# Towards Lifelong Learning in Optimisation Algorithms

Emma Hart
*School of Computing, Edinburgh Napier University, U.K.*

Keywords: Optimisation, Lifelong Learning.

Abstract: Standard approaches to developing optimisation algorithms tend to involve selecting an algorithm and tuning it to work well on a large set of problem instances from the domain of interest. Once deployed, the algorithm remains static, failing to improve despite being exposed to a wealth of further example instances. Furthermore, if the characteristics of the instances being solved shift over time, the tuned algorithm is likely to perform poorly. To counter this, we propose the *lifelong learning optimiser*, which autonomously and continually refines its optimisation algorithm(s) to improve with experience, and generates novel algorithms if performance drops. The approach combines genetic programming with an autonomous management method inspired by the operation of the natural immune system.

## 1 INTRODUCTION

Optimisation is an important activity for many businesses, providing better, faster, cheaper solutions to problems in areas including scheduling of people and processes, routing of vehicles and packing of containers. Metaheuristic algorithms provide a pragmatic way to tackle optimisation, providing high-quality solutions in reasonable time. Unfortunately, selection and tuning of an appropriate algorithm can difficult, often requiring an expert to design the algorithm, a software engineer to implement it, and finally application of automated tuning processes to refine the chosen algorithm. This is not only costly, requiring significant human-effort, but also results in software which can quickly become obsolete when it no longer matches the goals of a company or if the characteristics of the optimisation problems being solved changed substantially. In addition, although deployed optimisation software is exposed to a continual stream of new problem instances, unlike human learners, it fails to exploit this information. As a result, it does not improve its performance with experience, therefore wasting valuable information.

To counter this, a paradigm shift in desigining optimisation algorithms is required. This article describes the *life-long learning optimiser (L2O)* which when faced with a continual stream of problems to optimise: (a) refines an existing set of algorithms so that they improve over time as they are exposed to more examples, and (b) automatically generates new algorithms when faced with problem instances that are completely different from those seen before. The approach is inspired by ideas from the operation of the natural immune system, which exhibits many properties of a life-long learning system that can be exploited computationally, and uses genetic programming to automatically generate new algorithms.

## 2 LIFE LONG LEARNING

(Silver et al., 2013) propose that the time is now ripe for the AI community in general to move beyond learning algorithms to more seriously consider the nature of systems that are capable of learning over a lifetime. They suggest that algorithms should be capable of learning a variety of tasks over an extended period of time such that the knowledge of how to solve tasks is retained, and can be used to improve learning in the future. They name such systems lifelong machine learning, or LML systems, in accord with earlier proposals by (Thrun and Pratt, 1997).

(Silver et al., 2013) identify three essential components of an LML system. Firstly, that it should be able to retain and/or consolidate knowledge, i.e. incorporate a long-term memory. Second, they suggest it should selectively transfer prior knowledge when learning new tasks, i.e exploit existing learned information in an efficient manner. Finally, they note that to achieve this efficiently, a systems approach that ensures the effective and efficient interaction of the elements of the system is required.

Hart E.
Towards Lifelong Learning in Optimisation Algorithms.
DOI: 10.5220/0006810500010001
In *Proceedings of the 9th International Joint Conference on Computational Intelligence (IJCCI 2017)*, pages 7-9
ISBN: 978-989-758-274-5

We remarked in (Sim et al., 2015) that the natural immune system provides a obvious metaphor for building a system that meets the requirements of a LML as noted by (Silver et al., 2013). It exhibits memory that enables it to respond rapidly when faced with pathogens it has previously been exposed to; it can selectively adapt prior knowledge via clonal selection mechanisms that can rapidly adapt and improve existing antibodies (pathogen-fighting cells) to cope better with new variants of previous pathogens and finally, it embodies a systemic approach by maintaining a repertoire of antibodies that collectively cover the space of potential pathogenic material.

In the human immune system, immune cells are generated from *gene libraries*: the DNA encoding for the cells is constructed by random sampling from so-called V, D and J gene libraries which gives rise a huge diversity of cells due to the combinatorics of the process. A huge advantage of this process is that a very large number of cells can be constructed from a fixed repertoire of DNA. Shifting the focus to optimisation, we propose that *genetic programming* can provide an analagous function: from a fixed set of terminals and functions, a very large space of *algorithms* can be generated, thus providing the diversity required to achieve lifelong learning.

# 3 NELLI: AN L2O

In previous work, (Hart and Sim, 2014; Sim et al., 2015), we have combined the immune metaphor with genetic programming in a system dubbed *NELLI: Network for Lifelong Learning*. The system has been applied in bin-packing and job-shop scheduling domains. NELLI autonomously generates an ensemble of optimisation algorithms that are capable of solving a broad range of problem instances from a given domain. The size of the ensemble varies over time depending on the stream of instances that the system is exposed to: each algorithm generalises over some region of the area of instance space defined by the problems of interest. It has been demonstrated to improve its performance as it is exposed to more and more instances from a given family of problems, and generate new algorithms when faced with instances that exhibit very different characteristics from those previously seen. Finally, it is also shown to retain *memory*, in that if re-exposed to instances seen in the past, it quickly returns new algorithms which exhibit good performance.

# 4 CONCLUSIONS

NELLI represents the first steps towards creating L2O systems — optimisers that continue to adapt over time. However much work can be done in improving the system. The human immune system adapts over two time scales. Over an individual lifetime, new cells are generated from gene libraries as described above, while the gene libraries themselves adapt on an evolutionary timescale across generations, therefore changing their content. There is no reason why the same process canot be applied to Genetic Programming, with the functions/terminals that make up the algorithm — or even the operations of the GP process itself — evolving over time.

Another direction for future work concerns the manner in which the system reacts to change in instance characteristics. The current approach relies on trial and error, with newly generated algorithms competing against each other to remain in the system. The integration of machine-learning approaches to predict likely changes in instances offers the potential to pregenerate algorithms in anticipation of future demand, thereby increasing the efficiency of the system. Some efforts towards this have been described by (Ortiz-Bayliss et al., 2015) in relation to solving constraint satisfaction problems.

In conclusion, we argue for a shift in direction for the optimisation community: rather than focusing effort on developing more and more complex algorithms trained on large but static sets of data, a move towards developing systems that autonomously and continually generate specialised algorithms *on-demand* may bear considerable fruit.

# REFERENCES


Hart, E. and Sim, K. (2014). On the life-long learning capabilities of a nelli*: A hyper-heuristic optimisation system. In *International Conference on Parallel Problem Solving from Nature*, pages 282–291. Springer.

Ortiz-Bayliss, J., Terashima-Marn, H., and Conant-Pablos, S. (2015). Lifelong learning selection hyper-heuristics for constraint satisfaction problems. In *Advances in Artificial Intelligence and Soft Computing*.

Silver, D., Yang, Q., and Li, L. (2013). Lifelong machine learning systems: Beyond learning algorithms. In *AAAI Spring Symposium Series*.

Sim, K., Hart, E., and Paechter, B. (2015). A lifelong learning hyper-heuristic method for bin packing. *Evolutionary computation*, 23(1):37–67.

Thrun, S. and Pratt, L. (1997). *Learning to Learn*. Kluwer Academic Publishers, Boston, MA.

## BRIEF BIOGRAPHY

Prof. Hart gained a 1st Class Honours Degree in Chemistry from the University of Oxford, followed by an MSc in Artificial Intelligence from the University of Edinburgh. Her PhD, also from the University of Edinburgh, explored the use of immunology as an inspiration for computing, examining a range of techniques applied to optimisation and data classification problems. She moved to Edinburgh Napier University in 2000 as a lecturer, and was promoted to a Chair in 2008 in Natural Computation. She is active worldwide in the field of Evolutionary Computation, for example as General Chair of PPSN 2016, and as a Track Chair at GECCO for several years. She has given keynotes at EURO 2016 and UKCI 2015, as well as invited talks and tutorials at many Universities and international conferences. She is Editor-in-Chief of Evolutionary Computation (MIT Press) from January 2016 and an elected member of the ACM SIGEVO Executive Board. She is also a member of the UK Operations Research Society Research Panel.