Bridget Hallam November 26th, 2008

Good Experimental Papers in Robotics

Good papers are strong in several sections: experimental design, method reporting, results reporting, significance of results commented on, conclusions drawn. Primarily, ALL relevant FACTS must be reported; waffling and stupid statements (eg "the square on the screen moved in a lifelike manner") minimised; only VALID conclusions should be drawn; and alternative reasonable methods and possible explanations considered. With citations for everything, of course.

In particular, make it clear which phases were done in simulation and which on robotic hardware.

1 Experimental design

Experiments are designed to answer a question. Good experiments result from well-formulated questions. eg "Can a standard fork-lift truck with a standard webcam on top navigate in this warehouse using unmodified ceiling features?". By contrast, "I wonder what happens if we try this with a GA?" is not well formulated. It may be your real starting point, but it is not concrete enough to be suitable for underpinning a paper. Chose a question, design an experiment to answer this question, and report enough of your results to tell others not only the different answers that you found but also the sensitivities of your system to variations – in environment, program parameters, hardware components or whatever.

Having formulated a precise question to guide and drive your paper, you are ready to design your experiment. You are aiming to answer your question as unequivocally as possible. This means that you have to standardise as much as possible eg lighting, ground surface, battery levels, temperature ...(some of them at several different levels, to show how robust your system is); and consider alternative possibilities, eg "Why not simplify the navigation problem by using environmental markers/considering only vertical lines/using GPS signals/using an infra-red camera ...instead?"

As a minimum, you have to have a MEASURABLE way to show whether and to what extent the experiment "worked". So you need to decide what to measure, how to measure it, and what statistics to apply. If you don't have statistical training yourself, you will find your papers get much better if you consult with a statistician when your preliminary results are appearing, ie, BEFORE you think you have finished collecting data. The statistician can help you find variables which need controlling, and tell you how to make your results more significant. They may also help you to find a good "nul hypothesis" test, eg by running the robot with a random behaviour generator or the generally-accepted technique in place of the clever controller you designed, so you can PROVE that your system is better.

2 Introduction

Tell your readers what question you are asking and the main different ways this question could be answered. Outline the pros and cons of each (with refs) and explain why you chose your particular technique. The assumptions underpinning this technique should be clearly stated.

3 Method reporting

Here you state which variables you standardised to what values, and ideally why you chose these values. Also, which variables remained unstandardised. So maybe you did preliminary testing at random times, but always did your test runs at night with newly-charged batteries. If your preliminary results changed your experimental design, then this part of the design needs to be mentioned. eg "Preliminary tests showed that steering was adversely affected by damp weather, so all experiments took place when the humidity was under 50%."

This is also the right place to state how many times you repeated each phase of the experiment!

4 Results reporting

Presumably, your robot DID something, otherwise you would not be writing this paper. Good papers are precise and reasonably detailed about what the robot did. Bad papers are vague. Use measured values of anything relevant, with statistical analyses. You need to report results that do NOT support your conclusions, as well as those which do. A video of the robot successfully identifying and picking up a drinks can is a tremendous asset to your presentations, but for a scientific paper we need to know something about the success rates of each of the different phases and the ways in which failures occured. The video earns you a slot on YouTube, the analysis of success rates and failure modes gets you a serious scientific reputation. And enables others to try to imitate your system, which gets you citations and collaborations ...

Present detailed data with real error estimates and proper statistical analysis wherever appropriate.

5 Conclusions

You would like to be able to conclude by stating the answer to your question, but this is not always possible. Maybe the question was too big, or the system too sensitive to variables, or maybe your methodology doesn't seem to work in this instance. Whatever, you can still draw conclusions from your data. Just make sure that your conclusions are justifiable, otherwise serious researchers might despise you as unscientific. For example, concluding that your system CAN be used for its purpose is justifiable if it worked at all, but you need some disclaimers unless your system "works" most times.

Remember to point out places where more work would be beneficial. For example, maybe your system is sensitive to turbulence so you would want to make it more reliable and/or measure the limits within which the system is OK. Other future lines of research could also be mentioned here, eg trying the robot outdoors, or trying to reduce component prices.