

Coursework 2: Understanding Lions *or* Map Learning

Joanna J. Bryson

March 2, 2011

1 Introduction

All of your courseworks are designed primarily to give you experience in developing intelligent control and/or cognitive systems. However, the course is also intended to give you experience and feedback in writing about research. To this end, you will be writing research reports of at least one full page but not much more than two, using exactly this format. Submissions should be in the format specified for Coursework 1.

2 Approach

There are two different ways to do Coursework 2, and they will have very different learning outcomes. No one course can cover everything, and the exam will allow you to select between alternative questions catering to both sets of experiences. Nevertheless, it is probably worth thinking about how you would approach each of these problems, even though you will only actively work on one.

2.1 Option 1

There are many different types of social simulation, but in keeping with the emphasis of ICCS, you will be doing a *spatial simulation*, where you are meant to be simulating some approximation of:

- real time,
- real space,
- real animal capacities, e.g. motion, perception, metabolism.

The basic task is to replicate the results reported by Barry and Dalrymple-Smith (2005). This means you should briefly review the scientific reason for the model as well as checking the model's outcomes through replicating them in NetLogo. If you are still confused about the idea of replication, you may want to read King (1995) or Bryson et al. (2007).

2.2 Option 2

Extend from Coursework 1 to help your robot learn a map of a space it circumnavigates, such that the more times it goes around the space, the fewer times it hits its bumpers. First, this will require having your robot make a noise every time it bumps into something, so you can record its progress per time around the enclosure. Second, this will require using time, sonar or possibly odometric information to keep track of where the robot thinks it is. Third, this will probably be facilitated by using vision to identify landmarks, so that the robot can have more information to reduce its own uncertainty.

In the worst case, you can just give the robot a map as Polly had (Horswill, 1993). However, what we'd like to see is whether you can get the LEGO robots to learn. If you took Pattern Matching last year, you might want to revise notes you have on HMM and POMDP and use a probabilistic representation, but Polly used a deterministic one so you may be able to as well (though Polly didn't learn.)

3 Results

As with Coursework 2, we would like you to write this as if it were an experimental paper, that is we would like to see you propose and support hypotheses. The type of results will depend on which option you pursue.

3.1 Option 1

NetLogo provides a tool for running experiments (BehaviourSpace) and for drawing figures. Thus it is pretty easy to get NetLogo to run experiments for you, if you can think of parameters you might like to vary that would be experimentally interesting.

With respect to your own results, if you describe a reasonably-well working system in a comprehensible manner you will pass. If you can get the basic replication and describe it appropriately in the report, you will get around 55. Getting a mark over 70 requires demonstrating insight, creativity and/or understanding that goes beyond the basics laid out for you in this document. For example, you might try to extend the results to look at strategies for other species, strategies for the prey, or to try to create a stable ecosystem. You will probably want to do some more reading up on the animals you are modelling to do any of these (and obviously cite that reading).

3.2 Option 2

Getting this working is likely to be tricky, so rather than worrying about running quantitative experiments, we are more interested this time in a detailed account of the algorithms you try and the results you attain. The main thing you will want to graph is just the reduction of the bumper hits — presumably over a number of trials (think about using error bars.)

With respect to your own results, if you describe a reasonably-well working system in a comprehensible manner you will pass. If you can get the basic system running and describe it appropriately in the report, you will get around 55. Getting a mark over 70 requires demonstrating insight, creativity and/or understanding that goes beyond the basics laid out for you in this document. For example, you might try to finding some learning algorithms or even libraries on line and adopting them to your problem (properly cited of course).

4 Discussion and Conclusions

These are per the instructions in Coursework 1.

References

Barry, A. and Dalrymple-Smith, H. (2005). Visual communication and social structure — the group predation of lions. In Bryson, J. J., Prescott, T. J., and Seth, A. K., editors, *Modelling Natural Action Selection: Proceedings of an International Workshop*, pages 146–151, Edinburgh. The Society for the Study of Artificial Intelligence and the Simulation of Behaviour (AISB).

- Bryson, J. J., Ando, Y., and Lehmann, H. (2007). Agent-based models as scientific methodology: A case study analysing primate social behaviour. *Philosophical Transactions of the Royal Society, B — Biology*, 362(1485):1685–1698.
- Horswill, I. D. (1993). *Specialization of Perceptual Processes*. PhD thesis, MIT, Department of EECS, Cambridge, MA.
- King, G. (1995). Replication, replication. *PS: Political Science and Politics*, XXVIII(3):443–499. with comments from nineteen authors and a response, “A Revised Proposal, Proposal,”.