Getting Started in Undergraduate Research[1]

Kelly A. Shaw
kshaw@richmond.edu

Getting involved in technical research can help students decide which track they want their career to take. It can also give them the experience and skills needed to pursue their desired careers. Working on research projects not only helps individuals become deeply knowledgeable about a specific problem, but it also helps students gain exposure to new tools and ways of thinking. Students who participate in research also develop their technical skills with respect to using tools, creating experiments, analyzing results, and improving their work. All of these skills make students attractive to prospective graduate schools, research advisers, and employers.

Given all of these great advantages of becoming involved in research, one wonders why all students do not participate. While it shares similarities to coursework, doing research is significantly different. There is no guaranteed correct answer to research problems – many approaches may solve a problem with varying success. There is also no guaranteed time frame for a solution to be reached. Not everyone enjoys this open-endedness and therefore may choose not to do research.

Students interested in trying out research must develop the skills needed to succeed in their endeavors. Like learning to play the guitar or learning a new language, the research process is something that can be learned with practice and guidance. This document is intended to serve as guidance for students just starting or thinking about getting started on research.

What is research?

Many people do not understand what it means to do research. They envision solitary, brilliant scientists who almost magically come up with novel solutions to existing problems. In reality, research is a very collaborative and evolutionary process. Scientists work on problems for years, gradually expanding their knowledge about a problem and the advantages and disadvantages of proposed solutions. They learn from other scientists' results through published work and discussions, and they combine their expertise and resources through research collaborations.

A more accurate way to think of research is as a continuous learning experience. It is a process of figuring out which problems are important enough that they need solutions, learning about a problem and existing solutions, coming up with a new and unique solution, evaluating that solution, learning its limitations and possibly discovering new problems or new solutions, and then repeating that process on those new problems or with those new solutions. As a consequence, perfect solutions and closure are extremely rare. This is different from course work where there usually is a correct solution. When you do research, you are basically exploring a series of solutions to find a better solution than existing solutions. Typically, someone will later find an even better approach, but it is likely that the earlier research will have created understanding that led to the subsequent, better approach.

Given the evolutionary nature of research, it is important to realize that even small results can have an important impact on the progression of the research field. Each new result adds to the community's overall understanding of the area, making future advances possible. Undergraduate students can and do produce valuable results that advance research on challenging problems.

 DOI:10.1109/MPOT.2012.2195788

Finding a research project and adviser

Once you have decided that you want to experiment with doing research, you need to determine when you plan on doing it, what topic you want to work on, and who you want to work with.

When to start research

When students start working on research is largely dependent on the culture in their departments and the interest level of individual students. Students frequently start during their third year of college or during the following summer. By then, students have an understanding of a variety of subareas in their field of study, and they have developed discipline specific skills that will be needed for research tasks. Because the students still have one more year of college remaining, there is sufficient time for them to delve deeply into a meaningful project and potentially achieve a valuable outcome or result. By starting research before their last year of college, it also gives research advisers sufficient time to evaluate students' skills in order to write recommendations or provide references during the fall.

Some departments encourage students to start working on projects as early as their first year of college or whenever a student first expresses interest. Research advisers who work with undergraduate students know to tailor the research tasks to the knowledge and skill level of the individual student. If you are interested in working on research, the key is to ask a research adviser about the possibilities. The person may start working with you immediately or may suggest that you do some reading or take a course before you do research together; once you have completed the suggested tasks, broach the subject with the adviser again.

Discovering research opportunities

Students' research experiences actually begin with them investigating the different research projects available for participation. For most students, this will involve learning about the research projects done by faculty at their university. One way to begin this exploration is to look at faculty web pages. Most of these web pages will mention the specific subfields (e.g. computer architecture) a faculty member works in. Additionally, they may include a list of recent publications. Reading over the introductions of these papers can help students determine if a faculty's research appeals to them. Do not be concerned if you cannot understand the entire paper – learning how to read research papers is part of learning how to do research.

Other ways to learn about research in your department are to enroll in classes which interest you, attend research group meetings, and attend department talks. All of these introduce you to important questions being pursued by researchers. Group meetings and classes expose you to what specific research is being done locally as professors frequently teach classes in their areas of specialty. Although speakers at department talks are usually from other institutions, it is likely the faculty host has an interest in the same research area. If there is no clear host, you can ask a faculty member which local faculty members do research in a similar area.

An alternate approach for getting involved in undergraduate research is to do a research internship away from your university. National internship programs exist that fund students to work at universities on research during the summer. One such program that exists for many different fields is the National Science Foundation (NSF) Research Experiences for Undergraduates (REU); information about REU sites and their research projects can be found on a web page maintained by NSF. Specific fields frequently have their own national internship programs that can be found by searching online using the research field (e.g. electrical engineering) and the word "internship" or by searching the field's national association website (e.g. IEEE). Programs also exist for underrepresented groups which can be searched for by including the terms "women" or "underrepresented". Application deadlines are frequently in early spring.

Choosing a research adviser

When choosing a research project, it is important to choose both a topic that interests you and to choose a research adviser whose interaction style will work well with your personality. Most people realize that choosing a topic that interests them is vital to their enjoyment and productivity. Finding a good research adviser match is equally important for happiness and research progress. Advisers' interaction styles vary considerably from very hands-off to micromanaging. Some faculty members find more time to dedicate to undergraduate researchers while other faculty members rely on more senior students to oversee and help new research students.

To make your experience as positive as possible, it is important to talk to students working with a prospective research adviser to learn about the research advisers' manner of undergraduate student interaction. Find out how often and for how long students meet with the adviser. Ask if the adviser responds to email questions quickly or if they prefer students stop by their office with questions. Discover whether or not there are individuals such as postdoctoral scholars or other students who are available to answer your questions. Once you have collected all of this information, you should reflect on whether or not you will feel comfortable and have the support you need if you work with this individual.

Once you have decided you are interested in working with a specific researcher, ask the individual for a meeting to discuss the possibility of working on research with them. To make a good impression, be prepared to talk about your interest in the research and all you have learned about this topic already including any questions you have. You should expect to talk about when you hope to do the research (i.e. school year or summer) and whether or not you are hoping to receive credit or be paid for the work. If you are seeking to be paid, you should bring information about university funding opportunities and associated deadlines to discuss.

If the individual agrees to work with you, be prepared to discuss expectations. Do not be surprised if the adviser suggests a trial period that lasts a semester or a summer. A trial period will allow both of you to evaluate your working relationship and enable you to determine if you truly enjoy working on this research; it will give you a good point at which to discuss whether or not you want to continue working together on this project.

Learning about the research project

When you join a research project, you will first need to familiarize yourself with the research problem you will be exploring and the tools you will need to use. This will involve finding and reading resources that describe the problem and existing approaches. It will also include discovering what tools already exist that you can incorporate into your research and how to use those tools.

Finding written resources

Your research adviser will likely suggest you read several papers or textbook chapters. These should be viewed as a starting point for you to learn all the information you need to know. To fully comprehend the big picture for your research problem, you will need to supplement this material with published resources you find on your own.

Textbooks can be a great resource for basic questions. Your adviser or other members of the research group may provide suggestions on textbooks to examine. Another approach is to find out which textbooks are used in classes at your university or in courses taught at other universities by faculty actively pursuing research in this area. A bonus feature of looking online for textbooks used in courses is that you may find lecture slides or other course materials that may be helpful.

Textbooks do not generally include all of the most recent research and they do not go into details of a large variety of research approaches. You will need to search for conference or journal publications for the most up-to-date and detailed information. There are a number of strategies you can use for finding related work. If you can find out what the top conferences and journals are for this topic, you can look at recent issues or proceedings from these sources for papers related to your research. You can also use online search engines like CiteSeer or Google Scholar, using keywords for this research. Once you have a set of relevant papers, you can also follow citations forward and backward. Meaning, you can look for papers that cite the paper you have already found (forward) and you can look at papers that were cited by the paper you have already found (backward). Finally, if you know specific researchers who have worked in this area previously, you can look at their web pages for other related work they may have done.

After finding papers related to your research, you need to discern which papers are actually closely related to your research and worthy of you spending time reading them. To do this, start by reading the abstract. If that seems promising, read the introduction and conclusion. Based on these sections, you should be able to discern how closely related this work is and whether or not it would be useful for you to read the entire paper. If the paper does not have a strong link to your work, discard it and move on. If the paper is somewhat related to your work, file it away to possibly read at a later date. You want to only read entire papers when they are strongly related to your research.

Reading technical papers

Reading technical papers can be overwhelming initially. The papers will include technical words and phrases that you have never encountered and that are not explained. There will be tools and techniques used that you have never heard of. It will also seem like an infinite amount of detail is crammed into a small number of pages. With practice, you learn that you do not need to understand or remember all of the details in most papers.

When reading a paper, your goal is to first understand the problem and solution presented at a high level, not every detail about the paper's implemented approach. Once you understand the paper's contribution, you can decide whether it is worth your looking at the paper in depth and can determine how much time you think the paper is worth investing. Depending on the contribution and its relative importance to the problem you are investigating, a paper may deserve your spending more or less time reading it. Some papers will only merit your skimming the introduction and conclusions so you can acquire the basic gist of the work. Other papers strongly related to your work may demand that you study them for eight hours to absorb as much detail as possible. With experience, you get better at deciding how much time to invest on each paper.

Because papers will initially contain words and phrases you do not know, it is a good idea to skim the paper for these unfamiliar words, particularly in the introduction and solution description. You can then look these words up in related papers, textbooks, or online resources and write down their definitions in the paper's margins for easy access when you read the paper. Once that is complete, read the paper. As you read, mark or highlight the key ideas and write down your questions and comments. After reading the paper, see if you can find the answers to the questions by searching online or in written sources or by asking another research group member.

After reading the paper, you should write a brief summary of the paper to insure that you understood the main concepts and contributions. Some example questions are:

1. What problem is being solved?
2. Who cares about the solution and why? How important is the problem?
3. What is the insight motivating the solution?
4. What is the proposed solution?
5. How effective is the solution and what are its limitations?
6. What insights/observations are made with respect to the solution's limitations? What future work does this suggest?

7. How is this approach similar and different from your approach?

For more information on reading technical papers, see William Griswold's "How to Read An Engineering Research Paper".

Familiarizing yourself with existing infrastructure

It is likely that your research will involve using tools or other infrastructure developed by someone else. Familiarizing yourself with the features of these tools and learning how to use them will be something you will want to start doing at the same time you familiarize yourself with the technical literature.

Most existing tools or infrastructure come with some form of documentation (e.g. README) that provides a high level overview. Start by reading that to obtain an understanding of the resource's overall structure. It might be useful to create a diagram of all the subparts, describing briefly what each subpart does or represents and the relationships between the subparts. You should then be able to prioritize which of the tool's subparts will be relevant to your work. Once you have identified the subparts relevant to your project, begin examining those subparts in greater detail. Like technical papers, some subparts you will need to study in great detail because you will be using or modifying them while others you may only need a passing knowledge of what they do. Try to only spend time on the subparts relevant to your work.

Familiarizing yourself with the infrastructure will be much easier if you learn how to quickly search whatever system you are on for files and text. Searching will allow you to quickly find information about different subparts and uncover the relationships between the different subparts. Information about a subpart may be located unexpectedly in the documentation for another subpart's documentation; it would be hard to discover this information without searching or reading all of the documentation. On Unix systems, *find* and *grep* are two useful commands for searching. An Internet search should provide the names of the tools to use on your system.

Creating experiments and analyzing results

Entire books have been written on how to create a set of experiments to test a hypothesis. No one expects that you will be an expert at designing experiments and analyzing results when you first start doing research. Over time, your skill at doing both will improve as you are exposed to more research.

Designing experiments

When creating experiments, you want to answer the following questions:

1. What inputs will be used?
2. How will you configure your solution for each test?
3. What numbers will you collect in order to analyze the effectiveness of your approach and support or refute your hypothesis?

The inputs you choose will likely be determined by your research area. Many research areas have standardized inputs which are used by researchers in order to provide comparisons between different research solutions. Consequently, you may have limited choices. Looking at related work can help you discover accepted inputs for your field.

Depending on the type of research you are doing, there may be a single factor that you can configure or there may be many configurable factors where each factor can be assigned a set of values. Depending on the number of factors you are configuring and the number of values each of those factors can take, it may be infeasible to try every possible combination of the factors and values. Consequently, when designing experiments, you want to vary the smallest number of factors that allow you to still test your approach. You also want to limit the number of values you choose for each factor, and you want those selected values to make sense as a collection. For example, you may have a value that varies from 1 to 8 by powers of 2 or a factor that you increase by orders of magnitude. The goal is to keep your increments consistent so you can detect trends in collected results.

When collecting results, you want to collect not only numbers which represent the end result of you research (e.g. reducing execution time), but you also want to collect numbers that will help explain or support your argument about how your solution is achieving that improved end result. For example, your end goal may be to reduce the execution time of an application, but since your proposed solution is impacting the computer's memory system you also want to collect information about the number of memory requests, how long it took memory requests to complete, or how long requests had to wait for a highly used resource.

Analyzing results

After you have constructed your experimental infrastructure, you want to insure it is generating valid results before you try to start making conclusions from your data. It can be useful to create a set of small tests where you can determine exactly what the results should be. Each of these tests should be as simple as possible and still test the desired condition. Confirming correctness on these small tests cases will increase your confidence in the correctness of your results on real inputs.

In addition to using small tests to verify the correctness of your approach, it can also be helpful to perform some sanity checking of your results. This can be as simple as creating back-of-the-envelope estimates of what your results should be and comparing your experimental results to these estimates. An additional way of doing this is to create two oracles – one which calculates worst case results when the proposed solution is not helpful and one that calculates best case results when the proposed solution is maximally helpful. Any experimental results you collect should fall between these two extremes or it is likely an error has been made in the experimental setup.

Once you have checked that your results are in the appropriate range, you will want to visualize your results in order to make patterns easy to discern. When you examine these results, it is important to pay attention to not only general trends that coincide with your expectations but also anomalous results. If a result is strongly positive or negative, you need to understand why. It could indicate that your approach works particularly well in some situations or poorly in other situations. Both are really interesting observations. Anomalous results may also indicate an error in the experimental framework. Either way, you will want to investigate and explain what causes these outliers.

Tips

Finally, this series of practical tips will hopefully help make your research experience enjoyable and efficient.

- Organize all of your research materials (e.g. notes, papers, documents) into a single space. Researchers commonly maintain a physical research notebook for this purpose, but you can determine what works best for you. You might use a blog or a collection of files in a directory on your computer.
- When talking to your adviser, plan ahead. Determine what questions or topics you want to discuss before a meeting and follow up with those issues after the meeting if you do not have enough time to discuss them all. Also, make sure to analyze any results you have before your meeting – raw data is nice, but data that has been analyzed for trends and problems is even better.
- Make time in your schedule for research. During the academic year, it may be difficult to "find" time for doing research. You should "make" time for research, adding it to your weekly schedule. Also, spread your research time throughout the week so you have time to get help before a deadline or meeting.

- New researchers frequently have a hard time figuring out when they know enough to move forward. It can be hard to start an experiment if you do not feel like you're an expert on the research topic or if you have not learned all the facets of a tool. There will always be more to learn, so it is important for you to be willing to start even when you are not an expert. You can keep learning while you start doing research. Be willing to trust your adviser about when you know enough to proceed.
- Finally, research is all about the unknown which can be pretty daunting. You do not know if you will succeed, you are not sure what the next step is, you cannot talk to anyone who has done this before. It can be downright scary. It is important to acknowledge this truth and then to acknowledge that some of your ideas will fail and that is normal. When something does not work as you thought, you have just learned something important! That knowledge will help improve your research by changing your approach or allowing you to discover a different problem that needs to be solved. Don't worry about failure; just keep trying.

Summary

Research can be a lot of fun and can help you develop skills you can use in whatever career path you choose. This article has presented some helpful hints for getting started in research to help your experience begin as smoothly as possible. You will learn much more about the research process once you get started. The key is to give research a try and to stay enthusiastic and open to both positive and negative results.