Examiner's Report

Thesis: Proactive Evolutionary Algorithms for Dynamic Optimization Problems

PhD Candidate: Patryk Filipiak

Summary of Work

This thesis develops a series of strategies for implementing *proactivity* in evolutionary algorithms (EA) in order to improve the utility of EA for application in dynamic environments. The strategies are tested on a number of continuous dynamic benchmark problems, including some problems with constraints, and an inverse kinematics problem defined in an environment with moving obstacles.

In the thesis the candidate provides an initial chapter which motivates the work. A clear contrast is drawn between reactive and proactive approaches to dealing with dynamic environments. Specifically, it is noted that proactive approaches may have utility in cases where an environment is 'predictable', i.e., where it is changing slowly or cyclically - in other words where consecutive states of the environment are correlated – but are not changing in a random or chaotic fashion.

Three distinct strategies are outlined briefly in the first chapter, *anticipation of evaluations* which uses an ARIMA model to forecast future fitness values at a set of locations in the search space, *anticipation of optima locations* which seeks to forecast the location of optimal fitness, and *anticipation of landscape changes* which uses a Markov chain prediction mechanism to predict future environmental 'states'. The chapter also provides an outline of the rest of the thesis. The only items missing from the chapter are a very clear statement of the research objective, underpinned by a scope limits section and also a related section indicating the literatures which are potentially relevant to this work – with a specific statement concerning which literatures this thesis draws most heavily from.

The second chapter provides a comprehensive, well structured, review of current state of the art in application of EA for dynamic optimisation. The review does a good job in covering the key issues from this literature in an integrated fashion and it is evident that the candidate has a deep understanding of the relevant literature. A discussion is also provided concerning the use of benchmark functions for dynamic optimisation. A related discussion is provided concerning the metrics that can be used to assess optimisation performance in dynamic environments with an important distinction being drawn between optimality-based metrics, behaviour-based metrics and prediction accuracy metrics.

Chapter 3 provides an introduction to time series forecasting and to the operationalisation of the three proactive strategies implemented in this thesis. The concept of population segmentation (the allocation of population across exploration, exploitation and anticipating behaviours) is also introduced.

The following three chapters implement the three strategies and test their performance on a range of benchmark functions. The results in each case are discussed with the proactive versions of the EA generally outperforming reactive versions of the same algorithm. Chapter 7 provides an application of the anticipation of evaluations strategy to the inverse kinematics problem (which is introduced in the chapter) and again, the results indicate that the proactive version of EA outperforms a reactive algorithm.

A final chapter provides conclusions and suggests future work. Two appendices are provided which cover a background to time series analysis and elements of Markov Chains.

The thesis is well-written with only a small number of phrasing errors / typos noted on its reading. Over the course of the thesis, the candidate displays a clear knowledge of the field of EA and of key issues relevant to application of these methods in dynamic environment. I enjoyed reading the thesis.

The candidate, as well as presenting the strengths of the proposed strategies for implementing proactivity, recognises the limitations and underlying assumptions of the approaches. For example, a practical problem identified by the candidate in the implementation of some of the methods in this thesis is the parametrising of the relevant prediction models and the implicit choice of memory length used in the proactive process. Another identified problem is that the three methods may struggle (respectively) in higher-dimensional settings, multimodal settings, and cases where the environment transitions though a large number of states. Advances in performance in these cases is left for future work.

Academic and Practical Significance of the Work

Many of the most challenging problems facing researchers and real-world decision-makers are those with a dynamic nature. That is, the environment in which the solution exists, and consequently the optimal solution itself, changes over time. Examples of dynamic problems include trading in financial markets, time series analysis of gene expression data, and routing in telecommunication networks, to name but a few. Biological organisms inhabit dynamic environments and mechanisms have arisen to promote the `survivability' of biological creatures in these environments. These mechanisms are useful sources of inspiration in helping us to design computer algorithms to attack real-world problems in dynamic environments.

This thesis contributes to a significant existing literature concerning how best to implement evolutionary algorithms in dynamic environments. Hence, the academic and practical importance of this research is evident.

As the candidate correctly identifies, most existing methods for dealing with changing environments work on a ‘reactive’ basis in that they try to ‘catch up’ with a change in the environment after it occurs. Hence, the approach taken in this work, attempting to proactively ‘anticipate’ change before it occurs, is particularly novel. This is of course a very challenging problem in its own right, hence, it is not reasonable to expect any thesis to ‘solve’ this problem and the contribution of this work must be viewed in this

light. The thesis opens up new approaches to dealing with dynamic environments and provides plentiful scope for follow on research.

The thesis has both academic and practical significance making contributions to both areas. Academic contributions of the study include:

(i) An up to date review of prior literature on dynamic optimisation
(ii) Development of three strategies for implementation of proactive evolutionary algorithms namely, anticipation of evaluations, anticipation of optima locations, and anticipation of landscape changes. Crucially, these strategies are quite general and can be applied to any EA (more generally, the underlying concepts can be applied to a wide range of metaheuristics)
(iii) Empirical testing of methods and drawing of conclusions based on results as to how best to implement the methods
(iv) Suggestions for follow on work on the implementation of proactive strategies

Publications from this Work

It is noted that a total of 13 conference publications (11 are mentioned in the draft thesis, two more are mentioned in the supporting documentation) have resulted from the research undertaken for this thesis. In general, the conferences are of excellent international standing, have strong peer review processes, and would be quite selective in terms of their acceptance rates. This level of publication productivity from PhD research is excellent and it provides a strong assurance as to the quality of the work. While acknowledging these strengths, the candidate is encouraged to consider advancing some of the research work undertaken to journal publications.

Suggested Areas for Improvement / Future Work

In chapter one, the thesis should more clearly state its objective, including scope limits on this. While it is evident that the work covers EA application in dynamic environments, the author notes limitations on the scope of the work undertaken across a number of chapters. It would be easier for the reader if this material was included in a paragraph in the first chapter.

In chapter 3, the author correctly notes that the performance of each of the three proactive strategies is a joint function of the strategy, the underlying EA implementation (and of course the environment). This makes it tricky to make conclusive statements concerning the utility of specific strategies as potentially a different implementation of an EA (e.g., with different selection pressure or different replacement strategy) could produce alternative baseline performance levels in any given dynamic environment. Ideally, the thesis should include a more detailed discussion of this point when analysing the results for the various experiments and consider how this impacts on what conclusions we can actually draw from the experiments.

In the results presented in chapter 4 in table 4.2 to 4.4 (and similar presentation of results in subsequent chapters) it is not clear whether the differences between the best

results obtained from the various proactive and reactive methods are statistically significant. While each experiment is repeated 50 times with the reported best result being averaged over these 50 experiments it is not explicitly stated in the chapter whether the differences found are statistically significant.

In the more computationally intensive algorithmic versions, it is not absolutely clear whether these perform better than the simpler algorithms. To undertake a fair test here, an equivalent level of computational resources should be allocated to the simpler algorithms with the consequent results being compared.

One other point that was not clear to me is the likely scalability of the results. While I may have missed it, I can not find a clear reference in chapter 4 or 5 as to the dimensionality of the problems addressed there. In chapter 6, it is noted that the experiments are undertaken on relatively low dimensional 5d and 10d versions of the five test problems considered. As the dimensionality increases, the potential complexity of the state space could likewise increase. Likewise, the problem addressed in chapter 7 is a dynamic 2-d constrained IK problem. Hence, the question concerning the scalability of the results appears to be open.

Conclusion

The thesis presents a significant body of work and an original contribution which is clearly at the doctoral level. I am satisfied to recommend that the PhD degree be awarded for this work.

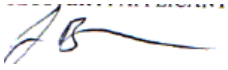

2 December 2016

Professor Anthony Brabazon
Associate Dean
Smurfit Graduate School of Business
University College Dublin Ireland