

B. How to write a research paper

The final deliverable of a robotics class often is a write-up on a “research” project, modeled after research done in industry or academia. Roughly, there are three classes of papers:

1. Original research
2. Tutorial
3. Survey

The goal of this chapter is to provide guidelines on how to think about your project as a research project and how to report on your results as original research.

B.1. Original

Classically, a scientific paper follows the following organization:

1. Abstract
2. Introduction
3. Materials & Methods
4. Results
5. Discussion
6. Conclusion

The *abstract* summarizes your paper in a few sentences. What is the problem you want to solve, what is the method you are employing, what are you doing to assess your work, and what is the final outcome.

B. How to write a research paper

The *introduction* should describe the problem that you are solving and why it is important. A good guideline to write a good introduction are the Heilmeier questions:

1. What are you trying to do? Articulate your objectives using absolutely no jargon.
2. How is it done today, and what are the limits of current practice?
3. What's new in your approach and why do you think it will be successful?
4. Who cares?
5. If you're successful, what difference will it make?
6. What are the midterm and final "exams" to check for success?

Originally conceived for proposal writing by the head of DARPA, there are additional questions including "What will it cost?", "How long will it take?", and "What are the risks and pay-off", which are left out for the purpose of writing a research paper. In the context of scientific research, the question "What are you trying to do?" is best answered in the form of a *hypothesis*, see below.

The *materials & matters* section describes all the tools that you used to solve your problem, as well as your original contribution, e.g., an algorithm that you came up with. This section is hardly ever labeled as such, but might consist of a series of individual section describing the robotic platform you are using, the software packages, and flowcharts and descriptions on how your system works. Make sure you motivate your design choices using conclusive language or experimental data. Validating these design choices could be your first results.

The *results* section contains data or proofs on how to solve the problem you addressed or why it cannot be solved. It is important that your data is conclusive! You have to address concerns that your results are just a lucky coincidence. You therefore need to run multiple experiments and/or formally prove

the workings of your system either using language or math, see also Section A.3.

The *discussion* should address limitations of your approach, the conclusiveness of its results, and general concerns someone who reads your work might have. Put yourself in the role of an external reviewer who seeks to criticize your work. How could you have sabotaged your own experiment? What are the real hurdles that you still need to overcome for your solution to work in practice? Criticizing your own work does not weaken it, it makes it stronger! Not only does it become clear where its limitations are, it is also more clear where other people can step in.

The *conclusion* should summarize the contribution of your paper. It is a good place to outline potential future work for you and others to do. This future work should not be random stuff that you could possibly think about, but come out of your discussion and the remaining challenges that you describe there. Another way to think about is that the “future work” section of your conclusion summarizes your discussion.

It is important not to mix the different sections up. For example, your result section should exclusively focus on describing your observations and reporting on data, i.e., facts. Don't conjecture here why things came out as they are. You do this either in your hypothesis — the whole reason you conduct experiments in the first place — or in the discussion. Similarly, don't provide additional results in your discussion section.

Try to make the paper as accessible to as many reader styles and attention spans as possible. While this sounds impossible at first, a good way to address this is to think about multiple avenues a reader might take. For example, the reader should get a pretty comprehensive picture on what you do by just reading the abstract, just reading the introduction, or just reading all the figure captions. (Think about other avenues, every one you address makes your paper stronger.) It is often possible to provide this experience by adding short sentences that quickly recall the main hypothesis of your work. For example, when describing your robotic platform in the materials section, it does not hurt to introduce the section by something like “In order

B. How to write a research paper

to show that [the main hypothesis of our work], we selected...”. Similarly, you can try to read through your figure captions if they provide enough information to follow the paper and understand its main results on their own. Its not a problem to be repetitive in a scientific paper, stressing your one-sentence elevator pitch (or hypothesis, see below) throughout the paper is actually a good thing.

B.2. Hypothesis: Or, what do we learn from this work?

Classically, a hypothesis is a proposed explanation for an observed phenomenon. From this, the hypothesis has emerged as the corner stone of the scientific method and is a very efficient way to organize your thoughts and come up with a one sentence summary of your work. A proper formulation of your hypothesis should directly lead to the method that you have chosen to test your hypothesis. A good way to think about your hypothesis is “What do you want to learn?” or “What do we learn from this work?”.

It can be somewhat hard to actually frame your work into a single sentence, so what to do if a single hypothesis seems not to apply? One reason might be that you are actually trying to accomplish too many things. Can you really describe them all in depth in a 6-page document? If yes, maybe some are very minor compared to the others. If this is the case, they are either supportive of your main idea and can be rolled into this bigger piece of work or they are totally disconnected. If they are disconnected, leave them out for the sake of improving the conciseness of your main message. Finally, you might feel that you don’t have a main message, but consider all the things you done equally worthy, and despite answering the Heilmeier questions you cannot fill up more than three pages. In this case you might consider picking one of your approaches and dig deeper by comparing it with different methods.

Being able to come up with a one-sentence elevator pitch framed as a hypothesis will actually help you to set the scope of the work that you need to do for a research or class project.

How good do you need to implement, design or describe a certain component of your project? Well, good enough to follow through with your research objective.

B.3. Survey and Tutorial

The goal of a *survey* is to provide an overview over a body of work — potentially from different communities — and classify it into different categories. Doing this synthesis and establishing common language and formalism is the survey's main contribution. A survey following such an outline is a possible deliverable for an independent study or a PhD prelim, but it does not lend itself to describe your efforts on a focused research project. Rather, it might result from your involvement in a relatively new area in which you feel important connections between disjoint communities and common language have not been established.

A different category of survey critically examines concurring methods to solve a particular problem. For example, you might have set out to study manipulation, but got stuck in selecting the right sensor suite from the many available options. What sensor is actually best to accomplish a specific task? A survey which answers this question experimentally will follow the same structure as a research paper (see above).

A *tutorial* is closely related to a survey, but focuses more on explaining specific technical content, e.g, the workings of a specific class of algorithms or tool, commonly used in a community. A tutorial might be an appropriate way to describe your efforts in a research project, which can serve as illustration to explain the workings of a specific method you used.

B.4. Writing it up!

Writing a research report that contains equations, figures and references requires some tedious book-keeping. Although technically possible, word processing programs quickly reach their limitations and will lead to frustration. In the scientific community \LaTeX has emerged as a quasi standard for typesetting

B. How to write a research paper

research documentation. \LaTeX is a mark-up language that strictly divides function and layout. Rather than formatting individual items as bold, italic and the like, you mark them up as emphasized, section head etc, and specify how things look elsewhere. This is usually provided by a template provided by the publisher (or your own). While \LaTeX has quite a learning curve compared to other word processing software, it is quickly worth the effort as soon as you need to start worrying about references, figures or even indices.

Further Reading

- W. Strunk and E. White. *The Elements of Style* (4th Edition). Longan, 1999.
- T. Oetiker, H. Partl, I. Hyna and E. Schlegl. *The Not So Short Introduction to $\text{\LaTeX} 2_{\epsilon}$* . Available online.