



# HOW TO SUCCEED IN GRADUATE SCHOOL: A GUIDE FOR STUDENTS AND ADVISORS

## PART 1 OF 2

by Marie desJardins

**T**his paper attempts to raise some issues that are important for graduate students to be successful and to get as much out of the process as possible, and for advisors who wish to help their students be successful. The intent is not to provide prescriptive advice—no formulas for finishing a thesis or twelve-step programs for becoming a better advisor are given—but to raise awareness on both sides of the advisor-student relationship as to what the expectations are and should be for this relationship, what a graduate student should expect to accomplish, common problems, and where to go if the advisor is not forthcoming.

### Introduction

This article originated with a discussion I had with several female professors about the problems women face in graduate school, and how more women could be encouraged to go to graduate school in computer science. Eventually, the conversation turned to the question of what these professors could do in their interactions with female students to support and encourage them. I volunteered that over the course of my graduate career I had collected a variety of papers and email discussions about how to be a good advisor, how to get through graduate school, and issues facing women. They were eager to get this material, and I told them I would sort through it when I got a chance.

After mentioning this project to a number of people, both graduate students and faculty—all of whom expressed an interest in anything I could give them—I realized two things: first, the issues that we were talking about really were not just women's issues but were of interest to all graduate students, and to all caring advisors. Second, in order to disseminate the information I had collected (and was starting to collect from others) it seemed to make more sense to compile a bibliography, and write a paper that would summarize the most useful advice and suggestions I had collected.

I solicited input from friends and colleagues via mailing lists and Internet bulletin boards, and collected almost an overwhelming amount of information. Sorting through it and attempting to distill the collective wisdom of dozens of articles and hundreds of email messages has not been an easy task, but I hope that the results provide a useful resource for graduate students and advisors alike. The advice I give here is directed towards PhD students in computer science and their advisors, since that is my background; but I believe that much of it applies to graduate students in other areas as well.

In my experience, the two main things that make graduate school hard are the unstructured nature of the process, and the lack of information about what you should spend your time on. I hope that this article will provide information for both graduate students and advisors that will help make the process less painful. I want to emphasize that graduate school is not easy, and these suggestions will not always be easy or even possible to follow (and they may not even be the ideal goal for you, personally, to strive for). You should not let that discourage you: start small, think big, and keep yourself focused on your ultimate goal, which should not just be to get through graduate school, but to

enjoy yourself, make progress towards being able to do what you want to do with your life, and learn something in the process.

I owe a debt of gratitude to David Chapman, whose paper [2] was an invaluable reference for me not only during the writing of this article, but during graduate school as well.

The goals of this article are to raise awareness of the need for a healthy and interactive graduate student-advisor relationship, to provide pointers and guidance for both advisors and graduate students in navigating the maze of a doctoral degree, and to give references and resources for those who hope to learn more.

### Before You Start

Many headaches can be avoided by doing some advance planning. First, why go to graduate school at all? The usual reasons given are that a PhD is required or preferred for some jobs, especially research and academic positions; that it gives you a chance to learn a great deal about a specific area; and that it provides an opportunity to develop ideas and perform original research. Wanting to delay your job hunt is probably not a good enough reason. Over the past decade, research and academic positions have become more difficult to find, and many recent PhDs end up “killing time” in a series of postdoctoral positions, or taking nonresearch jobs. Having a PhD is not a guarantee of finding a better job in and of itself! In addition, graduate school is a lot of work and requires strong motivation and focus. You have to really want to be there to make it through.

It helps to have a good idea of what area you want to specialize in, and preferably a couple of particular research projects you might like to work on, although many graduate students change their minds about research projects and even specialization field after they start school. Look for books and current journals and conference proceedings in your area, and read through them to get an idea of who is doing what where. You will be doing a *lot* of reading once you start graduate school, so you might as well get used to it. This is where advisors first enter the scene: faculty members ought to be willing to talk to undergraduates and help them find out more about research areas and graduate schools. Try to get involved in research: ask professors and TAs (teaching assistants) whether they need someone to work on an ongoing project, or start an independent research project with guidance from a faculty member.

Contact faculty members and graduate students at the schools you are interested in. Tell them about your background and interests and ask them what research projects they are working on. A good way to do this is email, which is much easier and quicker to respond to than a paper letter. A good advisor will be willing to answer these kinds of inquiries, although if they are busy they may give you only a brief answer or point you towards a graduate student. You will have to use your intuition to decide whether they are brushing you off or just busy. If you cannot get any answer at all, consider that that individual might not end up being a very accessible advisor. Asking these questions will help you narrow down your choices and may increase your chances of admission if the professors you contact become interested in working with you.

Your best bet is to find a school where there are at least two faculty members you would be interested in working with. That way, if one does not work out, or is too busy to take on a new student, you have a fallback position. Breadth of the graduate program (i.e., high-quality faculty in a broad range of subareas) is also a good thing to look for in a school, especially if you are not entirely certain what you want to specialize in.

It is also important to most people to feel comfortable with the community of graduate students. It pays to talk to some of the graduate students (both junior and senior) to find out how they like it, which advisors are good, and what kinds of support (financial and psychological) are available. Because there are so many students applying to each school, even highly qualified applicants are often rejected. You should apply to a range of programs—and do not take it personally if you do get rejected by some of them.

You can increase your chances of getting into graduate school by developing good relationships with your professors and work managers (this is very important for getting good recommendations), working on a research project, having a clear sense of what you want to work on (although it is always all right to change your mind later), having a broad background in your field and in related fields (for example, psychology classes are useful for AI students), getting good grades (especially in upper division classes in your area of interest), and getting a high score on the GRE if required. Also, it is a good idea to start thinking early about sources of funding: apply for an NSF fellowship, for example.

## Doing Research

For many new graduate students, graduate school is unlike anything else they have done. Sometimes it is hard to know exactly what it is you are supposed to be learning. Yes, you have to complete a dissertation, but how do you start? What should you spend your time doing?

Graduate school is a very unstructured environment in most cases. Graduate students typically take nine hours or less of coursework per semester, especially after the second year. For many, the third year—after coursework is largely finished and preliminary exams have been completed—is a very difficult and stressful period. This is when you are supposed to find a thesis topic, if you are not one of the lucky few who has already found one. Once you do find a topic, you can expect two or more years until completion, with very few landmarks or milestones in sight.

The following sections talk about the day-to-day process of doing research, criticism and feedback, working on the thesis, and financial support for research.

## The Daily Grind

Being a good researcher involves more than “merely” coming up with brilliant ideas and implementing them. Most researchers spend the majority of their time reading papers, discussing ideas with colleagues,

writing and revising papers, staring blankly into space—and, of course, having brilliant ideas and implementing them.

Part 2 of this article discusses the process and importance of becoming part of a larger research community, which is a critical aspect of being a successful researcher. This section contains ideas on keeping track of where you are going, and where you have been, with your research, staying motivated, and how to spend your time wisely.

Keeping a journal of your research activities and ideas is very useful. Write down speculations, interesting problems, possible solutions, random ideas, references to look up, notes on papers you have read, outlines of papers to write, and interesting quotes. Read back through it periodically. You will notice that the bits of random thoughts start to come together and form a pattern, often turning into a research project or even a thesis topic. I was surprised, looking back through my journal as I was finishing up my thesis, how early and often similar ideas had cropped up in my thinking, and how they gradually evolved into a dissertation.

You will have to read a lot of technical papers to become familiar with any field, and to stay current once you have caught up. You may find yourself spending over half of your time reading, especially at the beginning. This is normal. It is also normal to be overwhelmed by the amount of reading you think you “should” do. Try to remember that it is impossible to read everything that might be relevant: instead, read selectively. When you first start reading up on a new field, ask your advisor or a fellow student what the most useful journals and conference proceedings are in your field, and ask for a list of seminal or “classic” papers that you should definitely read. For AI researchers, a useful (if slightly outdated) starting point is Agre’s [1] summary of basic AI references. Similar documents may exist for other research areas—ask around, and cruise the Internet. Start with these papers and the last few years of journals and proceedings.

Before bothering to read *any* paper, make sure it is worth it. Scan the title, then the abstract, then—if you have not completely lost interest already—glance at the introduction and conclusions. (Of course, if your advisor tells you that this is an important paper, skip this preliminary step and jump right in!) Before you try to get all of the nitty-gritty details of the paper, skim the whole thing, and try to get a feel for the most important points. If it still seems worthwhile and relevant, go back and read the whole thing. Many people find it useful to take notes while they read. Even if you do not go back later and reread them, it helps to focus your attention and forces you to summarize as you read. And if you do need to refresh your memory later, rereading your notes is much easier and faster than reading the whole paper.

A few other points to keep in mind as you read and evaluate papers:

- Make sure the ideas described really worked (as opposed to just being theoretically valid, or tested on a few toy examples).
- Try to get past buzzwords: they may sound good, but not mean much. Is there substance and an interesting idea underneath the jargon?
- To really understand a paper, you have to understand the motivations for the problem posed, the choices made in finding a solution, the assumptions behind the solution, whether the assumptions are realistic and whether they can be removed without invalidating the approach, future directions for research, what was actually accomplished or implemented, the validity (or lack thereof) of the theoretical justifications or empirical demonstrations, and the potential for extending and scaling the algorithm up.

Keep the papers you read filed away so you can find them again later, and set up an online bibliography (BibTeX is a popular format, but any-

thing consistent will do). I find it useful to add extra fields for keywords, the location of the paper (if you borrowed the reference from the library or a friend), and a short summary of particularly interesting papers. This bibliography will be useful for later reference, for writing your dissertation, and for sharing with other graduate students (and eventually, perhaps, advisees).

### Staying Motivated

At times, particularly in the “middle years,” it can be very hard to maintain a positive attitude and stay motivated. Many graduate students suffer from insecurity, anxiety, and even boredom. First of all, realize that these are normal feelings. Try to find a sympathetic ear—another graduate student, your advisor, or a friend outside of school. Next, try to identify why you are having trouble and identify concrete steps that you can take to improve the situation. To stay focused and motivated, it often helps to have organized activities to force you to manage your time and to do something every day. Setting up regular meetings with your advisor, attending seminars, or even extracurricular activities such as sports or music can help you to maintain a regular schedule.

Chapman [2] enumerates a number of “immobilizing shoulds” that can make you feel so guilty and unworthy that you stop making progress. Telling yourself that you *should* have a great topic, that you *should* finish in *n* years, that you *should* work 4, or 8, or 12 hours a day is not helpful for most people. Be realistic about what you can accomplish, and try to concentrate on giving yourself positive feedback for tasks you do complete, instead of negative feedback for those you do not.

Setting daily, weekly, and monthly goals is a good idea, and works even better if you use a “buddy system” where you and another student meet at regular intervals to review your progress. Try to find people to work with: doing research is much easier if you have someone to bounce ideas off of and to give you feedback.

Breaking down any project into smaller pieces is always a good tactic when things seem unmanageable. At the highest level, doing a master’s project before diving into a PhD dissertation is generally a good idea (and is mandatory at some schools). A master’s gives you a chance to learn more about an area, do a smaller research project, and establish working relationships with your advisor and fellow students.

The divide-and-conquer strategy works on a day-to-day level as well. Instead of writing an entire thesis, focus on the goal of writing a chapter, section, or outline. Instead of implementing a large system, break off pieces and implement one module at a time. Identify tasks that you can do in an hour or less; then you can come up with a realistic daily schedule. If you have doubts, do not let them stop you from accomplishing something—take it one day at a time. Remember, every task you complete gets you closer to finishing. Even if you do not make any obvious progress, you will have learned something, although it may be “never waste your time on this task again!”

### Getting to the Thesis

The hardest part of getting a PhD is, of course, writing the dissertation. The process of finding a thesis topic, doing the research, and writing the thesis is different from anything most students have done before. If you have a good advisor and support network, you will be able to get advice and help in setting directions and goals. If not, you may need to be more independent. If this is the case, do not isolate yourself from the world: try to go out and find the resources and support you need from professors, other graduate students, mailing lists, friends, family, and publications like this one.

### Finding an Advisor

Finding the right advisor can help you immeasurably in successfully completing a thesis. You should ideally have selected the schools you applied to by identifying faculty members you would like to work with. If not, start looking around as early as possible. Of course, the ideal advisor will be in the area you are interested in working in, will actively be doing high-quality research and be involved in and respected by the research community, and (not least) will be someone you can get along with.

Read research summaries by faculty members, which are usually published by the department, go to talks they give, and attend or audit courses given by professors you might be interested in working with. Talk to other graduate students and recent graduates. Ask them how their relationships with their advisors are/were, how quickly the advisor’s students graduate, and how successful (well recognized, high-quality) their research is. What kinds of relationships do they have—frequent interactions, collaborative work, encouraging independence, handing out topics or helping students to create individual research areas, or a more hands-off style?

Other things to find out about potential advisors:

- What is the average time their PhD students take to finish their degrees? What is the dropout rate for their students?
- How long have they been on the faculty? There are advantages and disadvantages to being one of the first members of a new research group. On the positive side, you often have more freedom to choose your research topic and to influence the direction of the group’s research. On the negative side, you may be more isolated (since there will not be older graduate students in the group), your advisor will not have as much experience, and if they do not get tenure you may be scrambling for a new advisor several years into your thesis.

A good advisor will serve as a mentor as well as a source of technical assistance. A mentor should provide, or help you to find, the resources you need (financial, equipment, and psychological support); introduce you and promote your work to important people in your field; encourage your own interests, rather than promoting their own; be available to give you advice on the direction of your thesis and your career; and help you to find a job when you finish. They should help you to set and achieve long-term and short-term goals.

Once you identify one or more potential advisors, get to know them. Introduce yourself and describe the area you are interested in. Attend their research group meetings if they hold them regularly. Give them a copy of a research proposal if you have a good idea of what you want to work on, and ask for comments. Ask whether they have any TA or RA (research assistant) positions available, or if there are any ongoing research projects that you could get involved with. Read their published papers, and the work of their students. Drop by during office hours and ask questions or make comments. Offer to read drafts of papers—and do more than just proofread.

The type of relationship that each student needs with an advisor will be different. Some students prefer to be given more direction, to have frequent contact, and to be “checked up on.” Others are more independent. Some may need contact but be self-conscious about asking for it. Other things that vary include what kinds of feedback is preferred (lots of “random” ideas vs. very directed feedback (pointers)), working individually vs. in groups, working on an established research project vs. a new, independent effort; working in the same area as your advisor or doing an “outside” thesis.

You may find that your thesis advisor does not always give you all of the mentoring that you need. Multiple mentors are common and use-

ful; they may include other faculty members in your department or elsewhere, senior graduate students, or other colleagues. You may want to seriously consider changing thesis advisors if your advisor is inaccessible or disinterested, gives you only negative feedback, does not have the technical background to advise you on your thesis, or harasses you.

The most important thing is to ask for (i.e., demand politely) what you need.

### Finding a Thesis Topic

Doing a master's project is often a good idea (and is required by some schools). Although choosing an appropriately scaled-down topic may be difficult, having the ideal topic is also less important, since you will have the chance to move on after only a year or so. If you have a good idea of what you want to do your PhD dissertation on, choosing a master's project that will lead into the dissertation is wise: you will get a head start on the PhD, or may decide that you are not interested in pursuing the topic after all, thereby saving yourself a lot of work and grief farther down the road.

A good source of ideas for master's projects, and sometimes for dissertation topics, is the future work section of papers you are interested in. Try developing and implementing an extension to an existing system or technique.

Generally speaking, a good PhD thesis topic is interesting to you, to your advisor, and to the research community. As with many aspects of graduate school, the balance you find will depend at least in part on the relationship you have with your advisor. Some professors have well defined long-term research programs and expect their students to contribute directly to this program. Others have much looser, but still related ongoing projects. Still others will take on anyone with an interesting idea, and may have a broad range of interesting ideas to offer their students. Be wary of the advisor who seems willing to let you pursue any research direction at all. You probably will not get the technical support you need, and they may lose interest in you when the next graduate student with a neat idea comes along.

If you pick a topic that you are not truly interested in simply because it is your advisor's pet area, it will be difficult to stay focused and motivated—and you may be left hanging if your advisor moves on to a different research area before you finish. The same is true for choosing a topic because of its marketability: if you are not personally excited about the topic, you'll have a harder time finishing and a harder time convincing other people that your research is interesting. Besides, markets change more quickly than most people finish dissertations.

In order to do original research, you must be aware of ongoing research in your field. Most students spend up to a year reading and studying current research to identify important open problems. However, you will never be able to read everything that might be relevant—and new work is always being published.

Try to become aware and stay aware of directly related research—but if you see new work that seems to be doing exactly what you are working on, do not panic. It is common for graduate students to see a related piece of work and think that their topic is ruined. If this happens to you, reread the paper several times to get a good understanding of what they have really accomplished. Show the paper to your advisor or someone else who is familiar with your topic and whose opinions you respect. Introduce yourself to the author at a conference or by email, and tell them about your work. By starting a dialog, you will usually find that their work is not quite the same, and that there are still directions open to you. You may even end up collaborating with

them. Good researchers welcome the opportunity to interact and collaborate with someone who is interested in the same problems they are.

To finish quickly, it is usually best to pick a narrow, well defined topic. The downside of this approach is that it may not be as exciting to you or to the research community. If you are more of a risk-taker, choose a topic that branches out in a new direction. The danger here is that it can be difficult to carefully define the problem, and to evaluate the solution you develop. If you have a topic like this, it helps a lot to have an advisor or mentor who is good at helping you to focus and who can help you maintain a reasonably rigorous approach to the problem.

In the extreme case, if your topic is so out of the ordinary that it is unrelated to anything else, you may have difficulty convincing people of its worth. Truly innovative research is, of course, exciting and often pays back in recognition from the research community—or you could just be out in left field. If you have a far-out topic, be sure that people are actually *interested* in it, or you will never be able to “sell” it later, and will probably have trouble getting your work published and finding a job. In addition, it will be hard to find colleagues who are interested in the same problems and who can give you advice and feedback.

In any case, a good topic will address important issues. You should be trying to solve a real problem, not a toy problem (or worse yet, no problem at all); you should have solid theoretical work, good empirical results or, preferably, both; and the topic will be connected to—but not be a simple variation on or extension of—existing research. It will also be significant yet manageable. Finding the right size problem can be difficult. One good way of identifying the right size is to read other dissertations. It is also useful to have what Chapman [2] calls a “telescoping organization”—a central problem that is solvable and acceptable, with extensions and additions that are “successively riskier and that will make the thesis more exciting.” If the gee-whiz additions fail to pan out, you will still have a solid result.

A good way to focus on a topic is to write one-sentence and one-paragraph descriptions of the *problem* you want to address, and do the same for your proposed *solution*; then write an outline of what a thesis that solved this problem would look like (i.e., what chapters would be included, or if you are ambitious, what sections in each chapter).

Sometimes finding a small problem to work on and building on it in a “bottom up” fashion can work equally well, as long as you do not fall into the trap of solving lots of small unrelated problems that do not lead to a coherent, solid, substantial piece of research (i.e., a thesis).

Remember that a thesis is only a few years of your work, and that, if all goes well, your research career will continue for another 30 or 40. Do not be afraid to leave part of the problem for future work, and do not compare yourself to senior researchers who have years of work and publications to show for it. On the other hand, if you identify too much future work, your thesis will not look very exciting by comparison. Graduate students often pick overly ambitious topics and, in theory, your advisor will help you to identify a realistic size problem. Do not overestimate what other people have done. Learn to read between the lines of grandiose claims (something else a good advisor will help you to do).

Some schools may require that you write a thesis proposal. Even if they do not, this is a good first step to take. It forces you to define the problem, outline possible solutions, and identify evaluation criteria; and it will help you to get useful feedback from your advisor and other colleagues. Writing a good thesis proposal will take up to several months, depending on how much background work and thinking you have already done in the process of choosing the topic.

The proposal should provide a foundation for the thesis. First, you must circumscribe the problem and argue convincingly that it needs to

be solved, and that you have a methodology for solving it. You must identify and discuss related work: has this problem been addressed before? What are the shortcomings of existing work in the area, and how will your approach differ from and be an improvement over these methods?

Present your ideas for solving the problem in as much detail as possible, and give a detailed plan of the remaining research to be done. The proposal should include, or be structured as, a rough outline of the thesis itself. In fact, unless your final topic differs significantly from your proposed topic (which many do), you may be able to reuse parts of the proposal in the thesis.

You will probably have to take an oral exam in which you present and/or answer questions about your proposal. Be sure that your committee members are as familiar as possible with your work beforehand. Give them copies of the proposal, and talk to them about it. During the exam, do not panic if you do not know the answer to a question. Simply say, "I am not sure" and then do your best to analyze the question and present possible answers. Your examining committee wants to see your analytical skills, not just hear canned answers to questions you were expecting. Give a practice talk to other students and faculty members. Remember: you know more about your thesis topic than your committee; you are teaching *them* something for a change.

### Writing the Thesis

Graduate students often think that the thesis happens in two distinct phases: doing the research, and writing the thesis. This may be the case for some students, but more often, these phases overlap and interact with one another. Sometimes it is difficult to formalize an idea well enough to test and prove it until you have written it up; the results of your tests often require you to make changes that mean that you have to go back and rewrite parts of the thesis; and the process of developing and testing your ideas is almost never complete (there is always more that you *could* do) so that many graduate students end up "doing research" right up until the day or two before the thesis is turned in.

The divide-and-conquer approach works as well for writing as it does for research. A problem that many graduate students face is that their only goal seems to be "finish the thesis." It is essential that you break this down into manageable stages, both in terms of doing the research and when writing the thesis. Tasks that you can finish in a week, a day, or even as little as half an hour are much more realistic goals. Try to come up with a range of tasks, both in terms of duration and difficulty. That way, on days when you feel energetic and enthusiastic, you can sink your teeth into a solid problem, but on days when you are run-down and unmotivated, you can do some easier tasks.

### Getting Financial Support

Most graduate students (at least in the natural sciences) have a source of financial support that pays their tuition and a small living stipend. Although nobody ever got rich being a graduate student, you probably will not starve either. Sources of funding include fellowships (from NSF, universities, foundations, government agencies, and industry), employer support, research assistantships (i.e., money from a faculty member's research grant) and teaching assistantships. Kantrowitz [3] provides an extensive list of funding sources for math, science, and engineering graduate students.

Start looking for money early. Many schools arrange support in the form of an RA or TA position in the first year, but after that, you are on your own. Deadlines for applications vary, and if you miss one, you will probably have to wait another year. After you apply, it can take six

months or so to review the applications and several more months to actually start receiving money.

Ask faculty members (especially your advisor, who should be helping you to find support or providing support out of his or her grant money), department administrators, and fellow graduate students about available funding. Go to your university's fellowship office or its equivalent, and look through the listings in *The Annual Register of Grant Support*, *The Grant Register*, *The Chronicle of Higher Education*, and *Foundation Grants to Individuals*. Look into NSF grants (there are several different programs). Take advantage of your status as a woman or minority if you are one (this may be the only time when it actually is an advantage). Most universities have fellowship programs that may be administered through individual departments or may be campus-wide.

If you have not yet begun actively doing research, getting an RA position from a faculty member may be a good way to become involved in a research project. Working on an existing research project by maintaining or developing hardware or software, writing reports, and running experiments will give you a feel for what it is like to do research—and you may even find a thesis topic. Ask around to see what is available, and go talk to professors whose work you find interesting.

For a research grant or fellowship, you will probably have to write a proposal, so the more you have thought about potential thesis topics, the better off you will be. You may need to tailor your proposal to the interests and needs of the particular funding agency or program you are applying to, but stick to something you know about and are sincerely interested in.

Write for a general audience, since the people reviewing your application may not be in the same field. Emphasize your goals and why the project you propose to work on is important. Talk as much as you can about how you are going to solve the problem, and be sure that your proposed solution will satisfy the goals you have set forth. Follow the rules for format, page layout and length, or your application may not even be reviewed.

### References

1. Agre, P. E. 1982. What to read: A biased guide to AI literacy for the beginner. Tech. rep. working paper 239, MIT AI Lab.
2. Chapman, D. 1988. How to do research at the MIT AI lab. Tech. rep. AI working paper 316, MIT.
3. Kantrowitz, M. and Digennaro, J. P. 1994. *The Prentice Hall Guide to Scholarships and Fellowships for Math and Science Students*. Simon & Schuster.
4. Sloman, A. Notes on presenting theses. Available by anonymous ftp from <ftp.cs.bham.ac.uk>, in directory <pub/dist/poplog/teach>.

### Biography

Marie desJardins received her PhD in artificial intelligence from the University of California at Berkeley in 1992. She currently works with the Applied Artificial Intelligence Technology Program at SRI International, doing research in the areas of machine learning, planning, and intelligent tutoring systems. Dr. desJardins has taught numerous undergraduate courses, founded a student AI seminar series, and started the Big Sister program at Berkeley as president of Women in Computer Science and Engineering.

This article originally appeared in *Crossroads* 1.2 (Winter 1994), "Programming Languages."