Systems <http://www.journals.elsevier.com/robotics-and-autonomous-systems/>,
- IEEE Transactions on Robotics <http://www.ieee-ras.org/publications/t-ro>.

3.2 Conferences and Workshops

Some conferences specifically dedicated to machine learning, artificial intelligence, and robotics are the following:

- Neural Information Processing Systems (NIPS) <http://nips.cc/>,
- International Conference on Machine Learning (ICML) <http://icml.cc/>,
- International Conference on Artificial Intelligence and Statistics (AISTATS),
- International Joint Conferences on Artificial Intelligence (IJCAI) <http://ijcai.org/>,
- Robotics: Science and Systems (R:SS) <http://www.roboticsconference.org/>,
- IEEE International Conference on Robotics and Automation (ICRA),
- IEEE International Conference on Intelligent Robots and Systems (IROS),
- IEEE International Joint Conference on Neural Networks (IJCNN), every second year within the World Congress of Computational Intelligence (WCCI),

- European Symposium on Artificial Neural Networks (ESANN) <http://www.dice.ucl.ac.be/esann/>,
- European Conference on Machine Learning (ECML/PKDD) <http://www.ecmlpkdd.org/>,
- German Conference on Artificial Intelligence (KI) <http://konferenz.kuenstliche-intelligenz.de/>,
- GI Workshop on New Challenges in Neural Computation and Machine Learning (NC²) <http://www.techfak.uni-bielefeld.de/~bhammer/GINN/NC2/nc2.html/>.

3.3 Software and Link Collections

Some exemplary links to domains which summarise resources for autonomous learning or code are given below; there exist numerous further resources which are not covered in this list. Further, many publications are accompanied by downloadable code and many authors provide their code on request.

- AutoML covers resources which aim to put the human out of the loop in the design of machine learning techniques, among other resources, it links to AutoWEKA: <http://automl.org/>.
- Deep learning resources are collected at this site, among other code the software Theano is linked which also offers a step-by-step course to learn to use deep networks: <http://deeplearning.net/>.
- Shark is a C++ machine learning library featuring a wide range of learning algorithms including deep learning, besides useful optimisation tools: <http://image.diku.dk/shark/>.
- A page connected to kernel techniques, support vector machines, and Gaussian processes is the following: <http://www.kernel-machine.org/>.
- A page linking to various machine learning open source software is the following: <http://mloss.org/>.
- Torch provides a Matlab-like scripting language for machine learning algorithms based on a fast C implementation: <http://torch5.sourceforge.net/>.

References

1. Armstrong J (2004) The natural learning project. General economics and teaching, EconWPA. <http://EconPapers.repec.org/RePEc:wpa:wuwpgt:0412021>. Accessed 14 Sept 2015
2. Bengio Y, LeCun Y (2007) Scaling learning algorithms towards AI. In: Bottou L, Chapelle O, DeCoste D, Weston J (eds) Large scale kernel machines. MIT Press. http://www.iro.umontreal.ca/~lisa/pointeurs/bengio+lecun_chapter2007.pdf. Accessed 14 Sept 2015
3. Bertsekas DP, Tsitsiklis JN (1996) Neuro-dynamic programming, 1st edn. Athena Scientific, Nashua
4. Biehl M, Caticha N, Riegler P (2009) Statistical mechanics of online learning. In: Similarity-based clustering, recent developments and biomedical applications [outcome of a Dagstuhl Seminar], pp 1–22. doi:10.1007/978-3-642-01805-3_1. http://dx.doi.org/10.1007/978-3-642-01805-3_1
5. Bishop CM (2006) Pattern recognition and machine learning (information science and statistics). Springer-Verlag New York Inc., Secaucus
6. Bousquet O, Boucheron S, Lugosi G (2003) Introduction to statistical learning theory. In: Bousquet O, von Luxburg U, Rätsch G (eds) Advanced lectures on machine learning, lecture notes in computer science, vol 3176. Springer, pp 169–207
7. Dietterich T, Domingos P, Getoor L, Muggleton S, Dolan R (2008) Structured machine learning: the next ten years. *Mach Learn* 73:3–23
8. Douglas R, Sejnowski T et al (2008) Future challenges for the science and engineering of learning. Technical report, final NSF workshop report
9. Fort J (2006) SOM's mathematics. *Neural Netw* 19(6–7):812–816. doi:10.1016/j.neunet.2006.05.025
10. Friedrichs F, Igel C (2005) Evolutionary tuning of multiple SVM parameters. *Neurocomputing* 64:107–117. doi:10.1016/j.neucom.2004.11.022
11. Hammer B (2012) Challenges in neural computation. *KI* 26(4):333–340. doi:10.1007/s13218-012-0209-0
12. Hammer B, Saunders C, Sperduti A (2005) Special issue on neural networks and kernel methods for structured domains. *Neural Netw* 18(8):1015–1018. doi:10.1016/j.neunet.2005.07.004
13. Lang T, Toussaint M (2010) Planning with noisy probabilistic relational rules. *J Artif Intell Res (JAIR)* 39:1–49. doi:10.1613/jair.3093
14. Langley P (2011) The changing science of machine learning. *Mach Learn* 82(3):275–279
15. Mitchell TM (2006) The discipline of machine learning. Machine Learning Department technical report CMU-ML-06-108, Carnegie Mellon University, Pittsburgh
16. Toussaint M, Lang T, Jetchev N (2013) Kognitive Robotik—Herausforderungen an unser Verständnis natürlicher Umgebungen. *Automatisierungstechnik* 61(4):259–268. doi:10.1524/auto.2013.0027