Evolutionary function approximation for reinforcement learning

Lucian Buşoniu

Consider the evolutionary strategies for finding reinforcement learning (RL) solutions, described by Whiteson and Stone (2006).[1] Write a brief paper about this topic, integrating answers to the following questions in the logical flow of your paper. Do not restrict your reading to this single article (Whiteson and Stone, 2006); e.g., some of the questions below may require further reading.

- Why must the Q-function be approximated in RL? Why is it better to derive the approximator automatically, rather than construct it in advance and keep it fixed during learning?
- In what way are evolutionary optimization and RL combined in this article? Which of the techniques proposed in the article concerns just evolutionary optimization, without involving RL?
- Consider the *online* evolutionary optimization techniques proposed in this article. Explain in what sense they are 'online'.
- What are some important obstacles that you can see to the application of the techniques described in the article to the control of real-life systems involving economic costs and expensive hardware?
- In Figures 6 and 7, what are the reasons for the oscillation of the performance curves of online NEAT and online NEAT+Q? What is the period of the oscillation? Why is the performance curve of *offline* NEAT+Q oscillatory, even though the performance curve of offline NEAT is not?

References

Whiteson, S. and Stone, P. (2006). Evolutionary function approximation for reinforcement learning. *Journal of Machine Learning Research*, 7:877–917.

[1] While reading (Whiteson and Stone, 2006), you can safely skip Sections 5.4–5.6, and Section 7 (although it is of course not required to skip them).